Designing Sustainable and Trustworthy Markets: Models of Restaurant Operations and Platform Information Sharing

by

Yuwen Hu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in the University of Michigan
2026

Doctoral Committee:

Professor Ekaterina Astashkina, Co-Chair
Professor Izak Duenyas, Co-Chair
Professor Xiuli Chao
Professor Xun Wu

Yuwen Hu
yuwenhu@umich.edu
ORCID iD: 0009-0006-4724-0636

© Yuwen Hu 2026

ACKNOWLEDGEMENTS

This dissertation marks the end of a long, challenging, and unforgettable journey at the University of Michigan. Looking back, there were times when I felt lost, when research did not move forward, when I doubted myself, and when the future seemed unclear. But there were also moments of excitement, encouragement, and celebration that reminded me why this journey was meaningful. I would not have reached this point without the guidance, support, patience, and love of many people around me.

First and foremost, I would like to express my deepest gratitude to my advisors, Dr. Izak Duenyas and Dr. Ekaterina Astashkina. I feel truly fortunate to have had them as my advisors. Throughout my PhD journey, they have given me not only invaluable guidance in research, but also constant encouragement, trust, and care. They have always been there when I needed help—no matter where they were, whether they were traveling or in a different time zone. Whenever I faced a problem, they were willing to meet, think through it with me, and help me find a way forward.

I am especially grateful for their patience and for their belief in me. I still remember the difficult period when my first project on fashion rental needed substantial revision, and I was worried about whether I could pass candidacy. At that moment, instead of letting me stay trapped in fear and frustration, they encouraged me to move forward. Together, we started a new project on buffet operations and food waste, which later developed into one of the most promising works in my dissertation. Their confidence in me often came before my own confidence in myself, and that has meant more to me than I can fully express.

Beyond research, Izak and Kate have also shown me what it means to be thoughtful, caring, and grounded as people. I remember seeing Izak reading while dining in the dining hall, always making time to learn and think even in small moments of the day. I also remember him carefully looking into college courses for his daughter, with the same seriousness and care that he brings to research. These memories showed me a side of him beyond academia: a devoted father and a lifelong learner. I also remember Kate taking care of Mark with such warmth and patience. From her, I learned not only how to be a rigorous scholar, but also how to be gentle, responsible, and loving in everyday life. Their examples have shaped me in ways that go far beyond this dissertation.

I would also like to thank my cohort, Hyungchan and Renjun, and the senior students who have supported me throughout these years: Fan, Zijin, Joonho, and Clare. A PhD journey can sometimes feel lonely, but because of you, mine was filled with companionship and warmth. Hyungchan, thank you for all the advice you have given me, not only on finance and the stock market, but also on life, career, and the future. Renjun, thank you for showing me what courage looks like in daily life. Your "just do it" attitude—asking, trying, and moving forward even when there are doubts—has impressed me deeply. Fan, you have been my best friend at Ross. With you, I could talk, share, laugh, and feel sad, just like with close friends from high school. From life to study, you have never hidden anything from me and have always been willing to help. Your honesty, generosity, and companionship made these years much warmer. I also spent almost three years in the office with Zijin, Joonho, and Clare. You have always been the senior students who cared about me. You helped me when I faced obstacles in research, cheered me up when I felt down, and congratulated me at every milestone. When you moved out of the office, I felt a real emptiness. It was then that I realized how much your presence had been part of my everyday PhD life. Thank you all for making Ross feel less like a workplace and more like a community.

I am also deeply grateful to my friends Rex, Yusong, and Yuqing. You extended the radius of my life in Ann Arbor. Seeing you meant life outside research; it meant relaxation, good food, celebration, and the freedom to be myself without pretense.

I would also like to thank the mentors who have guided and encouraged me along the way. Stef, to me, you are a perfect example of a researcher. You are always efficient, always learning new things, always thinking, always positive, and always rigorous. You set a high standard in research, and I know I will continue learning from you throughout my life. Roman, thank you for encouraging me at every key step of my PhD journey. Your energy, humor, and positivity made even stressful moments feel lighter. I am also deeply grateful to Brian and Xiuli for their valuable advice on my research and the job market. Your guidance helped me broaden my view of how research can be connected to real-world problems and used to create meaningful impact beyond academic papers. Jane, thank you for your significant help during my job market. Your guidance and support were invaluable during one of the most important and difficult periods of my PhD life. I am also deeply grateful to Xinyu and Tony, who recommended me and helped send me onto the path of pursuing a PhD. Without your support at the beginning, this journey might never have started.

I am grateful to Ross and to the University of Michigan. Ross has been a big family to me. It is a great honor to be part of a university with so many leading scholars and outstanding schools across different fields. I have been fortunate to receive help and advice from many

people, including Mengzhenyu, Alys, Yao, Feng, and Zhuodong. Thank you for always being willing to share your experiences and offer valuable suggestions for my academic career.

I am also grateful to the staff members at the University of Michigan who have supported the university and made daily life here easier. The convenient bus system and great food on campus saved me a tremendous amount of time and money, especially as someone without a car. These forms of support may seem ordinary, but they made it possible for me to focus more fully on research and life in Ann Arbor.

Most importantly, I would like to thank my family for their continuing support. Your love has always been the foundation beneath everything I do.

To my wife, you are the one I owe the most. We were separated for five years, and I know this journey was difficult not only for me, but also for you. Every phone call with you—hearing your voice, seeing your face—made me feel relaxed and no longer lonely. The two months you spent with me in Ann Arbor were the happiest time of my entire PhD life. After you went back, I kept the habit of feeding the squirrels and birds; they also miss you. Thank you for loving me, waiting for me, and walking with me through these years, even from far away. I will always love you, and I promise that nothing will separate us again.

Table of Contents

ACKNOWLEDGEMENTS ii

LIST OF FIGURES vii

LIST OF TABLES ix

LIST OF APPENDICES x

ABSTRACT xi

CHAPTER

1 Introduction 1

2 Profitability and Food Waste in Buffet Restaurants 3

2.1 Introduction 3
2.2 Literature Review 7
2.2.1 Food Waste in Operations Management 7
2.2.2 Restaurant Operations and Buffet Design 9
2.3 Model Setup 12
2.4 Impact of Improved Convenience 15
2.4.1 Equilibrium, Per-diner, Expected Plate Waste 15
2.4.2 Equilibrium Expected Profit of the Buffet Operator 18
2.5 Impact of Waste-reducing Instruments 21
2.5.1 Penalty for Plate Leftovers 22
2.5.2 Consumer Education about Waste Consequences 25
2.5.3 Comparison of Instruments 27
2.6 Model Extensions 29
2.6.1 Partial Demand Learning 30
2.6.2 Insatiable Diners 31
2.6.3 Environmentally Conscious Diners 32
2.6.4 Impact of Plate Size Restriction 33
2.7 Directions for Future Research 35

3 Inventory Recommendations in Platform-Operated Fulfillment Service: Trust and Information Asymmetry 37

3.1 Introduction 37

3.2 Literature Review . . . 41
3.2.1 Platform Operations . . . 41
3.2.2 Fulfillment Service . . . 42
3.2.3 Information Sharing and Trust . . . 43
3.3 Model Setup . . . 44
3.3.1 Newsvendor Setup . . . 45
3.3.2 Signal-to-Noise Ratio (SNR) . . . 47
3.3.3 Trust Evolution and Static Recommendation . . . 48
3.4 Steady-state Analysis . . . 49
3.4.1 Static Recommendation . . . 49
3.4.2 Steady-state Trust Level . . . 50
3.4.3 Steady-state Profit . . . 53
3.5 The Case with Demand Censoring . . . 56
3.5.1 Bayesian-updating Sellers . . . 57
3.5.2 Steady-state Trust Level . . . 62
3.5.3 Platform's Optimal Strategy . . . 64
3.5.4 Comparison with the Baseline Optimal Profit . . . 65
3.6 Numerical Analysis . . . 67
3.6.1 Numerical Setup . . . 67
3.6.2 Supply Chain Coordination and the Value of Information . . . 68
3.6.3 Impact of Demand Censoring and Win–Win Scenarios . . . 72
3.6.4 Extension to General Trust Evolution . . . 74
3.6.5 Summary of Numerical Study . . . 78
3.7 Conclusion . . . 79

APPENDICES . . . 81

BIBLIOGRAPHY . . . 157

List of Figures

FIGURE

2.1 Impact of education on buffet's profitability. 27

3.1 Flow of events in cycle t 50
3.2 An example of sample paths of trust α_t and its cumulative average $\frac{1}{T}\sum_{t=1}^{T}\alpha_t$ for three values of DoS. 51
3.3 Platform and seller long-run average profits as a function of the platform's DoS 54
3.4 Probability that the seller increases trust, $\rho^c(\alpha,\eta)$, and the seller's long-run average trust, $\bar{\alpha}_c$, as functions of the DoS η. 61
3.5 Convergence of long-run average trust, the platform's profit, and the seller's profit in the baseline demand model. 68
3.6 Distribution of percentage gains from adopting the optimal recommendation for the platform, the seller, and the supply chain 69
3.7 Sensitivity analysis of (i) the percentage increase in supply-chain profit under the optimal recommendation (benefit of optimal recommendation) and (ii) the platform's share of that benefit, as functions of interest misalignment, the seasonal component, and the market-noise parameter 71
3.8 Average percentage increase in supply-chain profit, average percentage increase in the platform's profit, and the platform's share of the gain under different interest misalignment levels and signal-to-noise ratios (SNR) ($n = 100$). 71
3.9 Changes in the platform's and seller's long-run average profits under demand censoring, as functions of interest misalignment, the seasonal component, and the market-noise parameter 73
3.10 Probability that equilibrium long-run average profit increases under demand censoring. Resolution: $n = 100$ 74
3.11 Platform's long-run average profit as a function of the DoS under alternative seller trust behaviors (marker indicates the maximum). 76
3.12 The optimal DoS η^*, long-run average trust level $\bar{\alpha}$, and change in the platform and seller's long-run average profit when demand is censored $\Delta\bar{\Pi}$ with asymmetric learning rate. 77
3.13 The optimal DoS η^*, long-run average trust level $\bar{\alpha}$, and change in the platform and seller's long-run average profit when demand is censored $\Delta\bar{\Pi}$ with a conservative seller. 78

A.1 The optimal flat-rate penalty fee and the optimal education level across different combinations of self-service costs and unit production costs. 93

A.2 The marginal impact of consumer education on the buffet operator's expected profit across different values of post-education unit overage costs and unit production costs, for small and large values of the self-service costs. 95
A.3 Equilibrium expected profits before and after education across different combinations of post-education unit overage costs (o_e) and unit production costs (k) when the self-service costs are low, moderate, and high. 98
A.4 Per-diner expected plate waste under the optimal flat-rate penalty fee (ω_p^*) and the optimal education level (ω_e^*) across different combinations of self-service costs (ι) and unit production costs (k) when the unit overage costs (o) is small and large. 99
A.5 Comparison of per-diner expected plate waste under the optimal flat-rate penalty fee (ω_p^*) and the optimal education level (ω_e^*). 101
A.6 (When diners are environmentally conscious.) The marginal impact of consumer education on the buffet operator's expected profit for different combinations of values of this cost ($o_{G,e}$) and the values of the unit production cost (k), under small, moderate, and large self-service cost values. 113
A.7 (When diners are environmentally conscious.) Effectiveness of consumer education as an instrument and an example of waste comparison under the optimal flat-rate penalty ($\omega_{G,p}^*$) and the optimal post-education unit overage cost ($\omega_{G,e}^*$). 113

List of Tables

TABLE

3.1 Demand components and information access 46
3.2 Summary statistics of outcome variables (2,408 parameter combinations) 69
3.3 Benefits of the optimal recommendation for the platform, the seller, and the supply chain (baseline vs. censored demand) 70

List of Appendices

APPENDIX

A Proofs and Supplementary Materials for Chapter 2 81

Auxiliary Lemmas 81
Proofs for Results in Sections 2.4 and 2.5 83
Formal Analyses & Proofs for Model Extensions in Section 2.6 101
Auxiliary Results and Discussions 118

B Proofs in Chapter 3 127

Abstract

Modern market operations increasingly rely on sustainable practices and trustworthy platform-participant relationships to ensure long-term profitability and efficiency. This dissertation explores how thoughtful operational design and information sharing can align incentives, foster trust, and minimize inefficiencies in service and e-commerce markets.

The first chapter investigates the trade-off between profitability and food waste in all-you-can-eat buffet restaurants, addressing the critical challenge of post-consumer plate waste which accounts for over 70% of buffet waste. This research develops an analytical model to examine how operational levers—namely convenience, monetary penalties, and consumer education—impact both waste generation and firm profit. Preliminary findings show that enhancing convenience for plate refills reduces waste and boosts profit by helping diners better match their uncertain demands. Furthermore, carefully calibrated penalties consistently outperform education interventions in preserving profitability.

Chapter two takes a digital supply chain perspective, studying inventory recommendations within Platform-Operated Fulfillment (POF) services, where misaligned incentives and information asymmetry often hinder seller adoption. The research asks how platforms can optimally design recommendations to balance short-run profit extraction with long-run trust formation. The analysis reveals that sharing information through strategic recommendations yields a "win-win" outcome for both parties. Interestingly, under demand censoring (lack of visibility into unfulfilled demand during stockouts), an "informational shield" is created that systematically tilts trust formation in the platform's favor, enabling more aggressive recommendations and higher long-run platform profits in volatile markets.

In summary, this dissertation contributes to our understanding of operational mechanism design across distinct markets. By highlighting the roles of physical convenience and trust-building information sharing, this study offers practical policy recommendations for managers and platforms aiming to optimize economic and sustainable outcomes.

CHAPTER 1

Introduction

In today's highly competitive and resource-constrained markets, it is critical to leverage robust operations management and supply chain strategies to align incentives, manage uncertain consumer behavior, and foster trust. Both service operations and digital supply chains face significant inefficiencies due to misaligned objectives and information asymmetry, whether dealing with post-consumer waste in restaurants or optimal inventory control in e-commerce fulfillment. Through two essays, I study different manifestations of these challenges: optimizing food accessibility and waste-reduction interventions in all-you-can-eat buffets under uncertain consumer demand, and designing trust-building inventory recommendation policies for platform-operated fulfillment services facing asymmetric information. Ultimately, this dissertation contributes to the understanding of operational mechanism design across distinct markets, demonstrating how thoughtful system design and information sharing can optimize both economic and sustainable outcomes.

The first essay investigates the trade-off between profitability and food waste in institutional all-you-can-eat buffet operations. A major operational challenge in this context lies in accurately modeling diners whose demands are endogenous and uncertain, as physiological and psychological factors make it difficult for consumers to anticipate their true appetite before a meal. The core research question asks how improving convenience and deploying monetary or educational interventions impact diner plate waste and profitability. To answer these questions, I develop a stylized analytical model of a profit-maximizing buffet operator serving diners with uncertain demands, incorporating self-service cost as a proxy for convenience and framing the diner's consumption as a two-stage Newsvendor problem with reactive capacity via plate refills. The analysis reveals that greater convenience—such as easier plate refills—actually reduces plate waste and boosts profit by encouraging diners to take smaller initial servings to hedge against their uncertain appetites. Furthermore, the results show that optimal monetary penalties consistently outperform consumer education on profit, as education can unexpectedly hurt the operator's profitability when production

costs are low.

The second essay turns to the digital supply chain, studying inventory recommendations within Platform-Operated Fulfillment (POF) services, such as Fulfillment by Amazon (FBA). The fundamental challenge in this ecosystem is the lack of seller adoption of platform-generated recommendations, which stems from a classic conflict of interest over inventory risk and the platform's superior, private demand information. This makes seller trust a critical mediator for effective information sharing . The central research question asks how a platform should optimally design inventory recommendations to balance short-run profit extraction with long-run trust formation and seller adoption. Methodologically, I develop a repeated Newsvendor model in which an informed platform recommends order quantities to a seller who retains full control of ordering, while trust evolves endogenously based on the seller's perceived performance of past recommendations. Interestingly, I find that under demand censoring (when sellers cannot observe true demand during stockouts), an "informational shield" is created that masks the potential overstock costs of larger recommendations. This asymmetry systematically tilts trust formation in the platform's favor, allowing the platform to earn higher long-run profits under censored demand than under full transparency in volatile markets.

Collectively, these two essays underscore the importance of context-aware modeling and strategic policy design in overcoming operational misalignments, information asymmetry, and behavioral uncertainties. By highlighting the roles of physical convenience and trust-building information sharing, this research offers practical policy recommendations for managers and platform operators aiming to optimize performance. For real-world buffets, operators are advised to enhance layout convenience to minimize congestion, and to strategically implement monetary instruments rather than relying solely on awareness campaigns to achieve financial and environmental sustainability. For e-commerce platforms, the findings suggest that treating trust and transparency as operational design variables allows managers to strategically tailor recommendations, utilizing the informational shield to sustain adoption and capture a larger coordination surplus in uncertain market environments.

CHAPTER 2

Profitability and Food Waste in Buffet Restaurants

This chapter is based on joint work with Ekaterina Astashkina and Izak Duenyas.

2.1 Introduction

Restaurants and foodservice businesses generate 12.9M tons of surplus food yearly, more than 80% of which is sent to landfill or incinerated as waste, exacerbating climate change [108]. Buffet restaurants—defined as establishments where food is placed in a public area for diners to serve themselves—comprise 5,428 businesses in the United States as of 2022 and contribute to this statistic [65]. Among these are all-you-can-eat buffets—where diners pay a flat entrance fee for unlimited food—the most common buffet type, generating a sizable $8 billion in annual revenue [104]. These buffets span various industries, including hospitality, educational institutions, hospitals, retirement homes, large events, and street dining. Buffets have been labeled the culprit of food waste [56], where nearly half of the food prepared remains uneaten [78].

Buffet restaurants operate on thin profit margins—which makes efficient food waste management crucial for their survival and profitability. In line with the UN's Sustainable Development Goal #12 to halve global per capita food waste by 2030 [125], buffets must navigate the challenge of reducing waste while staying profitable. At a high level, we thus explore how all-you-can-eat buffet design variations impact diner plate waste and a buffet's profitability. We specifically focus on plate waste (i.e., post-consumer) because, in all-you-can-eat buffets, it accounts for more than a staggering 70% of total food waste [108]. This type of waste is also harder to tackle compared to kitchen waste (i.e., pre-consumer) and, unlike kitchen waste, it cannot be donated.

Buffet restaurants generally fall into two distinct categories. The first includes university and employee dining halls, where customers are repeat diners who value efficiency, face time constraints, and are familiar with the food offerings. These diners are typically not there

to explore. Think of a student grabbing lunch between back-to-back classes—on a tight schedule, mentally preoccupied, and focused on refueling and getting back to their day as opposed to savoring a drawn-out meal. The second category includes leisure buffets—such as those in resorts, cruises, and hotels—where diners are unfamiliar with the menu, seek variety, and view the meal as part of the overall experience. This paper focuses on the first category: institutional buffets where efficiency, not experimentation, drives diner behavior.

A key design factor for such buffets is the level of convenience for customers, especially the ease of access to food stations for plate refills [114]. Convenient designs feature short distances between seating and food stations with open layouts, reducing congestion and minimizing the time and effort needed for diners to get their food. In contrast, inconvenient designs make food less accessible, deterring frequent refills. Such setups often create significant congestion, especially during peak times, making it harder for diners to access the buffet and return to their seats [57]. Extremely inconvenient designs even feature barriers such as turnstiles that discourage diners from returning for more food (e.g., employee dining halls of Las Vegas Sands Corp. and numerous corporate cafeterias). Given that buffet designs vary in convenience, in the first part of the paper, we examine how improving convenience in all-you-can-eat institutional buffets impacts diner plate waste and a buffet's profitability.

To tackle plate waste, all-you-can-eat buffets have implemented and studied various interventions (see Section 2.2.2 for an overview); however, the effectiveness of some approaches remains under-explored. Two notable examples are: (i) monetary incentives, such as penalties for excess waste (common in some Asian buffet, e.g., [89]) or rewards for clean plates (used in institutional settings, e.g., [131]), and (ii) consumer education on the environmental and social impacts of food waste, widely promoted in university and corporate dining settings [28]. In the second part of the paper, we examine how these monetary and educational interventions influence plate waste and buffet profitability, and compare their effectiveness.

Model Framework

We develop a stylized analytical model of an all-you-can-eat buffet where a buffet operator sets an entrance fee to maximize expected profit. A key feature of the model is the self-service cost—the cost incurred by customers each time they fill their plates—which captures the level of convenience the buffet provides, with higher values reflecting less convenient buffet designs. Customers face varying outside option costs and uncertain food demand due to their unpredictable appetites. Each customer decides whether to join the buffet and how much food to take on their first plate. After eating, they may leave plate waste or, if plate is cleared, decide to refill their plate. Customers aim to satisfy their demand at the lowest expected cost. They incur a unit overage cost for generating plate waste and a unit underage

cost for falling short of satisfying their demand. The buffet operator incurs production costs based on the amount of food customers take, both consumed and left as plate waste.

We characterize the optimal decisions of the customers and the buffet operator that arise in equilibrium, and derive two key performance metrics: per-diner expected plate waste and the buffet operator's expected profit. To assess the role of convenience, we conduct sensitivity analyses on these metrics based on the self-service cost.

To assess the efficacy of the penalty intervention, we modify the base model to allow the buffet operator to impose a penalty fee whenever a diner generates plate waste. In this scenario, the operator sets both the entrance fee and the penalty fee. We compare the resulting equilibrium profit and plate waste to those from the base model, where no penalty is imposed. We evaluate the impact of rewards in a similar fashion.

To assess the efficacy of the education intervention, we adjust the base model by assuming that after education, each diner's unit overage cost increases. More intense education corresponds to a higher post-education unit overage cost. We compare the equilibrium profit and plate waste pre- and post- education intervention.

Findings, Contributions, and Practical Implications

Operational upside of convenience: Our first result reveals the operational upside of a more convenient buffet design. Specifically, greater convenience—such as easier refills—can reduce plate waste (Theorem 2.3) and increases profitability (Theorem 2.7). This happens because convenience helps diners navigate the inherent uncertainty of their appetites. By encouraging smaller initial portions, this convenience reduces plate waste and overall consumption, ultimately leading to lower production costs and enhanced profitability. The key takeaway for practitioners is that buffets that impede food access, whether intentionally or not, may miss out on these advantages.

Consumer waste as deadweight loss: Owing to the unique economics of buffets, a direct implication of our model is that consumer waste benefits no one. Customers pay a fixed price, so any uneaten food imposes a direct cost on both the operator and the diner—constituting a deadweight loss. Although this form of mutual loss is often acknowledged in practice, it has not been formalized in prior analytical models. This economic logic contrasts sharply with that of grocery retail settings, where firms profit from consumer over-purchasing—even when it leads to at-home waste—since every unit sold generates revenue.

Uncertainty can reduce waste: We show that increased demand uncertainty—typically viewed as a source of inefficiency—can actually reduce waste in buffet settings, where diners

adjust their intake after learning their true demand. This contrasts with most operations management contexts, where uncertainty tends to degrade performance. The mechanism arises from diners hedging against uncertainty by initially serving themselves smaller portions, thereby reducing overage.

Penalties & **rewards:** We show that imposing an optimal penalty for plate leftovers consistently reduces plate waste and increases the buffet operator's profit (Theorem 2.10). The largest waste reductions occur in buffets with high unit production costs, while the greatest profit gains arise in buffets with moderate production costs—where the trade-off between discouraging waste and maintaining customer demand is most favorable. Given the variability in unit production costs across buffet [32, 69], these findings offer relevant guidance for practitioners. We also show that other optimally designed rewards perform equally well.

Education: We further find that, like the penalty intervention, educating consumers consistently reduces per-diner plate waste. However, unlike penalties, it does not always increase the buffet operator's profit. Specifically, consumer education reduces profits if and only if the buffet's unit production costs are below a certain threshold (Theorem 2.13). The insight that education can reduce profits is especially novel and managerially relevant, urging buffet operators to rethink awareness campaigns unless paired with appropriate cost-sensitive strategies.

Comparison of interventions: Finally, we find that employing monetary instruments (penalties or rewards) consistently yields higher profits than educating consumers. However, optimal consumer education (i.e., at an intensity that maximizes profit), compared to imposing an optimal penalty or reward, can result in greater reduction in plate waste (Theorem 2.14). This occurs when the buffet operates under moderate unit production costs and diners' pre-education unit overage costs are low compared to their unit underage costs. The second condition refers to situations where, prior to education, diners prioritize avoiding shortages to a greater extent than avoiding waste, making waste a secondary issue—something likely to hold true in practice.

In sum, both monetary and educational instruments reduce plate waste, but only optimally calibrated penalties and rewards consistently improve profitability. To our knowledge, prior work has not compared these two interventions within a unified analytical framework.

Robustness

We analyzed several model variants, including diners with imperfect demand learning, insatiable appetites, or environmentally conscious preferences (Sections 2.6.1-2.6.3). We also examined scenarios where the amount of food a diner can take in each round is limited, and studied the implications of plate size restrictions (Section 2.6.4). We also extended the model to account for congestion feedback, where greater convenience attracts more diners and thus endogenously affects the self-service cost (Appendix A.4.3). Across all extensions, our core findings remain robust, with only minor deviations—and each variant offers additional insights.

Conceptual Contribution

Overall, prior research on restaurant plate waste has been largely empirical or qualitative, and has emerged from fields outside operations management (OM). While this body of work documents observable drivers of waste—such as portion sizes, menu variety, or behavioral nudges—it does not model the underlying operational design choices or cost structures, nor does it link these factors to profitability. As a result, existing literature lacks a formal theoretical account of how buffet system design affects consumer behavior, waste, and financial performance [109]. Our model addresses this gap. Additionally, our work provides a structured, operations-informed foundation for future analytical and empirical research within OM, as well as in adjacent fields such as sustainability, behavioral economics, public health, and marketing.

2.2 Literature Review

We contribute to the literature on (i) food waste in operations management and (ii) buffet design and restaurant operations.

2.2.1 Food Waste in Operations Management

Research on food waste in operations management is expanding, as highlighted by the OM-forum paper by [2], which calls for more investigation in this area.

Analytical papers on food waste cover a range of topics, including shipment policies in food retail [3], salesforce incentives [4], food quality standards [90, 36], inventory strategies for pre-made foods [100], product display design in grocery stores [9], retail promotions [133], variability propagation in perishable food supply chains [135], blockchain technology in food supply chains [75], among others.

Most studies in this literature emphasize food waste generated by retailers (i.e., pre-consumer waste). However, in developed countries—in both retail and restaurant contexts—consumer-generated food waste (i.e., post-consumer waste) is typically much larger [50]. Our paper focuses specifically on consumer-generated food waste, where the demand faced by operators is endogenous. This contrasts with retail-focused studies that treat retail demand as exogenous, simplifying profitability analyses.

A few papers, like ours, study consumer-generated food waste under endogenous demand, but focus on grocery retail and household waste (e.g., [12, 8, 77, 59]). In contrast, we analyze restaurant plate waste, which involves distinct behaviors: buffet diners can adjust consumption mid-meal, refill instantly, and face flat pricing instead of per-item costs. These features fundamentally shift waste dynamics. Our model incorporates them to reveal how convenience shapes waste, affects profitability, and influences the effectiveness of interventions—insights not addressed in prior work.

Among prior work, [11] is particularly relevant for linking convenience to food waste: higher store density lowers travel costs, leads to more frequent shopping, and reduces household waste. Our finding yields a similar waste-reducing effect of convenience, but in a fundamentally different setting with distinct consumer behavior, cost structures, and pricing. More importantly, their model does not address profitability—and if included, the conclusions would likely diverge from ours. In retail, lower consumer waste reduces overage volume of sales, which can hurt retailers. In contrast, buffets benefit: under flat pricing, less waste lowers production costs. These fundamental differences in pricing structures and incentive alignment limit the applicability of retail-based models and insights to buffet settings. To address this gap, we develop a newsvendor-style model with reactive capacity, where repeat self-service is shaped by buffet convenience.

Our model also yields novel insights on demand uncertainty. Prior work finds that greater uncertainty increases waste in standard inventory settings (e.g., [7]). [11] similarly links higher demand variability to more waste on the retail end, but models consumer demand as a Poisson process, where mean and variance are linked—precluding independent analysis of uncertainty. In contrast, we model individual demand using a uniform distribution, allowing direct variation in uncertainty. This reveals a surprising result: sometimes, greater uncertainty can reduce plate waste. Crucially, this effect hinges on buffet convenience, highlighting a nuanced and operationally important interplay between access friction, consumption variability, and the likelihood of plate refill—one not captured in prior work.

Finally, to our knowledge, only three papers in operations management directly address food waste in restaurants. The first, [136], analytically explores how platforms like TooGoodToGo help clear restaurants' leftover inventory, but it focuses on consumer waste at home

rather than consumer waste in a restaurant, making its questions and models fundamentally different from ours. The other two—[93] and [19]—focus on kitchen waste in restaurants and are empirical. In contrast, our work is analytical and focuses on plate waste, which, as noted earlier, presents a particularly important challenge in buffet settings.

2.2.2 Restaurant Operations and Buffet Design

Analytical papers on restaurant operations cover various topics, including self-ordering technology [47], restaurant revenue management [14], coordination contracts in food delivery [29], strategies to reduce shipment delays and improve delivery accuracy [124], among others. Our work stands apart by focusing specifically on all-you-can-eat buffets, a unique context that previous research on restaurant operations has not addressed.

Existing research on buffet design has explored various interventions to reduce plate waste, but remains entirely empirical or qualitative. To our knowledge, this is the first analytical model to examine how buffet design affects both plate waste and profitability. Below, we outline the literature on the specific designs and interventions analyzed in this paper.

Accessibility: The first part of our paper focuses on a key operational lever: the accessibility of food in buffet settings. While empirical research has shown that increased accessibility can reduce plate waste (e.g., [72]), our work advances this insight by unpacking the operational mechanisms that drive this effect and identifying when the relationship holds. Specifically, we find that the effectiveness of accessibility depends critically on the ease of replenishment. Greater accessibility prompts smaller initial servings only when diners (who face uncertain demands) can conveniently return for refills—that is, when the self-service cost is sufficiently low—ensuring a non-zero probability of plate refill. This qualification helps explain why some empirical studies report mixed or context-dependent results.

Some scholars question whether increased accessibility consistently reduces waste. For instance, [71] suggest that easier access may backfire by reducing the perceived value of food, potentially encouraging overconsumption (although they are not able to test this). Moreover, their study is rooted in hotel buffets—settings designed for leisure and indulgence, where abundance and variety-seeking are central to the dining experience. Such contexts are fundamentally different from institutional buffets, like those in universities or workplace cafeterias, where convenience, speed, and functionality dominate. In these environments, the marginal value of an extra trip to the buffet is not leisure but hassle, and waste is more likely driven by uncertainty in appetite rather than decadence or novelty-seeking. Our model is calibrated to this reality—the reality of institutional buffets. We show that when refills are convenient, diners are more likely to begin with a conservative serving and use the

subsequent round to better match their true demand—thereby reducing waste by improving information.

Further, while our model abstracts from behavioral complexities that are unlikely to arise in institutional buffet settings, it offers a robust framework for analyzing how buffet design interacts with diner behavior—particularly under uncertainty about individual appetite. We focus on this uncertainty in part because it is closely aligned with findings from the Theory of Planned Behavior (TPB) literature. [34], applying TPB to restaurant settings, identify perceived behavioral control (PBC) as the most influential driver of plate waste. In this context, PBC refers to how confident a diner feels in their ability to avoid waste—an ability that depends on how well they can anticipate their appetite and select an appropriate portion. These dimensions of perceived uncertainty are captured in our model through stochastic diner demand, and help explain the difficulty of making informed self-service decisions.

To date, existing studies on plate waste have not offered analytical models that isolate these mechanisms cleanly, often due to confounding behavioral or contextual factors that are inherently present in empirical research. As summarized by [109], the lack of theoretical grounding and methodological limitations in prior studies on plate waste reduction tools make it difficult to infer clear cause-and-effect relationships, and their results should be interpreted with caution. Our work aims to close this gap and can also serve as a foundation for future empirical studies seeking to test and build on our theoretical insights.

Waste-reducing instruments: In addition to exploring the relationship between convenience and plate waste, we investigate two classes of waste-reducing interventions: monetary (i.e., penalties and rewards) and non-monetary (i.e., consumer education about waste consequences). While previous research has shown that monetary fines can significantly reduce plate waste (e.g., [79] report a 54% drop in average waste with penalties), our work refines this understanding. We delve into the conditions that maximize the efficacy of penalties on both waste reduction and profitability, revealing that their impact varies based on the underlying food production costs.

Past empirical research on educational campaigns has shown mixed results (e.g., [43, 130]). Our analytical model offers an explanation: some prior interventions may have failed to increase diners' perceived cost of overage, resulting in little impact on waste. We further compare education to monetary instruments and show that while neither consistently outperforms the other in reducing waste, monetary instruments dominate in terms of profitability.

The analytical nature of our work enables rigorous comparison of intervention efficacy, yielding insights that are difficult to obtain through empirical methods alone (e.g., [27]). For example, we show that the relative effectiveness of penalties versus education hinges

on the consumer's overage and underage costs, as well as the buffet's production costs. Such insights into the mechanisms governing consumer behavior and restaurant operators' decisions—shaped by the cost structures in buffet settings—are often overlooked in existing research.

In addition to the interventions covered in our paper, several other strategies have been explored in past research regarding their effectiveness in reducing plate waste. Notable interventions include reducing plate size, which has been shown to be effective in multiple studies [72, 128, 55], and reducing portion sizes of food offered at buffet stations, which not only decreases plate waste but also reduces overall food intake [45]. Other strategies, such as removing trays from the buffet area [76, 6] and employing gamification approaches to incentivize less waste [41], have also shown significant results.

Additional interventions, such as offering takeaway boxes for leftovers [35], streamlining menus and reducing food variety [132, 103], food pre-tasting [99, 70], and improving food labeling [99, 70], have been suggested as potentially effective, but their efficacy has not been thoroughly investigated to date.

In summary, our analytical research on food waste reduction in buffet settings complements and extends prior empirical work by illuminating the underlying mechanisms of consumer behavior through an operational lens—offering insights that empirical approaches are not always equipped to uncover. In addition, our study fills a critical gap in the literature by examining how different buffet designs and waste-reducing interventions affect profitability—an aspect that, to the best of our knowledge, has not been explored in existing research.

The rest of the chapter is organized as follows. In Section 2.3, we set up our base model which captures the game between the profit-maximizing all-you-can-eat buffet operator and a unit mass of cost-minimizing customers. We also introduce the definition of individual plate waste, expected non-monetary costs, and other terms that will be used in the paper. In Section 2.4, we provide the equilibrium outcomes of the game defined in the base model. We also provide a sensitivity analysis of the self-service cost on the individual plate waste and the buffet's expected profit. The discussion of the impacts on the individual plate waste and the buffet's expected profit from different waste-reducing instruments is provided in Section 2.5. Finally, Section 2.6 extends the baseline model to include imperfect demand learning, insatiable appetites, and environmentally conscious preferences, as well as plate size restrictions.

2.3 Model Setup

We build a stylized analytical model of a buffet operator offering an all-you-can-eat dining experience to a unit mass of infinitesimally small customers. Customers are heterogeneous in their outside option cost c_o, which is uniformly distributed on $[0, \bar{c}_o]$ across the population. A defining feature of the buffet setting is that a flat entrance fee, paid upfront, grants customers unlimited access to the available food.

Customer Demand. We model each consumer's demand—that is, the amount they would ultimately want to eat in the buffet—as a random variable D, distributed uniformly on $[0, \bar{d}]$. (In Appendix A.4.2, we show that our results still hold if demand is instead distributed as $D_a \sim U[\underline{d}, \bar{d}]$ with $\underline{d} > 0$.) We assume that demands are independently and identically distributed across the population. This captures a key behavioral fact: food intake is inherently uncertain and varies across meals, even for the same individual. Research in nutrition science, physiology, psychology, and marketing identifies two main sources of this variability—internal and external—both of which our model accommodates [60, 18, 112, 33].

Internal uncertainty stems from physiological and psychological factors that shape appetite. Hunger fluctuates with energy needs, hormones, stress, sleep, time of day, and mood—drivers that shift daily and are hard to anticipate [48]. Externally, food intake is influenced by how the meal tastes once it begins. Factors like freshness, seasoning, and preparation can't be evaluated beforehand but directly affect how much consumers eat [112].

As a result, consumers rarely know their true demand upfront. They arrive with a vague sense of hunger and adjust as the meal unfolds, guided by evolving internal cues and real-time sensory feedback. To reflect this, we assume that the realization of their demand D is learned only after the first plate. In Section 2.6.1, we extend this to a partial learning setting, where consumers know if they're still hungry but not by how much. This mirrors evidence that satiety builds gradually and is difficult to assess even mid-meal [119].

Timeline and Decisions. The buffet operator (he) moves first by choosing an entrance fee φ to maximize his expected profit. Given the entrance fee, each customer (she) decides (i) whether to join the buffet and (ii) the initial amount of food to put on her plate, q, which we refer to as the size of the customer's first plate. As the customer consumes the food, she learns her demand realization d and consumes the amount equal to the minimum of her demand realization and the amount on her plate: $\min\{d, q\}$. She then makes her final (third) decision: (iii) to leave the buffet or stay. If the customer leaves, she either incurs the per-unit cost of underage u for every food unit she is short of from satisfying her demand

in full (this occurs if she took too little: $q < d$) or the per-unit cost of overage o for every unit on her plate she did not consume (this occurs if she took too much: $q > d$). If the customer stays, she refills her plate with the refill amount $(d-q)^+$, which fully satisfies her demand and leaves no extra food on her plate. Here, X^+ is equal to X when $X > 0$ and is 0 otherwise.

(In Appendix A.4.1, we show that if diners endogenously choose their refill amounts, cost-minimizing behavior leads them to take exactly $(d-q)^+$. Thus, assuming this behavior in the base model simplifies exposition without affecting results. Second, the two-round structure also reflects reality, as field evidence shows buffet diners typically take two to three servings [127]. Extending the model to more rounds would add complexity without changing core insights.)

To decide whether to join the buffet, the customer compares the costs of each alternative: the expected total cost of dining at the buffet g (explained in detail next) against her idiosyncratic cost of an outside option c_o (which subsumes the option of dining elsewhere). The customer chooses the alternative that corresponds to the lower cost. Her subsequent decisions are made to satisfy her demand at the lowest possible cost, as we present next. For brevity, we call any customer who chooses to join the buffet a diner.

Diner's Costs. The unit underage cost, u, subsumes negative emotional effects [88]. The unit overage cost, o, captures the disposal cost of plate leftovers and emotional disutility such as regret, shame, and guilt about not finishing all the food [116]. In practice, we expect that $o < u$, and hence we assume this in the base model. As we show in Section 2.6.3, our key results continue to hold even when customers exhibit the reverse cost structure, where $o \geq u$.

Every time a diner fills her plate, she incurs an inconvenience cost of self-service, denoted by ι (henceforth referred to as a self-service cost), which comprises the costs of traveling from the seating area to the food stations and back, waiting in the buffet queues, and the hassle associated with browsing, picking, and loading food onto the plate. Thus, the expected total cost that the diner incurs at the buffet is given by:

$$g(\varphi, q) := \varphi + \iota + \mathbb{E}_D[o \cdot (q - D)^+] + \mathbb{E}_D[\min\{\iota, u \cdot (D - q)^+\}]. \tag{2.1}$$

The first two terms of the above expression are the entrance fee φ that the diner pays upfront and the self-service cost ι that the diner incurs to fill her first plate. The third term is the expected overage cost $\mathbb{E}_D[o \cdot (q - D)^+]$, which describes situations when the diner does not finish food on her first plate. The last term $\mathbb{E}_D[\min\{\iota, u \cdot (D - q)^+\}]$ describes situations when the diner finishes her first plate. At this point, the diner has observed her actual demand, d, and must choose between two alternatives, selecting the one with the lower cost

(as indicated by the minimum operator). Specifically, she either leaves the buffet without satisfying her demand in full (and, thus, incurs the cost of underage $u \cdot (d-q)^+$) or stays in the buffet to get a plate refill with the amount that satisfies her demand in full (and, thus, incurs the self-service cost ι for the second time).

We decompose the expected total cost incurred by the diner at the buffet into two components: monetary and non-monetary:

$$g(\varphi, q) = \underbrace{\varphi}_{\text{monetary cost}} + \underbrace{\iota + \mathbb{E}_D[o \cdot (q-D)^+] + \mathbb{E}_D[\min\{\iota, u \cdot (D-q)^+\}]}_{\text{exp. non-monetary cost } \Delta(q)}, \tag{2.2}$$

where the expected non-monetary cost $\Delta(q)$ includes all costs a diner expects to incur inside the buffet, except for the payment. This decomposition will be a key reference throughout the paper.

Buffet Operator's Profit. The buffet operator collects φ from every diner and incurs a unit production cost k for every unit of food taken by diners. Let $\phi(\varphi, q)$ denote the fraction of customers who chose to join the buffet and τ denote the expected total amount of food that a diner takes in the buffet (both defined in Section 2.4). The operator's expected total profit is given by:

$$\pi(\varphi, q) = \phi(\varphi, q)\,(\varphi - k\tau(q))\,. \tag{2.3}$$

We assume that the buffet operator supplies food in quantities that exactly match diners' demand. While, in practice, buffets often prepare food in slight excess, this assumption serves as a reasonable approximation. Empirical evidence suggests that pre-consumer (kitchen) waste from overproduction is significantly smaller than post-consumer (plate) waste in buffet settings [108]. This simplification allows us to focus on the more pressing issue of post-consumer waste.

In our analyses, we focus only on cases where the buffet operator generates non-negative profit, as guaranteed by the following assumption.

Assumption 2.1. *The self-service cost is not too large: $\iota < \bar{\iota}$, and the unit production cost k is not too large: $k \leq \bar{k}(\iota)$; where $\bar{\iota}$ and $\bar{k}(\iota)$ are defined in Appendix A.2.1.*

The upper bound on the self-service cost ensures at least one customer finds the buffet dining attractive, while the upper bound on the unit production cost ensures the buffet remains profitable.

Plate Waste The expected individual plate waste ω generated by a diner at the buffet when taking q units of food on her first plate is defined as the food taken but not consumed:

$$\omega(q) = \mathbb{E}_D\left[(q - D)^+\right]. \tag{2.4}$$

Given the model's sequential structure, we characterize the equilibrium using standard backward induction (full derivations are in Appendix A.2.1). We now present the results and key insights.

2.4 Impact of Improved Convenience

We study the impact of improved convenience in all-you-can-eat buffets on two performance metrics of interest: the (i) per-diner expected plate waste (Section 2.4.1) and (ii) expected profit of the buffet operator (Section 2.4.2), that arise in equilibrium. The proofs of all results presented in this section are provided in Appendix A.2.1.

2.4.1 Equilibrium, Per-diner, Expected Plate Waste

We begin by analyzing a diner's optimal in-buffet behavior: the size of the first plate and whether she stays in the buffet for a plate refill, and how these choices change with improved convenience (Prop. 2.2). We then use these insights to draw conclusions about the impact of increased convenience on per-diner expected plate waste (Theorem 2.3), that arises in equilibrium. We refer to decisions made before demand is realized as ex-ante, and those made after as ex-post.

Ex-ante, the diner selects the first plate size $q^*(\varphi)$ to minimize expected total cost of in-buffet dining: $q^*(\varphi) = \arg\min_q g(\varphi, q)$, where g is defined in Equation 2.1. The decision whether to stay for a plate refill is made ex-post: the diner weighs the cost of underage (incurred if she leaves the buffet without a refill) against the self-service cost (incurred if she does refill) and selects the option with the lower cost. The next result summarizes the diner's optimal behavior in the buffet and how it is shaped by convenience.

Proposition 2.2.

(i) A diner's optimal first plate size is given by $q^*(\varphi) \equiv q^* = \min\{\iota, \zeta\}/o$. *A plate refill occurs if and only if the realization of the diner's demand is sufficiently large* ($d > \iota \cdot (\frac{1}{o} + \frac{1}{u})$).

(ii) As long as the self-service cost is not prohibitively high ($\iota < \zeta$), *a higher level of convenience leads to a smaller first plate size* ($dq^*/d\iota > 0$) *and a more likely plate refill* ($d\mathbb{P}(\textit{refill})/d\iota < 0$). *Otherwise, convenience has no effect on these measures.*

Here, $\zeta \equiv \bar{d} \cdot (\frac{1}{o} + \frac{1}{u})^{-1}$ *and* $\mathbb{P}(\textit{refill}) = 1 - \min\{\iota/\zeta, 1\}$.

The above proposition shows that the optimal first-plate size is independent of the entrance fee: $q^*(\varphi) = q^*$. This is because the fee is a sunk cost, simply adding to the diner's total expected cost. (Section 2.7 discusses how sunk cost fallacy might alter this.) Thus, the decisions in the proposition are both best responses and equilibrium outcomes.

Further, due to the option of a plate refill, which comes with a self-service cost ι, a diner may operate in one of two regimes ex-ante. The specific regime is determined by a unique threshold value ζ for the self-service cost ι, as specified in the proposition above.

Regime 1: If the self-service cost is prohibitively high (i.e., $\iota \geq \zeta$), the customer behaves as if the refill option does not exist—no level of underage cost justifies paying for an expensive refill. We refer to such values of ι as prohibitively high, since they effectively eliminate the refill option. In this regime, the customer is a standard newsvendor, choosing her first (and only) plate size to balance expected overage and underage costs, given by: $q^*|_{\iota \geq \zeta} = F^{-1}\left(\frac{u}{u+o}\right) = \bar{d} \cdot \frac{u}{u+o} \equiv \zeta/o$. ($F^{-1}(.)$ denotes the inverse CDF of D.) Naturally, here, the optimal plate size is independent of ι.

Regime 2: If plate refilling is not too costly (i.e., the self-service cost is below a threshold: $\iota < \zeta$), then, ex-ante, there is a non-zero probability of a refill. The diner anticipates this when choosing the size of her first plate. In this regime, the optimal first plate size balances the expected overage cost against the expected underage cost, capped by the self-service cost. The optimal solution is given by $q^*|_{\iota < \zeta} = \iota/o$. This first plate size increases in ι, reflecting the idea that in more convenient buffets (lower ι), diners, on average, serve themselves less initially. (Regime 2 resembles a newsvendor setting with reactive capacity in the form of plate refills. However, unlike standard reactive sourcing models—where each unit of additional capacity incurs a per-unit cost—our buffet model imposes a fixed cost for any refill, regardless of quantity. This makes the optimal first portion size distinct from the classical critical fractile formula.)

In sum, when refills are feasible, the diner takes less on the first plate compared to the standard newsvendor case, since the option to refill offers flexibility. The size of this reduction depends on how convenient the refill is. In equilibrium, the diner refills only if her realized demand is high enough: $d > \iota \cdot \left(\frac{1}{o} + \frac{1}{u}\right)$. This threshold increases with ι, meaning diners are more likely to refill when the buffet is more convenient. We now present our first main result.

Theorem 2.3. *As long as the self-service cost is not prohibitively high ($\iota < \zeta$), a higher level of convenience leads to a strictly lower equilibrium per-diner expected plate waste ($d\omega^*/d\iota > 0$). Otherwise, convenience has no effect ($d\omega^*/d\iota = 0$).*

In our base model, the diner learns her demand while consuming the first plate, so waste arises only if she doesn't finish it. Naturally, a larger first plate size q^* leads to higher expected plate waste ω^*, as more food increases the chance of leftovers. As shown in point (ii) of Prop. 2.2, greater convenience reduces the optimal first plate size (q^*)—and thus the expected plate waste per diner (ω^*)—unless the self-service cost ι is prohibitively high. In that case, neither the first plate size nor expected waste responds to changes in convenience.

Three additional insights emerge from our analyses of waste in equilibrium:

(i) The Disproportionate Impact of Plate Size on Waste. We find that a larger first plate size leads to disproportionately higher plate waste (as captured by the quadratic relationship between the first plate size and per-diner expected plate waste $\omega^* = (q^*)^2/(2\bar{d})$). This non-linear relationship arises because a larger plate size increases both the likelihood of generating plate waste and the amount of plate waste produced.

(ii) Waste Reduction Peaks at Moderate Convenience Levels. We find that it is specifically at moderate convenience levels that increasing convenience yields the most significant reduction in plate waste (see Appendix Lemma A.2). This is because, when convenience levels are too low (i.e. ι is at the prohibitively high levels: $\iota \geq \zeta$), convenience has no impact on plate waste (Theorem 2.3). Conversely, when convenience is too high, optimal first plate size is already small, so there isn't much change in plate waste with increased convenience.

(iii) Unexpected Benefits of Demand Uncertainty. We find that greater demand uncertainty can, somewhat counterintuitively, lead to lower plate waste under certain conditions:

Lemma 2.4. *(Informal) In sufficiently convenient buffets, greater demand uncertainty reduces the equilibrium per-diner expected plate waste.*

This finding stands in contrast to prior research, which typically associates higher demand uncertainty with increased overage in classical newsvendor models (e.g., [7, 102]) or links greater retail demand variability to higher food waste (e.g., [11]). In our setting, however, the opposite can occur. The reason lies in a buffet-specific feature: the ability to refill plates. When the buffet is reasonably convenient, diners facing more uncertain demand tend to serve themselves more cautiously on their first plate, relying on the option to refill if needed. This behavioral shift lowers expected plate waste. To isolate the effect of increased uncertainty while holding the mean fixed, we adjust the demand distribution's support to

begin at a strictly positive lower bound. The full details, along with a more formal statement and proof of the above lemma, appear in Appendix A.4.2 (see Lemma A.25). Notably, in less convenient buffets—where refills are more burdensome—greater uncertainty does lead to more waste, consistent with prior findings.

We have analyzed how increased convenience affects equilibrium per-diner expected plate waste. We now examine its impact on the buffet operator's equilibrium expected profit.

2.4.2 Equilibrium Expected Profit of the Buffet Operator

We first describe the remaining equilibrium decisions, including the customer's joining decision and the buffet operator's pricing decision. We then analyze the operator's equilibrium expected profit and determine how improvements in convenience influence this profit.

In equilibrium, any customer who faces a sufficiently high outside option cost joins the buffet (Appendix Lemma A.3) and behaves inside the buffet as outlined in Section 2.4.1.

The expected profit that the buffet operator generates under optimal customer in-buffet behavior $\pi(\varphi, q^*)$ is concave in an entrance fee φ (expression for $\pi(\varphi, q)$ is provided in Equation 2.3); therefore there exists a unique equilibrium profit-maximizing entrance fee $\varphi^*(q^*) \equiv \varphi^* = \arg\max_\varphi \pi(\varphi, q^*)$ (see Appendix Lemma A.3 for details). The optimal entrance fee is the result of a typical trade-off that the buffet operator faces: a higher entrance fee translates into a higher profit margin made on every customer who joins but attracts fewer customers. One might expect that a more convenient buffet would charge a higher entrance fee, but in some cases, lowering the fee with improved convenience is actually the optimal choice, as we establish in the next lemma.

Lemma 2.5. *The equilibrium entrance fee decreases with a higher level of convenience if and only if the unit production cost is large enough and the level of convenience is in the moderate range. Formally, $\forall \iota \in [\iota_\varphi, \zeta]$, $d\varphi^*(\iota)/d\iota > 0$ if and only if $k > \kappa_\varphi$, where ζ is defined in Prop. 2.2, and the expressions for ι_φ and κ_φ are given in Appendix A.2.1.*

The above lemma identifies cases where the optimal entrance fee decreases with improved convenience, rather than increasing. As discussed on page 17, at moderate convenience levels, greater convenience yields the largest reduction in per-diner expected plate waste, directly lowering production costs. When unit production costs are high, these savings per diner become especially substantial. In this case, each diner becomes less costly to serve, making it more profitable to attract more customers by lowering the entrance fee. The operator benefits more from a larger customer base paying a lower fee than from charging more and deterring diners, as savings from reduced waste outweigh lost revenue.

Although the equilibrium pricing isn't always monotonic in convenience (Lemma 2.5), the buffet operator's expected profit can be decomposed into two components, each changing monotonically with improved convenience. We formalize this next.

Proposition 2.6. *The equilibrium expected profit of a buffet operator can be expressed as follows:*

$$\pi^* = \frac{1}{4\bar{c}_o}(\bar{c}_o - \Delta^* - k\tau^*)^2. \tag{2.5}$$

Here, the equilibrium expected total amount of food that a diner takes in the buffet, τ^*:

$$\tau(q^*) \equiv \underbrace{\tau^*}_{\text{exp. amount taken}} = \underbrace{\mathbb{E}[D] - \overbrace{h^*}^{\text{exp. shortage}}}_{\text{exp. consumption}} + \underbrace{\omega^*}_{\text{exp. plate waste}} = \mathbb{E}[D] + \left(1 - \left(\frac{o}{u}\right)^2\right)\omega^*, \tag{2.6}$$

decreases in ι if $\iota < \zeta$ and, otherwise, is not impacted by it. Moreover, Δ^* decreases in ι. Here, $h^* = (o/u)^2\omega^*$ is the equilibrium per-diner expected shortage; ω^* is the equilibrium per-diner expected plate waste; and $\Delta^* \equiv \Delta(q^*)$ is the equilibrium per-diner expected non-monetary cost of in-buffet dining. The expressions for ω^*, $\Delta(q)$, and q^* are given on page 17, in Equation 2.2, and in Prop. 2.2, respectively.

The decomposition of the optimal profit function in Equation 2.5 reveals that, in equilibrium, the buffet's profitability is primarily driven by two factors: (i) the per-diner expected total amount of food taken, τ^*, and (ii) the per-diner expected non-monetary cost of in-buffet dining, Δ^*. Moreover, we establish that both of these factors decrease monotonically with the self-service cost. In the following, we elaborate on these two components and explain the reasoning behind their monotonicity properties.

(i): First, breaking down τ^* in Equation 2.6 reveals that the expected total amount of food a diner takes, τ^*, can be directly linked to per-diner expected plate waste ω^*.

In particular, the first equality in Equation 2.6 illustrates that the total amount of food a diner takes, τ^*, can be expressed as the sum of the expected amount the diner consumes and the expected plate waste, ω^*. The per-diner expected amount consumed can be further decomposed into the diner's expected demand, $\mathbb{E}[D]$, minus the expected shortage, h^* (i.e., the number of units by which the diner falls short of satisfying her demand). Intuitively, the per-diner expected total amount of food taken is equal to the diner's expected demand plus the expected mismatch, manifesting as either positive plate waste or negative shortage.

The second equality in Equation 2.6 arises from the fact that the diner's optimal first plate size balances the expected underage and overage costs. Consequently, the equilibrium levels of per-diner expected plate waste and shortage align according to the ratio $\frac{\omega^*}{h^*} = \left(\frac{u}{o}\right)^2$. This relationship between the per-diner expected total amount taken and plate waste, $\tau^* = \mathbb{E}[D] + \left(1 - \left(\frac{o}{u}\right)^2\right)\omega^*$, demonstrates that as long as the unit overage cost is lower than the

unit underage cost ($o < u$), an increase in per-diner expected plate waste corresponds to a larger expected total amount of food taken by the diner. Furthermore, based on Theorem 2.3, as long as the self-service cost is not prohibitively high ($\iota < \zeta$), the per-diner expected plate waste decreases with greater convenience, and, consequently, so does the per-diner expected total amount of food.

(ii): The per-diner expected non-monetary cost of in-buffet dining $\Delta^* \equiv \Delta(q^*)$ (defined in Equation 2.2) consists of three terms: the self-service cost ι, the expected overage cost $\mathbb{E}_D[o \cdot (q^* - D)^+]$, and the expected underage cost $\mathbb{E}_D[\min\{\iota, u \cdot (D - q^*)^+\}]$ capped from above by the self-service cost ι which captures the costly option of the plate refill. Δ^* is increasing in the self-service cost, which captures an intuitive observation: the less costly plate refill gets, the more the diner relies on it, which translates into lower expected demand mismatch costs that she incurs in equilibrium. Therefore, with more convenience offered in a buffet, it is easier for the buffet operator to attract customers. Our second main result now follows.

Theorem 2.7. *A higher level of convenience in a buffet results in a higher equilibrium expected profit for the buffet operator:* $d\pi^*/d\iota < 0$.

The result of Theorem 2.7 stems from Prop. 2.6. Specifically, increased convenience leads to diners taking less food τ^*, thereby reducing the per-diner expected production cost for the buffet operator: $k \cdot \tau^*$. Additionally, greater convenience makes it easier to attract customers, as they are willing to pay more for entry when the non-monetary cost of in-buffet dining Δ^* is lower. Together, these effects result in higher expected profit with increased convenience.

Theorem 2.7 provides a key practical implication: buffet operators aiming to boost profits should avoid designs that impose high costs on customers for accessing food or refilling plates. Reducing the distance between food stations and seating, eliminating turnstiles or barriers, and avoiding inconvenient utensils can help. Additionally, operators should encourage plate refills by making the process easier, redesigning the area to minimize queues at buffet stations, and optimizing the floor plan to reduce travel time.

Taken together, the results in this section suggest that improving buffet convenience can create a win-win: reducing plate waste (Theorem 2.3) while also increasing profitability (Theorem 2.7).

Consumer waste as deadweight loss A broader insight from our model is that, in buffet settings, consumer waste imposes real costs on both parties. Because diners pay a fixed price, any uneaten food creates a direct cost for the operator (via production expenses) and a cost for the diner (via the overage cost). This mutual loss delivers no offsetting

benefit—making it a form of deadweight loss. While this inefficiency is often acknowledged anecdotally (e.g., [31] reports that the average benefit-cost ratio for food waste reduction was nearly 7:1 over a three-year time frame based on the data from hotels in 15 countries), it has not been formally captured in prior analytical models. Notably, this stands in contrast to grocery retail settings, where firms benefit from consumer over-purchasing, as each unit sold generates revenue—even if it is ultimately wasted at the consumer's home.

Endogenous Convenience and Congestion Feedback Note that we treat the self-service cost ι as exogenous in the baseline model. In reality, however, convenience is endogenous. For example, when the buffet operator improves convenience (e.g., by adding food stations), this may attract more diners, leading to congestion that offsets the initial gains. Thus, the actual convenience experienced by diners is shaped jointly by operator choices and by how many people join. In our extended model (see Appendix Section A.4.3), we capture this feedback by letting the effective self-service cost be $\iota \equiv \iota_c \cdot \phi(\iota_c)$, where ι_c is the operator-controlled inconvenience level, and $\phi(\iota_c) \equiv \phi(\varphi(\iota_c), q(\iota_c))$ reflects the resulting diner participation rate which depends on ι_c. Although this extension adds analytical complexity, it does not alter our main insights about the role of convenience.

We now turn to other mechanisms aimed at reducing waste and improving buffet performance—specifically, penalty fees and consumer education.

2.5 Impact of Waste-reducing Instruments

In this section, we assess the effectiveness of two waste-reducing instruments: (i) imposing a monetary penalty for plate waste (Section 2.5.1) and (ii) educating diners on waste consequences (Section 2.5.2). Specifically, we modify our base model to gauge the individual impacts of each instrument on plate waste reduction and buffet profitability. We also compare the efficacy of these instruments from the waste reduction and profitability standpoints (Section 2.5.3). The proofs of all results presented in this section are provided in Appendix A.2.2.

For notational simplicity, we use subscripts p and e to denote the versions of the base model that include optimized penalty and education instruments, respectively (i.e., where the buffet operator selects the values of these instruments along with the entrance fee as the other decision component). The absence of a subscript indicates the base model without any specific instrument.

2.5.1 Penalty for Plate Leftovers

We let the buffet operator collect a penalty fee, denoted by v_p, every time a diner generates a non-zero amount of plate waste. The diner's expected total cost at a buffet that penalizes plate waste (when taking q_p units of food on her first plate), similar to the base model, comprises both monetary and non-monetary components:

$$g_p(\varphi_p, v_p, q_p) := \underbrace{\varphi_p + \overbrace{v_p \mathbb{P}(D < q_p)}^{\text{exp. penalty}}}_{\text{exp. monetary cost}} + \underbrace{\iota + \mathbb{E}_D[o \cdot (q_p - D)^+] + \mathbb{E}_D[\min\{\iota, u \cdot (D - q_p)^+\}]}_{\text{exp. non-monetary cost}}. \tag{2.7}$$

The diner's expected monetary cost of in-buffet dining includes the entrance fee φ_p and the expected penalty for generating leftovers, $v_p \mathbb{P}(D < q_p)$, where $\mathbb{P}(D < q_p)$ represents the likelihood of plate waste. As in the base model, the diner selects the size of her first plate to minimize the expected total cost: $q_p^*(\varphi_p, v_p) = \arg\min_{q_p} g_p(\varphi_p, v_p, q_p)$, and she joins the buffet only if her outside option is sufficiently high. Appendix Lemma A.4 characterizes the optimal first plate size and the equilibrium number of customers joining the buffet.

In this model variation, the buffet operator makes two decisions—on the size of the entrance fee φ_p and the size of the penalty fee v_p—to maximize the expected profit:

$$\pi_p(\varphi_p, v_p, q_p) = \phi_p(\varphi_p, v_p, q_p)(\underbrace{\varphi_p + v_p \mathbb{P}(D < q_p)}_{\text{exp. per-diner revenue}} - k\tau(q_p)). \tag{2.8}$$

Here, $\phi_p(\varphi_p, v_p, q_p)$ denotes the fraction of customers that join the buffet. In contrast to the base model, the per-diner expected revenue consists not only of the entrance fee but also the expected penalty amount, providing an additional source of revenue. The next lemma provides the characterization of the pricing decisions of the buffet operator in equilibrium.

Lemma 2.8. *In the buffet where customers are penalized for leaving plate leftovers, the optimal penalty fee v_p^* and entrance fee φ_p^* are given by:*

$$v_p^* = \arg\min_{v_p} \Delta(q_p^*(v_p)) + k \cdot \tau(q_p^*(v_p)), \tag{2.9}$$

$$\varphi_p^* = \frac{1}{2}(\bar{c}_o - \Delta(q_p^*(v_p^*)) + k \cdot \tau(q_p^*(v_p))) - v_p^* \mathbb{P}(D < q_p^*(v_p^*)), \tag{2.10}$$

where the expression for the optimal first plate size $q_p^*(v_p^*)$ under penalty is provided in Appendix A.2.2.

Lemma 2.8 is best understood by first assuming the buffet operator imposes a penalty fee without adjusting the entrance fee. This increases diners' expected total cost at a buffet (including, the non-monetary and monetary components), making the buffet less attractive for diners. To avoid penalty, diners take less food on their first plate thereby lowering their expected plate waste. Consequently, the buffet operator benefits from this reduction in plate

waste, by incurring lower per-diner expected production costs. As a result, compared to the no-penalty scenario, each customer becomes more valuable while the buffet becomes less appealing—leading to a warranted reduction in the entrance fee when a penalty is imposed.

This reduction isn't just to maintain the same per-diner revenue as in the no-penalty case. In fact, the no-penalty equilibrium entrance fee is actually higher than the combined total of the equilibrium post-penalty entrance fee and the expected penalties: $\varphi^* > \varphi_p^* + v_p^* \mathbb{P}(D < q_p^*(v_p^*))$. The buffet operator intentionally lowers a diner's expected total cost of dining in the buffet to attract more customers because their in-buffet behavior is corrected by the penalty, making them more desirable.

Finally, Equation 2.9 demonstrates how the optimal penalty is set. If set too high, diners' non-monetary costs rise excessively, and the savings in production costs aren't enough to offset the loss in attractiveness. If the penalty is set too low, it does not meaningfully reduce plate waste or production costs. The optimal penalty finds a sweet spot, lowering production costs while keeping diners' non-monetary costs manageable (as captured by the optimization problem in Equation 2.9).

Naturally, in equilibrium, a buffet that imposes a penalty fee for plate waste performs strictly better waste-wise compared to one that does not: $\omega_p^* < \omega^*$ (Appendix Lemma A.5). The next lemma provides insights into when imposing an optimal penalty results in a larger reduction in plate waste.

Lemma 2.9. *The (equilibrium) waste gap between a buffet that charges a penalty fee for plate leftovers and a buffet that does not, $\omega^* - \omega_p^*$, strictly increases in the unit production cost k.*

Lemma 2.9 states that, imposing a penalty fee on plate leftovers is associated with a bigger reduction in plate waste in buffets that operate under higher unit production costs. This is because a larger unit production cost necessitates the buffet operator to impose a higher penalty fee, which, in turn, "corrects" diner's in-buffet behavior more: diners take less on their first plate, which directly translates into lower per-diner expected plate waste.

We next compare profitability between a buffet that does not impose a penalty fee and one that does.

Theorem 2.10. *In equilibrium, a buffet that charges a penalty fee for plate leftovers is strictly more profitable than a buffet that does not impose a penalty fee: $\pi_p^* > \pi^*$.*

Theorem 2.10 states that, in equilibrium, a buffet that imposes a penalty fee for plate waste performs strictly better profit-wise compared to one that does not. The improvement

arises because two-part pricing (entrance and penalty fees) is more potent than a single pricing instrument (entrance fee alone). In particular, while the entrance fee only influences the number of customers that join the buffet, it does not impact their behavior inside the buffet. In contrast, the penalty affects diners' behavior inside the buffet (recall the discussion below Lemma 2.8): it lowers per-diner expected plate waste and, consequently, reduces the buffet's per-diner expected production costs. This reduction allows the buffet operator to decrease per-diner expected revenue, attracting more diners whose behavior is now corrected by the penalty, ultimately increasing the buffet's overall profitability.

Next, we establish when imposing penalties leads to a greater improvement in profits.

Lemma 2.11. *If the self-service cost is not prohibitively high ($\iota < \zeta$), the profit gap between a buffet that charges a penalty fee for plate leftovers and a buffet that does not, $\pi_p^* - \pi^*$, is non-monotonic in the unit production cost k: it first increases and then decreases in k. Formally, $\exists k_p^\pi > 0$ such that $d(\pi_p^* - \pi^*)/dk > 0$ if $k < k_p^\pi$ and $d(\pi_p^* - \pi^*)/dk < 0$ if $k > k_p^\pi$.*

The above lemma states that, as the unit production cost increases, the profit gap between a buffet that imposes a penalty fee and one that does not first increases and then decreases. Buffets with moderate unit production costs benefit the most from imposing a penalty fee, while those with very low or very high production costs benefit the least. At the low end, when unit production costs are small, penalties yield minimal improvement because reducing plate waste results in only small financial savings. In this case, the optimal penalty fee is low, and the two-part pricing closely mimics single pricing. At the high end, when unit production costs are large, buffets struggle to stay profitable regardless of the pricing scheme (with or without penalty fee). In this case, a buffet that penalizes customers sets the penalty fee high to achieve significant savings in per-meal production costs. However, high penalties reduce revenues as few customers are willing to risk penalties (diners barely create any waste, so not much revenues are collected on penalties), and entrance fees must be lowered to attract diners—all in all, leading to insufficient expected total revenues.

In sum, we find that imposing an optimal penalty fee leads to a win-win scenario by reducing plate waste and increasing profits. The greatest waste reduction from imposing penalties occurs in buffets with high unit production costs. However, it is only the buffets with moderate unit production costs that can achieve substantial profit gains from imposing penalties, as they can balance reducing waste and attracting customers effectively.

Penalties vs. Rewards In practice, buffet operators employ a variety of monetary instruments to reduce plate waste. Some impose a flat-rate penalty for leftovers [39], others use per-weight penalties based on the amount of uneaten food [107], and some instead offer

a monetary reward for returning a clean plate [131]. For expositional clarity, our main analysis focuses on the flat-rate penalty. However, our theoretical framework extends to other monetary instruments as well.

Lemma 2.12. *(Informal) When chosen optimally, a per-weight penalty, a flat-rate penalty, and a flat-rate reward yield equivalent outcomes: they result in the same reduction in per-diner expected plate waste and the same improvement in expected profit.*

The formal version of this result and its proof are provided in Appendix A.4.4, Lemma A.30. This finding is practically valuable: while penalties may generate resistance or dissatisfaction among diners—due to perceived unfairness or added stress—rewards tend to be more positively received. Our model does not capture such non-monetary behavioral effects, but, in practice, they may meaningfully influence customer experience and buffet reputation. The equivalence result, therefore, offers a powerful insight: operators can implement a reward-based system that is just as effective as penalties in reducing waste and improving profits, while potentially avoiding the negative sentiment often associated with punitive measures. For the formal definitions of these instruments and their analytical equivalence, see Appendix A.4.4.

2.5.2 Consumer Education about Waste Consequences

In this section, we study the impact of educating consumers about the consequences of waste on buffet profits and plate waste. In our model, educating a consumer means increasing their unit overage cost. Specifically, in the absence of education, each consumer has a unit overage cost o. Once the education instrument is applied, every consumer's unit overage cost increases to a higher value o_e. In practice, education varies in form and impact. Buffets may use posters, infographics, or digital displays, while more intensive approaches include videos with quizzes, workshops, or campus lectures on sustainability. A simple banner is likely less effective than a full educational session—with banner corresponding to a lower value of o_e in our model.

We assume that educating consumers comes at no cost to the buffet operator. (This assumption does not affect the results in this section or in Section 2.5.3, where we compare performance of different interventions. This is because, as we will show, education is strictly inferior to monetary instruments in terms of profitability—hence, accounting for its cost only strengthens this conclusion.) We also assume that the unit overage cost post-education is bounded from above by the unit underage cost: $o_e \leq u$, to capture the fact that there is a limit to the effectiveness of consumer education.

Naturally, in our model, any amount of education applied by the buffet will necessarily reduce per-diner expected plate waste: $\forall o_e > o$, $\omega^*(o_e) < \omega^*(o) \equiv \omega^*$ (Appendix Lemma A.6). This occurs because post-education, the diner's unit overage cost rises, leading them to take less food on their first plate, thus lowering expected plate waste. However, while the education instrument is effective in reducing plate waste, it is not always profitable, as demonstrated in our next result.

Theorem 2.13. *Consumer education o_e strictly reduces a buffet's expected profit if and only if the buffet's unit production cost k is sufficiently low: $\forall\ o_e > o$, $\pi^*(o_e) < \pi^*(o) \equiv \pi^*$ if and only if $k < k_e(o_e)$.*

The above result states that, surprisingly, educating consumers lowers a buffet's profit whenever a buffet operates under sufficiently low unit production costs. That is, buffets with such cost profiles, are better off by not educating consumers. This is because educating consumers—while reducing plate waste and lowering buffet's expenses—raises customers' expected non-monetary costs of in-buffet dining. Because diners now incur higher costs when they create plate waste, post-education, the buffet operator must lower the entrance fee to attract customers, compared to pre-education. And, under low enough unit production costs, $k \leq k_e(o_e)$, reduction in entrance fees leads to a greater revenue loss than the cost savings achieved from reduced plate waste.

A key practical insight from this analysis is that raising consumer awareness about the consequences of plate waste can result in a win-lose scenario—reducing both plate waste and profits.

In Appendix Lemma A.8, we further examine the monotonicity properties of the threshold value in unit production cost $k_e(o_e)$ (defined in Theorem 2.13) as a function of the education provided at the buffet (specifically the magnitude of the unit overage cost, post-education, o_e). Our findings are illustrated in Fig. 2.1, where the shaded region in the (o_e, k) coordinates indicates areas where education reduces profit.

We find that only buffets with high unit production costs (i.e., $k > u/2$) can educate consumers without reservations, as any amount of education in these buffets necessarily increases profits. This is indicated by the absence of shaded points, for $k > u/2$, in Fig. 2.1. Buffets with low unit production costs (i.e., $k < \underline{k}_e$) are better off avoiding consumer education, as any effort in this direction detracts from profits. This is shown by all shaded points, for $k < \underline{k}_e$, in Fig. 2.1. Lastly, buffets with moderate unit production costs (i.e., $k \in [\underline{k}_e, u/2]$) benefit from education only if the intensity is neither too high nor too low. This is illustrated in Fig. 2.1 by both shaded and unshaded areas, with shading occurring at extreme values of o_e, for $k \in [\underline{k}_e, u/2]$. However, note that determining the precise amount

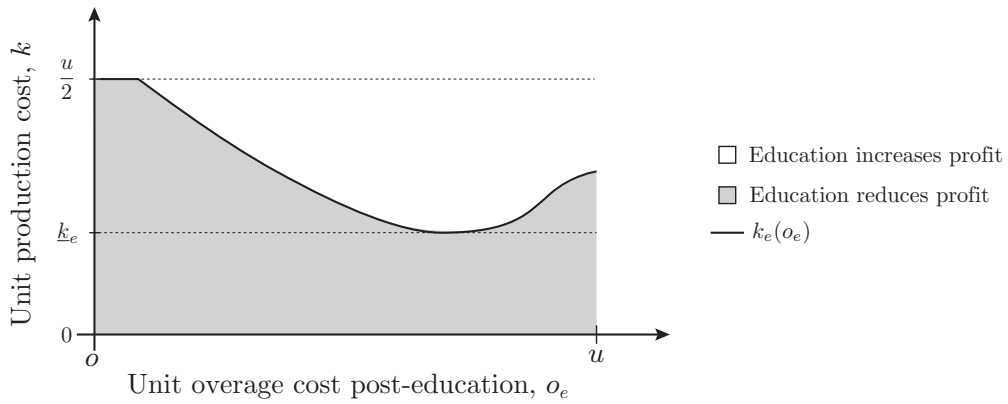

Figure 2.1: Impact of education on buffet's profitability.

Note: This figure shows that, in buffets with (i) high enough unit production costs, any amount of education increases profit; (ii) moderate unit production costs, too much or too little education reduces profits; and (iii) low enough unit production costs, any amount of education reduces profit. See Appendix A.2.2 for details.

of education needed (i.e., the exact value or range of unit overage cost for the consumer) is likely to be challenging in practice.

2.5.3 Comparison of Instruments

In this section, we compare the effectiveness of a monetary instrument (penalty) with a non-monetary instrument (consumer education).

As in Section 2.5.1, where we discuss the efficacy of penalty fees, we assume the buffet operator sets the penalty optimally. Specifically, the buffet operator chooses the optimal penalty, v_p^*, alongside the optimal entrance fee, φ_p^*, to maximize expected profit, yielding the equilibrium per-diner expected plate waste, ω_p^*, and buffet operator's expected profit, π_p^*, as characterized earlier.

Unlike Section 2.5.2, which explores the impact of consumer education at varying intensity levels, we assume here that the buffet operator can also choose the optimal level of consumer education. While precisely controlling education intensity (i.e., the exact post-education unit overage cost, o_e) may be challenging in practice, we make this assumption to enable a fair comparison between instruments. (Optimizing only one instrument would unfairly favor that instrument, which we seek to avoid.) Thus, in this section, the buffet operator selects the optimal post-education unit overage cost, o_e^*, and corresponding entrance fee, φ_e^*, to maximize profit. This leads to equilibrium per-diner expected plate waste, $\omega_e^* = \omega^*(o_e^*)$, and the buffet operator's expected profit, $\pi_e^* = \pi^*(o_e^*)$. Full characterization and results for the model under optimal education levels are provided in Appendix A.2.2. The resulting comparison of the effectiveness of instruments is given next.

Theorem 2.14 (Comparing the Performance of Instruments Set Optimally)**.** *Imposing a penalty for leftovers in a buffet always delivers higher expected profit than educating consumers. However, education is more effective in reducing per-diner expected plate waste if and only if buffet unit production costs are moderate and pre-education unit overage cost is low enough compared to the unit underage cost. Formally:* $\pi_p^* > \pi_e^*$*; however,* $\omega_p^* > \omega_e^*$ *if and only if* $k \in (\underline{k}^{\omega}, \bar{k}^{\omega})$ *and* $o < u/2$*.*

The above result is derived from Theorem 2.10 and Theorem 2.13, which separately analyze the effectiveness of optimal penalty and education interventions against a benchmark buffet with no such measures, as well as Appendix Lemma A.7, which characterizes and evaluates the performance of the optimal education level, o_e^*, that maximizes buffet operator's profit.

The first point that the above theorem establishes is that, from a profitability standpoint, the (monetary) penalty instrument consistently outperforms the (non-monetary) education instrument. This is because the components of the education instrument—the entrance fee and the post-education unit overage cost—serve distinct functions. Specifically, the entrance fee affects revenue collection, while the increased unit overage cost reduces production costs by lowering plate waste. In contrast, the components of the penalty instrument—the entrance fee and the penalty fee—cross-function. While the entrance fee still only affects revenue collection, the penalty fee not only reduces production costs but also acts as an additional revenue source. This dual role of the penalty fee enhances the penalty instrument's effectiveness in balancing production cost reduction and revenue generation, leading to higher profitability.

The second point established by the above theorem is that, from a waste reduction perspective, the education instrument outperforms the penalty instrument if and only if the buffet operates under moderate unit production costs and the pre-education unit overage cost is relatively low compared to the unit underage cost ($o < u/2$). Here is why.

When the unit overage cost is not sufficiently low ($o \geq u/2$), diners view waste and shortages as nearly equal concerns, leaving little room for education to make a noticeable impact. In this case, penalties are more effective at reducing waste. Conversely, when $o < u/2$, diners prioritize avoiding shortages over reducing waste, making waste a secondary concern. Here, education can significantly shift their focus and effectively reduce plate waste. The relative unit production costs then determine which instrument performs better from a waste reduction perspective. Specifically, for low production costs, any amount of education strictly reduces the buffet's profit (as established in Theorem 2.13), making it optimal to forgo education. This results in worse waste outcomes compared to the positive waste reduction achieved with a penalty instrument (Lemma 2.9). For high production

costs, both instruments are highly effective due to significant savings from waste reduction. However, the penalty instrument ultimately outperforms the education instrument because the latter reaches a saturation point (as is also supported by empirical evidence—see, e.g., [22]).

To sum up, the penalty instrument has an added financial benefit by generating extra revenue, which education instrument cannot replicate. This inherent revenue-generating advantage gives monetary instruments like penalties an edge over non-monetary instruments like education in terms of profitability. However, when it comes to reducing diner plate waste, neither instrument consistently outperforms the other. Their effectiveness hinges on the buffet's cost structure and how closely the pre-education unit overage cost aligns with the unit underage cost.

On Sub-optimal Education: In Appendix A.2.2, we compare the efficacy of penalty and education instruments for waste reduction without assuming that the education level is set optimally (i.e., we analyze the case where $o_e \in (o, u)$ is set exogenously). We focus on combinations of education levels and unit production costs that ensure profit post-education is no lower than pre-education (i.e., the white region in Fig. 2.1). Interestingly, we find that increasing education intensity (i.e., larger o_e) does not always broaden the range of unit production costs where education is more effective in reducing plate waste (Appendix Lemma A.8). The reason for this is that while higher education intensity leads to greater waste reduction, it also deters customers from joining the buffet by increasing their non-monetary dining costs. For the buffet to remain more profitable with intense education than without, the savings from reduced plate waste must be significant—something that only occurs when the cost of producing each unit of food is high. This restricts the set of unit production costs where the buffet makes more money post-education, thereby, potentially, shrinking the range of unit production costs where educating consumers is more effective than imposing an optimal penalty.

2.6 Model Extensions

We analyze several model variants, including (i) partial demand learning, (ii) insatiable consumers, (iii) environmentally conscious consumers, and (iv) plate size restrictions. This section summarizes each extension, confirms that our main findings largely hold with minor deviations, and highlights additional insights. The formal model setup, theorem statements, and proofs for each of these extensions are provided in Appendix A.3.

2.6.1 Partial Demand Learning

In our base model, we assumed that diners learn their demand perfectly with their first plate. Realistically, pinpointing demand quickly and with certainty is often infeasible. In this extension, we relax the assumption of perfect learning and instead assume that each plate that a diner takes allows them to only partially learn their demand. Specifically, in this setup, if a diner takes more food than needed, she immediately realizes it—knowing she is too full to finish her plate. However, if the amount is insufficient, she remains uncertain about the exact amount needed until she leaves the buffet. Appendix A.3.1 contains a formal set-up for this model extension.

Recall that, in the base model, diners perfectly observed their demand after the first consumption round, and if they opted for a second plate, they fully satisfied their demand. In this extension, however, diners may still have unsatisfied demand even after consuming their second plate, implying they might want a third plate if not restricted. However, to maintain tractability and consistency with the base model, we exogenously limit the number of plates to a maximum of two in this extension. (Allowing more rounds would add complexity but likely only strengthen our results: as self-service cost declines, diners would take smaller portions and thereby reduce waste.)

The formal analysis of the impact of improved convenience and the penalty instrument on buffet performance (along with the corresponding theorem statements for this model extension) is detailed in Appendix A.3.1. In what follows, we discuss these results, reporting any minor adjustments needed compared to the results obtained in the base model. Where appropriate, we clarify why such modifications remain largely inconsequential and do not alter the core insights of the paper.

Under partial demand learning, the results of Theorem 2.3 still hold: improved convenience reduces per-diner expected plate waste if the self-service cost is not prohibitively high (i.e., $\iota < \zeta_d$); otherwise ($\iota \geq \zeta_d$), it has no effect. In this setting, the convenience threshold ζ_d—which separates regime 1 (no refills) from regime 2 (occasional refills)—is lower than in the base model (i.e., $\zeta_d < \zeta$). Thus, with partial demand learning, the range of convenience levels that reduce waste narrows, and the range where convenience has no effect widens. This is because refills, while still useful, only partially reduce demand mismatch rather than eliminating it.

Theorem 2.7 remains valid under partial demand learning, but only if the diner's unit overage cost is not too high (i.e., $o < o_d < u$). This condition is stricter than in the base model ($o < u$), but we believe it still holds in practice.[1] Under this assumption, improved

[1]Fortunately, the threshold value for the unit overage cost is given by $o_d = (\sqrt{2} - 1)u > 0.4u$, making

convenience continues to raise the buffet operator's expected profit. We require this tighter range for the unit overage cost because, without it, partial demand learning could lead diners to take more food overall in expectation when refills are not prohibitively expensive (i.e., if $\iota < \zeta_d$) compared to when they are (i.e., if $\iota \geq \zeta_d$). Finally, our main result on the effectiveness of the penalty instrument (Theorem 2.10) also continues to hold.

Overall, adapting the base model to include partial demand learning provides a more realistic representation of consumption patterns and allows us to assess whether our core findings remain robust under less ideal conditions. As we demonstrated above, our results on the impact of improved convenience and the efficacy of the penalty instrument continue to hold with some minor modifications. This extension importantly underscores that the assumption of consumers perfectly learning their demand in the first round is not the primary driver behind our key results.

2.6.2 Insatiable Diners

In our base model, we assume that all diners have the same demand distribution. However, empirical evidence suggests that some diners have significantly higher appetites and such diners are often criticized for destroying the buffet's profitability [37]. This extension, thus, incorporates the presence of "insatiable diners," who consume more than regular diners.

Specifically, the unit mass of consumers is now divided into a fraction m of insatiable consumers and $(1 - m)$ of regular consumers, with the latter group remaining as in the base model. Insatiable consumers have a deterministic demand of $\beta\bar{d}$, where β is an integer greater than 1, and incur an outside option cost of $\beta\bar{c}_o$. Each consumer can take up to $\bar{d}$ units of food per round and there is no limit on the number of refills they can get. A complete formal model setup is provided in Appendix A.3.2.

Insatiable diners experience no uncertainty in their demand, hence they don't generate any plate waste; they leave the buffet only once they had satisfied their demand in full. We find that, as long as the proportion of insatiable diners is sufficiently low ($m < \hat{m}$), the buffet remains profitable. We discover that regular diners end up subsidizing insatiable diners, paying a premium over what they would in the base model.

It is worth emphasizing that this extension does not attempt to model adverse selection. While there is anecdotal evidence that particularly hungry individuals self-select into buffets, our goal is not to model that selection mechanism. Rather, we take the consumer mix as

the assumption $o < o_d$ relatively lenient. Specifically, while accurately measuring unit overage and underage costs can be challenging in practice, with o_d exceeding $0.4u$, we expect the assumption $o < o_d$ to be practically reasonable, as unit underage costs in buffet dining are likely to be significantly higher than unit overage costs [70, 122].

given and study its implications for waste, pricing, and interventions. That said, our results are consistent with such selection: hungrier consumers are more likely to find the buffet attractive and enter (all else equal), and some buffets may indeed end up serving that segment disproportionately. We view that as a market phenomenon—not a problem we aim to solve.

All findings in Theorems 2.3-2.13 remain intact, with formal result statements and proofs provided in Appendix A.3.2. The only additional condition required for the results related to waste-reducing instruments to hold (Theorems 2.10-2.13) is that the proportion of insatiable consumers must be small enough, as these consumers, within the scope of this model extension, are unaffected by waste-reducing interventions. Thus, the instruments are only effective if the fraction of insatiable consumers (who are immune to their impact) remains low. Fortunately, this requirement is not overly restrictive, as one of the most conservative estimates reports that only about 1 in 20 buffet diners consume significantly more than the average customer [37], with the actual proportion likely being even lower.

Interestingly, when the penalty is employed, insatiable diners actually derive even more value from the buffet. Since they aren't directly impacted by this interventions, they benefit indirectly from the reduced entrance fee. As a result, insatiable diners continue to enjoy the same amount of food, but now at a discount—essentially eating lots for even less.

2.6.3 Environmentally Conscious Diners

In our base model, diners prefer unmet hunger over generating plate waste. In practice, however, some "environmentally conscious" diners may prioritize waste avoidance, exhibiting waste-aversion. This extension explores how such preferences impact buffet operations, the role of convenience, and the effectiveness of interventions.

We adjust the model to represent diners who value waste avoidance more than fully satisfying their demand—that is, their unit overage cost is at least as high as the unit underage cost ($o \geq u$). To assess the impact of this preference shift, we assume all diners are environmentally conscious and examine whether this extreme case alters our results.

For environmentally conscious diners, the relationship between convenience and plate waste is preserved (i.e., Theorem 2.3 continues to hold). However, the relationship between the expected total amount of food taken and expected plate waste reverses: with less plate waste, diners end up taking more food, contrary to the base model, where less waste was associated with taking less (recall Equation 2.6). This reversal highlights a key discrepancy between the extension and the base model.

Therefore, when diners are environmentally conscious, they tend to take more food as

convenience increases. If production costs are high, this can reduce profitability—contrasting with Theorem 2.7, which shows that convenience generally raises profits in the base model. Notably, with environmentally conscious consumers, we show that convenience harms profit only when both inconvenience is extreme and production costs are high. Formal analysis and proofs are provided in Appendix A.3.3. It is surprising that increased convenience can hurt profits only when diners are environmentally conscious. One might expect that regular diners would be the ones to impact profitability negatively if the buffet operator were to offer more convenience. Yet, it is the environmentally conscious diners who pose the risk. As discussed above, greater convenience leads them to take more food overall—raising production costs and potentially reducing profitability.

Environmentally conscious diners still respond to penalties and educational efforts to reduce waste. So, our main results related to the effectiveness of the waste-reducing instruments (Theorems 2.10-2.14) continue to hold (see Appendix A.3.3 for proofs and formal statements).

In sum, the only result that might be at risk in this model extension is Theorem 2.7. Fortunately, we don't consider this a major issue. This is because this extension represents an extreme scenario where all consumers are highly environmentally conscious, whereas, in practice, most diners are likely not. Thus, while a few environmentally conscious consumers might use convenience in ways that could reduce profitability, their impact is unlikely to be significant unless they constitute the majority of the customer base.

2.6.4 Impact of Plate Size Restriction

In the baseline model, diners can take any amount of food per round, reflecting settings without physical constraints. However, in real-world buffets, plate size often limits how much food can be taken at once—and these sizes vary across venues. This extension factors this in.

We introduce an upper bound M on the amount of food a diner can take in any single round (either first plate or refill), and refer to M as the plate size. A diner is allowed one plate per round, and we assume M is large enough to allow any demand realization to be fulfilled in two rounds (i.e., $2M \geq \bar{d}$), preserving the two-round structure of the baseline model. Full derivations for the results of this extension are provided in Appendix A.3.4.

This extension confirms that our core findings related to per-diner plate waste, buffet profitability, and the role of refill convenience are preserved under this model extension (Theorems 2.3 and 2.7). With the plate size restriction in place, a more convenient buffet still can reduce plate waste and increases profitability. In addition, because plate size is a

lever that buffet operators actively experiment with, we further examine how such restrictions influence buffet outcomes. Below, we highlight the most interesting insights from this investigation.

Buffets with no refills ($\iota \geq \zeta$) We show that reducing plate size may result in lower plate waste and increased profits. This occurs because the plate size reduction prevents diners from taking their ideal quantity of food, effectively limiting overage and thus reducing expected waste. When unit production costs are sufficiently high, this reduction in waste translates into substantial cost savings and, overall, higher profits for the buffet operator. However, under lower production costs, interestingly, reducing plate size too much may backfire and result in lower profits. The reason is that cost savings from reduced food preparation are minimal, while the impact on diners is substantial: smaller plates force them to take well below their desired amount, sharply increasing expected underage. With refills being prohibitively expensive, this leaves diners unsatisfied, compelling the buffet to lower its entrance fee significantly. As a result, profits fall—even as waste decreases.

Buffets with refills ($\iota < \zeta$) When refills are possible, reducing plate size can have the unintended effect of increasing both plate waste and reducing profits. This counterintuitive outcome stems from how diners manage uncertainty about their appetite. In an ideal setting with no plate size restrictions, diners would take a small amount initially, learn their true demand, and then refill accordingly—akin to a tasting round before committing. However, with a binding plate size, this strategy becomes risky. For example, if each plate can hold 600 calories and a diner takes only 100 calories initially, they are limited to a total of 700 calories across two rounds—even if they later realize they want 1,000. To avoid falling short, diners hedge by taking more upfront, which leads to over-serving in the first round, increased plate waste, and detrimental impact on profits.

At a first glance, one might think that this effect is driven by our assumption that diners can take at most two rounds—a reasonable approximation for institutional settings, such as workplace or school cafeterias—where time constraints discourage multiple refills. However, we expect for the effect to persist even when more rounds are allowed, as long as some uncertainty about appetite remains before the penultimate round. That is, if diners are still unsure of their true demand in round 2 of a 3-round set-up, the plate size constraint may still prompt precautionary over-serving. The logic remains: with limited refill capacity, diners worry they won't be able to make up for under-serving later, and thus take more food than necessary earlier—leading to increased waste.

The highlighted dynamics underscore that reducing plate size is not universally beneficial:

while it may lower waste in some contexts, it can backfire in highly convenient buffets by discouraging refills, increasing first-plate portions, and ultimately raising food waste and hurting profits.

Plate size as a decision Our analysis also shows that, depending on the level of convenience, there may exist an optimal plate size that simultaneously minimizes waste and maximizes profit. However, such alignment is not guaranteed. In some cases, profit-maximizing plate sizes result in inefficiently high waste, and conversely, waste-minimizing choices may hurt profits. The practical implication is that buffet operators should carefully consider the interplay between plate size, convenience, and production costs when making design decisions.

Section. 2.2.2 reviews additional waste-reducing interventions studied in the literature. Our finding offers an operational and analytical perspective for understanding the mixed findings in evaluating the effectiveness of plate size reduction in the previous research [111, 128, 85, 58], and our contribution adds a critical nuance: even when smaller plates succeed in reducing waste, they do not always enhance profitability. Understanding when these goals align—and when they diverge—is essential for effective buffet design.

2.7 Directions for Future Research

Convenience in All-you-can-eat Buffets: Our study provides key insights into the interplay between convenience, plate waste, and profitability in all-you-can-eat buffets, but several aspects were left unaddressed, offering rich opportunities for future research.

First, while our analysis primarily focuses on operational dynamics, it does not account for behavioral factors. One such factor is the sunk cost fallacy, which manifests in two ways. The first is related to the flat fee diners pay upon entry: the sunk cost of this payment may drive diners to over-consume. Interestingly, this effect could intensify or weaken depending on the level of convenience, due to the non-monotonic relationship between entrance fees and convenience that we established in our paper. The second form of the sunk cost fallacy is tied to the inconvenience of refilling plates. If it takes longer or requires more effort to refill a plate, diners may overestimate the value of that effort and take more food. However, this behavior diminishes as convenience increases, therefore, in all likelihood, reinforcing the benefits outlined in our analysis. Future research could explore these dynamics further.

Second, our model does not account for the cost of offering increased convenience. While we expect this cost to be relatively trivial, in reality, making a buffet overly convenient may become prohibitively expensive for operators. A natural extension would be to model the cost

of convenience as a quadratic function, where too much convenience could erode profitability beyond a certain point. Although this is not the focus of our study, we acknowledge that, in practice, there may be diminishing returns to convenience, a point that could predictably modify the insights from Theorem 2.7.

Finally, our model assumes homogeneous food offerings—an abstraction that simplifies the analysis but omits the role of variety. This assumption is reasonable in institutional settings like corporate or university cafeterias (which our paper targets), where diners are primarily there to eat and leave—rather than to explore or savor the dining experience—and variety-seeking behavior is less prominent. In such contexts, our findings are expected to hold. One could formally extend our model to include two food types, allowing diners to choose both items simultaneously in the first round, allocate quantities based on expected taste match, and potentially refill with the preferred item. As long as diners behave rationally and use available information efficiently, increased convenience should not overturn our main results. We did not pursue this extension, as it would require introducing substantial additional structure not present in our baseline model, with limited added value: the core insights would likely remain unchanged.

In contrast, in leisure-oriented buffets—where variety is central—examining this dimension would likely yield valuable insights. In such settings, convenience may interact with variety in more complex ways. On one hand, convenience may reduce how much is taken per trip; on the other, it may encourage experimentation across more dishes, increasing the likelihood of uneaten portions. It may also drive overconsumption. Understanding these trade-offs would help clarify how variety influences waste and satisfaction in multi-option leisure buffets, which lie beyond the scope of this paper.

Waste-Reducing Interventions: Alongside convenience, we explored two waste-reducing instruments—monetary penalties and consumer education—both of which are under-researched in buffet settings. However, there is value in studying a broader range of waste-reduction strategies. For example, field evidence suggests that trayless dining or gamification might be effective interventions. Comparing these strategies with the penalty and education instruments analyzed in this paper would provide deeper insights into which approaches are most effective in reducing waste while preserving profitability.

This line of questions overall presents an intriguing future research direction that could offer buffet operators additional tools to optimize both sustainability and financial outcomes. By studying how these different waste-reduction strategies interact with buffet design and customer behavior, we can continue to build on the foundation laid by our work and expand the operational toolkit for buffet management.

CHAPTER 3

Inventory Recommendations in Platform-Operated Fulfillment Service: Trust and Information Asymmetry

This chapter is based on joint work with Fan Zhou and Stefanus Jasin.

3.1 Introduction

Platform-operated fulfillment (POF) services—such as Fulfillment by Amazon (FBA), Walmart Fulfillment Services (WFS), JD Logistics, and Alibaba's Cainiao—have emerged as a core operational backbone of modern e-commerce. The global e-commerce fulfillment market was valued at over USD 114 billion in 2024 and is projected to approach USD 280 billion by 2033 [106]. This rapid growth is fueled in large part by services like FBA, which are now used by approximately 82% of Amazon marketplace sellers—a segment that accounts for more than 60% of total platform sales [46, 92]. These services allow independent sellers to outsource warehousing, packing, and shipping to the platform. By tightly integrating marketplace functions with sophisticated logistics networks, platforms enhance delivery speed and reliability for consumers while simultaneously reducing operational complexity for sellers. The operating model is straightforward: sellers send their products to a platform-managed warehouse, and in exchange for fees covering storage and fulfillment, the platform manages the entire order cycle. This arrangement enables sellers to concentrate on core business activities such as product development and marketing [121].

Despite their logistical strengths, POF ecosystems face a fundamental operational challenge: platforms exercise limited direct control over sellers' replenishment decisions. In practice, platforms deploy advanced decision-support tools to recommend order quantities, such as Amazon's Automatic Replenishment system [26]. Nevertheless, as researchers from Amazon's FBA Science team observe, "sellers are their own best judges and have the flexibility to set their own policies" [94]. This autonomy has resulted in limited adoption of platform-generated recommendations. Although systematic evidence is scarce, industry analyses and

seller forums consistently point to a disconnect between platform tools and sellers' needs. Commonly cited shortcomings include inadequate adjustments for seasonality and failure to incorporate the effects of promotional events [86, 110]. These frictions represent a critical operational bottleneck, constraining the efficiency gains that platforms aim to achieve.

Two issues are particularly salient. First, objective misalignment arises from differing cost structures. Sellers bear the full Cost of Goods Sold (COGS) and the risk of inventory obsolescence, which incentivizes conservative stocking policies to mitigate the risk of overstocking [110]. Platforms, by contrast, earn revenue from each sale without holding inventory risk, and therefore prefer higher stock levels to maximize sales volume and commission revenue—a classic "conflict of interest." Second, informational asymmetry persists because platforms possess richer and timelier demand signals (e.g., market-wide trends, competitor performance, and aggregate customer behavior) than individual sellers, who typically operate with limited analytical resources [68, 140, 40]. This imbalance implies that even well-intentioned, data-driven recommendations may be rationally disregarded by sellers. A central factor in overcoming this barrier is seller trust. As [94] emphasize, "Amazon's seller assistance is only effective if it earns seller trust," underscoring its pivotal role in achieving alignment between platforms and sellers.

One possible remedy to misaligned incentives is contract redesign, which seeks to realign service-level incentives through adjustments to commission and fee schedules—that is, the standardized tables of commissions, fulfillment charges, and storage fees published by the platform. In practice, however, this approach is rarely feasible. Major platforms adopt standardized fee structures at a high strategic level [91, 44], with changes occurring only infrequently, as documented in empirical studies [30]. These uniform schedules apply to a vast and heterogeneous seller base with widely varying cost structures and risk preferences. Indeed, designing and dynamically tailoring contracts for millions of sellers would be operationally prohibitive, underscoring the inherent limitations of contract-based approaches to incentive alignment. It is worth noting that, in the academic literature, fee parameters are likewise treated as exogenous, reflecting their role as instruments of strategic governance rather than outcomes of seller-level negotiation [120, 80].

Given the intractability of contract redesign, a more practical approach is to refine the platform's recommended order quantities [94]. Rather than prescribing a platform-centric optimal quantity—which may appear overly aggressive to sellers—the platform can strategically nudge recommendations toward levels that are more acceptable from the seller's perspective. This preserves the benefits of the platform's superior demand forecasting while explicitly accounting for seller adoption behavior. As sellers observe consistent positive outcomes from following such guidance, their trust in the platform is reinforced, which in turn strengthens

future incentives for adoption. Trust thus becomes the critical mediator: it enables platforms to trade off some degree of short-term optimality in exchange for greater long-term adoption, fostering a more collaborative equilibrium and mutually beneficial outcomes [98].

This motivates a set of targeted research questions:

1. Should the platform provide order-quantity recommendations, and do such recommendations generate sustained long-run benefits for both the platform and the seller? (Answered in Proposition 3.3.)

2. How should the platform optimally trade off short-run profit extraction against long-run trust formation when designing its recommendation policy? (Answered in Proposition 3.4.)

3. Under the platform's profit-maximizing recommendation stance, how large is the resulting coordination surplus, and how is this surplus divided between the platform and the seller? (Answered in Section 3.6.2.)

4. How does demand censoring—arising from the lack of visibility into the realized demand following stockouts—alter the platform's optimal strategy, trust dynamics, and long-run outcomes? (Answered in Sections 3.5 and 3.6.3.)

5. How sensitive are the platform's optimal recommendations and long-run outcomes to different seller trust behaviors? (Answered in Section 3.6.4.)

To address these questions, we develop a repeated inventory–recommendation model in which a platform with superior demand information advises a seller on order quantities, while the seller retains ultimate control. The seller's willingness to adopt the platform's advice depends on her evolving level of trust: if following a recommendation improves her profit relative to her own decision, trust increases; otherwise, it declines. The platform thus faces a fundamental trade-off. It can recommend quantities that are more favorable to the seller in order to build trust and encourage future adoption, or it can prioritize its own short-term profit at the risk of eroding trust and diminishing adoption over time. We formalize this tension by studying policies under which the platform commits to a fixed "stance" in its recommendations, lying between the seller's interests and the platform's own.

We find that the interaction between the platform's recommendations and the seller's responses converges to a steady state. In this steady state, the seller's long-run trust settles at a level that reflects how frequently the platform's advice delivers higher profit than the seller's own decision. This trust level, in turn, governs the likelihood of future adoption and

forms the basis for our main results. From this steady-state characterization, several insights emerge.

First, our analysis shows that, wherever the platform's stance in recommendations lies, sharing information through recommendations can always benefit both parties (relative to no recommendation), yielding a "win–win" outcome.

Second, when sellers fully observe realized demand in each period, their response to platform advice depends on their stocking orientation. Conservative sellers, who are highly averse to excess inventory, lose trust when faced with aggressive recommendations. By contrast, aggressive sellers, who are more willing to stock deeply, gain trust from aggressive recommendations because these help them avoid shortages. The platform's optimal stance therefore varies with seller type: with aggressive sellers it benefits from pushing strongly toward its own preferred recommendations, whereas with cautious sellers it may need to moderate its advice—sometimes even recommending the exact quantity the seller would have chosen independently.

Surprisingly, we also find that a stronger private demand signal for the platform does not always increase its steady-state profit. When sellers occasionally disregard recommendations, the opportunity cost of non-adoption is greatest precisely when the ignored advice would have been most valuable.

Third, we quantify the coordination surplus generated by the platform's profit-maximizing recommendation stance and examine how it is split between the platform and the seller. The resulting surplus is larger when the platform's informational advantage is stronger and when the conflict of interest is more pronounced. However, its allocation is sensitive to the information environment: the platform tends to capture a larger share in low signal-to-noise settings, whereas the seller captures more value when demand signals are relatively precise.

Finally, the dynamics change markedly when sellers and the platform do not observe true demand in the event of stockouts, a situation commonly found in practice. For this case, we consider Bayesian sellers who form beliefs about how the recommendation would have performed under censored outcomes. Stockouts create what we call an "informational shield." When demand is fully observable, sellers can assess both sides of the recommendation: whether ordering more would have avoided shortages, but also whether it would have led to costly overstocking. Under censored demand, however, sellers lose visibility of the downside. They cannot tell whether the platform's larger recommendation would have generated excess inventory, but they still see the lost sales from understocking. This asymmetry systematically tilts trust updates in the platform's favor, making aggressive recommendations appear safer than they truly are. As a result, even cautious sellers tend to build trust under more aggressive recommendations, and the platform often adopts a fully self-interested

stance. Only when the conflict of interest is very large do sellers revise their beliefs sharply at a tipping point, forcing the platform to moderate just enough to preserve trust. Perhaps most counterintuitively, we find that the platform can sometimes earn higher long-run profits under censored demand than under full observability. In volatile markets, the trust-building effect of the informational shield can outweigh the value of complete transparency.

Importantly, the two information regimes we study—fully observed demand and censored demand—are both practically relevant. While full demand observability can be challenging in the presence of stockouts, sellers' ability to infer latent demand varies substantially across products and contexts. In some settings, sellers can partially reconstruct true demand using auxiliary feedback signals, such as customer inquiries, backorder requests, or marketplace diagnostics, making the fully observed-demand regime a reasonable approximation. In other settings, lost sales remain largely unobserved, and demand is effectively censored. Our results therefore highlight that the platform's optimal recommendation stance should account not only for incentive alignment and information asymmetry, but also for the degree of demand observability faced by the seller.

The remainder of the chapter is organized as follows. Section 2 reviews the related literature. Section 3 introduces the model, detailing the inventory framework, information structure, and trust dynamics. Section 4 develops the steady-state analysis under fully observed demand, while Section 5 extends the analysis to settings with censored demand. Section 6 concludes with a discussion of implications and directions for future research.

3.2 Literature Review

This study relates to three main streams of research: (i) platform operations, (ii) fulfillment services, and (iii) information sharing and trust. Across these literatures, our contribution is distinct in modeling the design of platform-generated inventory recommendations under both information asymmetry and endogenous trust, and in deriving the resulting steady-state implications for adoption and profit.

3.2.1 Platform Operations

Research on platform operations has examined selling formats, governance, pricing, competition, and information mechanisms. A central theme concerns the choice between agency selling, reselling, and hybrid designs. Agency selling is preferred when cross-channel spillovers are negative, absent, or small, whereas reselling is favored when spillovers are strongly positive [1, 53]. Pricing control—centralized versus decentralized—balances the benefits of reg-

ulating competition with the need to tailor prices; fee structures can decentralize pricing without reducing profit [20]. Our setting aligns with the agency model: sellers retain control over inventory decisions but are influenced by platform recommendations.

Competition and market dynamics highlight both vertical and horizontal tensions. With asymmetric information, competition can reduce trade relative to monopoly, though multi-homing alleviates this effect [54]. Platforms that sell alongside merchants must manage encroachment incentives [113]. Empirical evidence shows that Amazon's entry targets successful products, discourages seller growth, but simultaneously raises demand and reduces shipping costs [140]. Peer-to-peer markets show elastic supply and linear scaling even without network effects [38], while greater market thickness can reduce matches due to search frictions [82].

Matching and reputation mechanisms also play a central role. Collaborative consumption shifts ownership and usage equilibria while benefiting consumers [13]. In ride-hailing, inefficient driver movements reduce welfare, but pricing and matching mechanisms can mitigate this [25]. Search design can improve equilibria by withholding some information [74], while partial reputation disclosure can sometimes dominate full disclosure [115]. Recent syntheses of platform analytics and omnichannel experiments highlight complementarities and research opportunities [139].

3.2.2 Fulfillment Service

A second stream of research examines fulfillment services in e-commerce settings. Equilibrium analyses highlight complementarities between fulfillment and sales channels, with platform-provided fulfillment often part of win–win outcomes [61]. Traditional retailers typically rely on distribution centers for efficiency, though store networks can reduce last-mile costs [66]. Sharing logistics with rivals may also be optimal, depending on third-party provider quality and market size [105].

The FBA-specific literature emphasizes redistribution and strategic effects. FBA benefits low-service sellers and consumers but reduces welfare for high-service sellers, while platforms consistently gain [67]. It intensifies service competition and can mitigate price competition, generating win–win–win outcomes under certain product conditions [80]. Inventory commingling advantages platforms and lower-quality sellers but may hurt higher-quality sellers through endogenous fee increases [129]. Empirically, FBA has been shown to increase sales but also returns, leading to ambiguous net effects; risk-adjusted models recommend shifting "safe" products into FBA [117, 118]. Offering fulfillment services to rivals benefits dominant retailers absent competition, but creates ambiguity under direct rivalry [83]. In

inter-platform competition, subsidized FBA can encourage voluntary adoption and welfare gains, though excessive subsidies risk coercion or predation [84].

Finally, research has begun to examine the role of recommendations within fulfillment services. Amazon's recommendation systems improve seller outcomes, with quantified efficacy shown to boost adoption [94]. Our study complements this literature by endogenizing sellers' adoption of fulfillment recommendations under evolving trust.

3.2.3 Information Sharing and Trust

The third stream examines information provision and trust. We note that in this paper, the "information sharing" is implemented in an indirect sense. Rather than directly disclosing its private demand signal or forecast, the platform embeds this information in an inventory recommendation (i.e., a suggested order quantity), and the seller decides whether to act on it.

Platforms disclose information strategically: ambiguous forecasts can be credibly shared, while precise ones often cannot [10]. Selective demand disclosure may raise platform profits but reduce consumer welfare [16], while delaying the revelation of supplier quality can accelerate market discovery [17]. Disclosure incentives depend on market dominance [62], with asymmetric sharing frequently optimal [87].

Information asymmetry also shapes competition and encroachment. Integrated platforms leverage informational advantages to strengthen market power [30]. Selling formats interact with disclosure: sharing is more likely under agency contracts with high commissions [49]; supplier encroachment reduces wholesale prices and complements information sharing [51]; and equilibrium outcomes can range from full to selective disclosure [81, 138].

Operational research highlights the mechanisms that support credible information exchange. Nonlinear contracts incentivize truthful sharing more effectively than linear ones [52], and the distribution of benefits depends on seller cost structures [64]. Information sharing reduces uncertainty but creates misreporting risks in the absence of penalties [134]. Transparency encourages agency adoption and can deliver Pareto improvements [123]. Forecast-sharing incentives depend on manufacturer efficiency and may lead to an "accuracy trap" [137]. Service capability influences both contractual format and disclosure [15], while commission rates shape incentives to share [63].

Complementing these structural mechanisms, trust is essential for credible information exchange. Laboratory experiments show that cheap-talk forecasts become informative once trust is embedded in preferences [98]. Evidence also suggests that the form of assistance matters: information sharing builds more trust than advice or delegation [95]. Broader syn-

theses identify values, infrastructure, and processes as antecedents of trust [96, 97]. Dynamic trust models capture how retailers update the weight placed on external recommendations as a function of observed precision [42].

Relative to these literature, our contribution is twofold. Compared to studies that assume receivers fully trust senders, we model trust as an evolving state that mediates whether platform recommendations are adopted and whether informational advantages translate into realized value. We analyze a repeated inventory–recommendation game in which the platform commits to a static stance that blends platform- and seller-optimal quantities. Sellers update trust from realized (or censored) outcomes, and steady-state trust reflects the probability that platform recommendations outperform the seller's benchmark. Our analysis identifies conditions under which greater self-centering raises or lowers trust, and shows how demand censoring shields the platform, enabling more aggressive recommendations. Relative to the platform disclosure literature [e.g., 16, 87, 10], we shift the focus from what to disclose to how to recommend, demonstrating when softer recommendations build trust and when aggressive ones are optimal.

3.3 Model Setup

We study a repeated newsvendor setting in which each cycle replicates a single-period newsvendor problem. The platform (indexed by P and referred to as "he") possesses superior demand forecasting capabilities and recommends an order quantity to the seller (indexed by S and referred to as "she") to meet upcoming demand. The final order decision, however, remains with the seller.

Profits from fulfilled orders are shared under an asymmetric profit-sharing contract that favors the platform. This arrangement reflects common practice on major e-commerce platforms such as Amazon and JD, which charge sellers a substantial referral fee on each sale while passing most storage and holding costs onto the sellers themselves [126, 5]. Upon receiving a recommendation, the seller decides whether to adopt it or to rely on her own judgment, based on her trust in the platform. This trust evolves endogenously according to the historical performance of the platform's recommendations in improving the seller's profit relative to her own benchmark decisions.

The remainder of this section describes the underlying newsvendor framework and profit-sharing scheme, the dynamics of trust evolution, and the recommendation policies available to the platform.

3.3.1 Newsvendor Setup

Each unit is sold at a fixed price p and incurs a unit cost c. The platform retains a fraction $\gamma \in [0,1]$ of the revenue and covers a fraction $\beta \in [0,1]$ of the cost. Accordingly, the platform's margin per fulfilled unit is $\gamma p - \beta c$, while the seller's margin is $(1-\gamma)p - (1-\beta)c$. Under the assumed asymmetric contract, where $\gamma > \beta$, the platform captures a larger share of revenue while bearing a smaller share of cost. Within the standard newsvendor framework, this asymmetry implies that the platform's optimal service level (r_P) is strictly higher than the seller's (r_S).[1]

That is,

$$\gamma > \beta \implies \underbrace{1 - \frac{\beta c}{\gamma p}}_{r_P} > \underbrace{1 - \frac{(1-\beta)c}{(1-\gamma)p}}_{r_S}. \tag{3.1}$$

Demand in period t is specified as

$$d_t = \underbrace{\mu}_{\text{baseline demand}} + \underbrace{\xi_t}_{\text{seasonal component}} + \underbrace{\epsilon_t}_{\text{randomness}}. \tag{3.2}$$

Demand in period t is decomposed into three components: a baseline level μ, a seasonal state ξ_t, and a random variation ϵ_t. Here, μ is a constant baseline demand known to both the platform and the seller. The seasonal component ξ_t is privately observed by the platform, reflecting its informational advantage from access to richer and timelier marketplace signals (e.g., aggregate consumer traffic, competitor performance, and platform-wide promotional activity) that are not fully available to individual sellers at the time of ordering.

For simplicity, we assume that ξ_t follows a symmetric two-point distribution that is common knowledge: $\xi_t \in \{-\xi, \xi\}$ with $\Pr(\xi_t = -\xi) = \Pr(\xi_t = \xi) = 0.5$. We interpret $\xi_t = \xi$ as a high season and $\xi_t = -\xi$ as a low season. The term ϵ_t captures residual market-level randomness, independently drawn from a mean-zero uniform distribution, $\epsilon_t \sim U[-\epsilon, \epsilon]$. This distribution is known to both parties ex ante, but its realization is not observed at the time of ordering.

Table 3.1 summarizes the definitions of the demand components and specifies which party has access to each. For reference, F_ξ and F_ϵ denote the cumulative distribution functions (c.d.f.) of ξ_t and ϵ_t, respectively.

[1] In the standard newsvendor setting, the service level refers to the probability of meeting demand, defined as the underage cost divided by the sum of underage and overage costs.

Table 3.1: Demand components and information access

Variable	**Definition**	**Information Access at the Beginning of Period t**
μ	Baseline demand	Known to both platform and the seller
ξ_t	Seasonal component	F_ξ known to both; ξ_t observed by platform only
ϵ_t	Random noise	F_ϵ known to both platform and the seller

Given the platform's recommendation, the seller chooses an order quantity q_t, leading to the following profit functions:

$$\Pi^P(q_t, d_t) = \gamma p \cdot \min(q_t, d_t) - \beta c \cdot q_t, \tag{3.3}$$

$$\Pi^S(q_t, d_t) = (1-\gamma) p \cdot \min(q_t, d_t) - (1-\beta) c \cdot q_t. \tag{3.4}$$

Following the standard newsvendor framework, we define three order quantities: the platform's optimal order quantity $q_t^P(\xi)$, the seller's optimal order quantity $q_t^S(\xi)$, and the seller's perceived optimal order quantity $\hat{q}_t^S$:

$$q_t^P = q_P(\xi_t) := \arg\max_q \mathbb{E}_{F_\epsilon}\Big[\Pi^P(q, \mu + \xi_t + \epsilon_t \mid \xi_t)\Big] = \mu + \xi_t + F_\epsilon^{-1}(r_P), \tag{3.5}$$

$$q_t^S = q_S(\xi_t) := \arg\max_q \mathbb{E}_{F_\epsilon}\Big[\Pi^S(q, \mu + \xi_t + \epsilon_t \mid \xi_t)\Big] = \mu + \xi_t + F_\epsilon^{-1}(r_S), \tag{3.6}$$

$$\hat{q}_t^S := \arg\max_q \mathbb{E}_{F_{\xi+\epsilon}}\Big[\Pi^S(q, \mu + \xi_t + \epsilon_t)\Big] \quad = \mu + F_{\xi+\epsilon}^{-1}(r_S). \tag{3.7}$$

Here, $F_{\xi+\epsilon} := F_\xi * F_\epsilon$ denotes the convolution of F_ξ and F_ϵ, i.e., the c.d.f. of the combined uncertainty $\xi+\epsilon$. In the above, $q_t^P(\xi_t)$ and $q_t^S(\xi_t)$ represent the actual optimal order quantities for the platform and the seller, respectively, when both account for the seasonal component ξ_t. Due to asymmetric revenue sharing ($r_P > r_S$), the platform favors a higher stocking level than the seller, i.e., $q_t^P > q_t^S$.

As for $\hat{q}_t^S$, it denotes the seller's perceived optimal order quantity under her limited information. Specifically, without observing the seasonal component ξ_t, she optimizes her decision based only on the distribution of the combined uncertainty $\xi_t + \epsilon_t$, leading to $\hat{q}_t^S = \mu + F_{\xi+\epsilon}^{-1}(r_S)$. This rough estimate can deviate from her true profit-maximizing quantity: $\hat{q}_t^S$ may be either higher or lower than q_t^S depending on the realized seasonality, as $q_t^S - \hat{q}^S = \xi_t \in \{+\xi, -\xi\}$.

3.3.2 Signal-to-Noise Ratio (SNR)

The seasonal component $\xi_t \in \{-\xi, \xi\}$ is privately observed by the platform, while the random noise $\epsilon_t \in [-\epsilon, \epsilon]$ is symmetrically unknown to both parties. The relative importance of the signal to the noise, as perceived by the platform, is captured by the ratio ξ/ϵ, which we refer to as the Signal-to-Noise Ratio (SNR) of the market environment. The SNR provides a natural measure of the platform's informational advantage over the seller. Two extreme cases illustrate its interpretation. When $\xi = 0$ and $\epsilon > 0$, the seasonal component disappears and both parties face the same uncertainty—effectively "walking in the same darkness." Conversely, when $\xi > 0$ and $\epsilon = 0$, the platform observes demand perfectly, gaining complete foresight relative to the seller.

To avoid trivial cases and focus on meaningful regimes, we impose assumptions that restrict attention to an intermediate range of SNR values.

Assumption 3.1. *The seller may experience stockouts under her perceived optimal order quantity in both the low and high seasons, with stockouts more likely in the high season. Formally,*

$$0 < \Pr(d_t > \hat{q}_S \mid \xi_t = -\xi) < \Pr(d_t > \hat{q}_S \mid \xi_t = \xi) < 1. \tag{3.8}$$

In addition, there is a nonzero probability that the seller's perceived optimal order quantity exceeds the platform's optimal order quantity in at least one seasonal state:

$$\hat{q}_S \;\geq\; \min_{\xi_t \in \{-\xi, \xi\}} q_P(\xi_t). \tag{3.9}$$

Together, these conditions are equivalent to

$$(3.8), (3.9) \quad \Longleftrightarrow \quad 2(r_P - r_S) \;\leq\; \frac{\xi}{\epsilon} \;<\; 2\min\{r_S,\, 1 - r_S\}. \tag{3.10}$$

Condition (3.8) requires that the seller's perceived order quantity $\hat{q}_S$ is neither too small that stockouts occur with certainty in both seasons, nor too large that stockouts never occur. Instead, stockouts must be possible in both the low and high seasons, with a strictly higher probability in the high season. This reflects realistic retail settings, where sellers occasionally face shortages in both off-peak and peak periods, though shortages are naturally more common during peak demand. It also implies that the two seasonal demand distributions partially overlap: the lowest possible demand in the high season is smaller than the highest possible demand in the low season.

Condition (3.9) bounds the degree of interest misalignment between the platform and the seller. It rules out degenerate cases in which the platform always recommends a strictly

higher order quantity than the seller's benchmark, regardless of the seasonal state. Instead, the condition requires that the seller's perceived optimal order $\hat{q}_S$ can coincide with or exceed the platform's recommendation in some states. This assumption reflects common practice. Although sellers are often more conservative than platforms on average, they may occasionally overstock relative to platform guidance due to risk preferences, inventory targets, or idiosyncratic beliefs about demand.

Taken together, conditions (3.8) and (3.9) pin down a nontrivial range of the signal-to-noise ratio (SNR), as given in (3.10): the lower bound corresponds to (3.9), while the upper bound corresponds to (3.8). Beyond their practical interpretations discussed above, these bounds are also needed to avoid degenerate trust dynamics. Specifically, the two conditions ensure that under the seller's benchmark $\hat{q}_S$, the performance differential $\Delta\Pi_t^S$ is genuinely stochastic and takes both positive and negative values with positive probability. This stochastic variation is essential for trust to evolve endogenously rather than converging mechanically to full compliance or complete rejection.

If condition (3.8) fails, then in at least one seasonal state the seller's benchmark order induces a deterministic outcome (i.e., either permanent overstock or permanent understock), so the sign of $\Delta\Pi_t^S$ becomes (almost) fixed. In this case, trust updates no longer reflect meaningful learning and instead drift monotonically toward an extreme. If condition (3.9) fails, the platform's recommendation strictly exceeds $\hat{q}_S$ in all seasons. This removes the low-season overstock-mitigation channel that disciplines aggressive recommendations, collapses the platform's optimal stance to a corner solution, and eliminates the interior trade-offs between profit extraction and trust formation that are central to our analysis.

3.3.3 Trust Evolution and Static Recommendation

At period t, the seller's trust in the platform is represented by $\alpha_t \in [0, 1]$, which can be interpreted as the probability that the seller adopts the platform's recommendation rather than relying on her own forecast. Let q_t^R denote the order quantity recommended by the platform. The seller's actual order decision is determined as follows:

$$q_t = \begin{cases} q_t^R, & \text{with probability } \alpha_t, \\ \hat{q}_S, & \text{with probability } 1 - \alpha_t. \end{cases} \tag{3.11}$$

After demand is realized, which we assumed to be publicly observed in the baseline model, the seller updates her trust by evaluating the performance of the recommendation, measured by the profit differential between following the recommendation and relying on her

own decision:

$$\Delta_t \Pi^S := \Pi^S(q_t^R, d_t) - \Pi^S(\hat{q}_S, d_t). \tag{3.12}$$

Trust then evolves according to an exponential-smoothing update rule:

$$\alpha_{t+1} = \begin{cases} \delta \cdot 1 + (1-\delta) \cdot \alpha_t, & \text{if } \Delta_t \Pi^S > 0, \\ \alpha_t, & \text{if } \Delta_t \Pi^S = 0, \\ \delta \cdot 0 + (1-\delta) \cdot \alpha_t, & \text{if } \Delta_t \Pi^S < 0, \end{cases} \tag{3.13}$$

where $\delta \in (0, 1)$ is the seller's learning rate. In each period, trust is updated as a convex combination of the current trust α_t and an extreme belief (1 if the recommendation outperformed, 0 if it underperformed). Similar formulations using exponential smoothing to model trust evolution have been widely applied in the literature [see, e.g., 23, 42, 101]. A positive experience shifts trust upward, a negative experience shifts it downward, and a neutral experience leaves it unchanged. The learning rate δ determines the speed of adjustment: when δ is close to 1, the seller heavily weights the most recent outcome, making trust volatile and quick to change; when δ is small, trust evolves gradually, reflecting accumulated history. This formulation captures the idea that trust is earned or lost through repeated interactions, rather than determined by any single experience alone.

3.4 Steady-state Analysis

As noted earlier, in this paper, we focus on a time-stationary recommendation policy—what we term a static recommendation—that drives the system toward a long-run stationary state.

3.4.1 Static Recommendation

Under a static recommendation policy, the platform's recommendation is fixed and does not depend on the seller's current trust level. Specifically, we consider a family of policies where the recommended order quantity is a convex combination of the platform's and the seller's optimal orders:

$$\bar{q}_t^R(\eta) = \eta q_t^P + (1-\eta) q_t^S, \tag{3.14}$$

where $\eta \in [0, 1]$ denotes the degree of self-centering (DoS). This parameter governs the stance of the platform's recommendation: $\eta = 0$ corresponds to a fully seller-centered recommendation, $\eta = 1$ corresponds to a fully platform-centered recommendation, and intermediate values reflect a compromise between the two. Economically, the DoS captures the extent to

which the platform is willing to sacrifice short-term profit in order to improve the likelihood of seller adoption and build long-run trust. The platform commits to a fixed DoS throughout the game; however, the seller cannot observe either η or the underlying quantities q_t^P and q_t^S.[2]

We focus on a time-stationary (static) recommendation policy for two strategic reasons. First, commitment to consistency is a prerequisite for trust formation. A dynamic policy creates a time-inconsistency problem in which the platform is incentivized to "harvest" reputation by switching to selfish recommendations once trust is established. A static policy therefore serves as a credible commitment device, e.g., [73], ensuring that the platform cannot exploit the seller's accumulated trust. Second, considerations of robustness and learnability motivate a static policy. For the seller to update her trust meaningfully, the recommendation strategy must constitute a stable signal-generating process. A dynamic policy would instead present the seller with a "moving target," confounding the learning process [98]. Moreover, the mechanism design literature suggests that simple, time-invariant mechanisms are significantly more robust to model uncertainty (e.g., unknown seller learning rates) than complex, state-dependent policies, e.g., see, [21, 24].

The sequence of events in period t (the t-th cycle of the single-period newsvendor problem) is summarized in the flowchart below.

Figure 3.1: Flow of events in cycle t

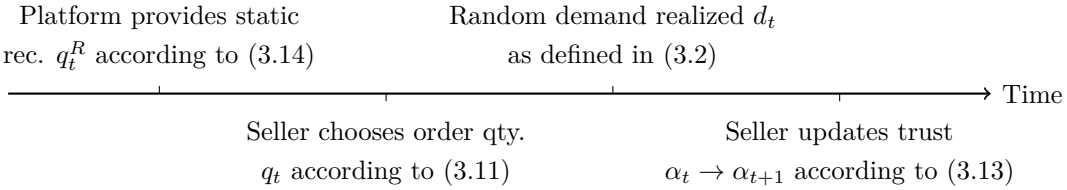

3.4.2 Steady-state Trust Level

The trust dynamics described in Section 3.3 evolve stochastically over time as the seller updates her belief based on the relative performance of the platform's recommendation. While individual trust realizations may fluctuate from period to period, our interest lies in the long-run behavior of the system under a fixed recommendation stance. In particular, we ask: to what level does trust settle on average, and how volatile is it in the long run?

The following proposition characterizes the steady-state trust level. It shows that the expected trust converges to the probability that the static recommendation outperforms

[2]The seller only observes the platform's recommendation $\bar{q}_t^R(\eta)$, but not its decomposition into q_t^P, q_t^S, or η.

the seller's perceived optimal order, which we denote by $\rho(\eta)$. This probability summarizes the credibility of the platform's recommendation in the long run under a given degree of self-centering.

Proposition 3.2. *Under a static recommendation policy, the expected trust level converges to the probability that the recommended order quantity outperforms the seller's perceived optimal decision. Formally, we have:*

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\alpha_t=\lim_{t\to\infty}\mathbb{E}[\alpha_t]=\Pr\left[\Pi^S(\bar{q}_t^R(\eta),d_t)>\Pi^S(\hat{q}_S,d_t)\right]:=\rho(\eta),$$

where the probability of outperformance is

$$\rho(\eta)=\frac{1}{2}+\frac{\xi}{4\epsilon}-\frac{(1-2r_S)(r_P-r_S)}{2}\,\eta. \tag{3.15}$$

The steady-state variance of the trust level is

$$\lim_{t\to\infty}\mathrm{Var}(\alpha_t)=\frac{\delta\rho(\eta)\left(1-\rho(\eta)\right)}{2-\delta}.$$

Figure 3.2: An example of sample paths of trust α_t and its cumulative average $\frac{1}{T}\sum_{t=1}^{T}\alpha_t$ for three values of DoS.

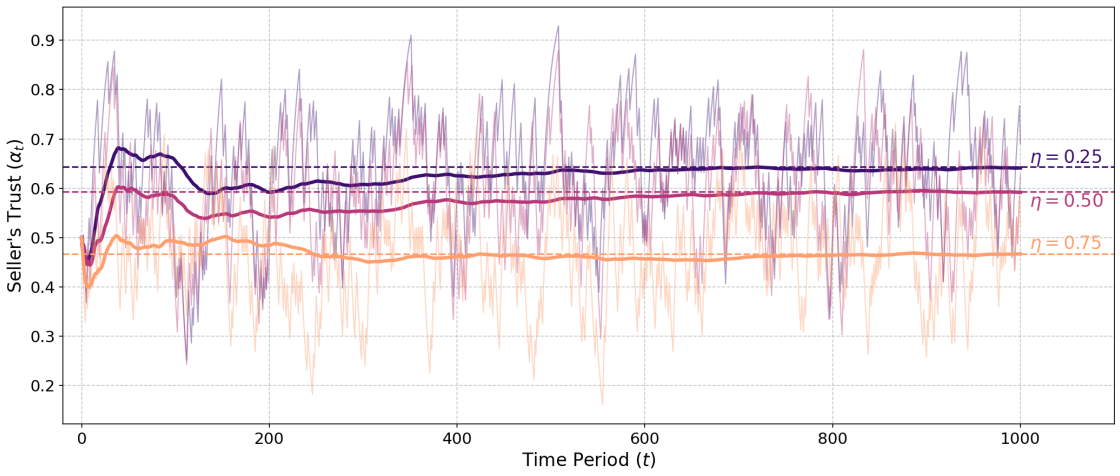

Note: $\eta \in \{0.25, 0.5, 0.75\}$, under $r_S < 1/2$

Figure 3.2 illustrates the convergence behavior characterized in Proposition 3.2 by plotting representative sample paths of trust α_t and its cumulative average $\frac{1}{T}\sum_{t=1}^{T}\alpha_t$. As shown in the figure, although trust fluctuates over time, its long-run average converges. Notably, for sellers with $r_S < 1/2$, a higher degree of self-centering (η) leads to a lower steady-state trust

level.

The probability $\rho(\eta)$ arises from two mutually exclusive "good news" scenarios. The first scenario, which we call High-season Understock Mitigation, occurs when $\xi_t = \xi$. In this case, the seller's benchmark order $\hat{q}_S$ tends to be too small, leading to stockouts. The recommendation improves performance if it raises the order enough to avoid lost sales, but not so much that it induces excessive leftover inventory. The second scenario, which we call Low-season Overstock Mitigation, occurs when $\xi_t = -\xi$. Here, the seller's benchmark order $\hat{q}_S$ tends to be too large, leading to costly overstocking. The recommendation improves performance if it reduces the order enough to cut inventory waste, but not so much that it creates new stockouts.

Under the asymmetric contract ($r_P > r_S$), increasing the degree of self-centering η shifts the recommendation upward. The net effect of this shift can be measured by (3.15), which yields:

$$\rho'(\eta) = -\frac{1}{2}(1 - 2r_S)(r_P - r_S).$$

Accordingly, platform self-centering erodes credibility when the seller is cautious ($r_S < 1/2$) but enhances credibility when the seller is aggressive ($r_S > 1/2$). Hence, whether more platform-centric recommendations build or undermine trust depends critically on the seller's stocking posture.

The variance expression in Proposition 3.2 also provides important insights into the stability of trust in steady state. Two forces shape this stability. The first is the learning rate δ. When δ is small, updates are incremental and variance approaches zero, as the seller places little weight on any single realization. In contrast, when δ is close to one, the seller reacts strongly to each outcome, and variance reaches its maximum level of $\rho(\eta)(1 - \rho(\eta))$. The second force is recommendation performance, captured by $\rho(\eta)$. Variance is maximized when $\rho(\eta) = \frac{1}{2}$, i.e., when recommendations are correct only about half the time. In this regime, trust fluctuates the most because each new outcome has the potential to either build or erode trust. Variance is minimized when $\rho(\eta)$ is close to 0 or 1, since consistently poor or consistently strong recommendations drive trust to the extremes, leaving little room for oscillation. Together, these forces imply that volatility in trust arises only when the platform delivers mixed performance and the seller is highly responsive to new information. In practice, stability is achieved either when sellers update their trust slowly (small δ), or when the platform's recommendations are either sufficiently reliable or unreliable.

3.4.3 Steady-state Profit

In steady state, the expected profits for the platform (indexed by $i = P$) and the seller (indexed by $i = S$) are denoted by $\bar{\Pi}^i(\eta)$ and defined as

$$\bar{\Pi}^i(\eta) := \lim_{t\to\infty} \mathbb{E}\left[\Pi^i(q_t, d_t)\right]. \tag{3.16}$$

Incentive compatibility of inventory recommendations. The following result demonstrates that information sharing through inventory recommendations (henceforth abbreviated as "information sharing" or "inventory recommendations" interchangeably, for ease of composition) both parties benefit when the platform provides recommendations.

Proposition 3.3. *Under Assumption 3.1, both the seller and the platform benefit from information sharing. Formally, for any $\eta \in [0, 1]$,*

$$\bar{\Pi}^S(\eta) > \underbrace{\mathbb{E}\left[\Pi^S(\hat{q}_S, d_t)\right] := \Pi_0^S}_{\text{Seller's profit without info. sharing}}, \qquad \bar{\Pi}^P(\eta) > \underbrace{\mathbb{E}\left[\Pi^P(\hat{q}_S, d_t)\right] := \Pi_0^P}_{\text{Platform's profit without info. sharing}}.$$

This result shows that both sides gain in steady state, despite the contract's asymmetry. The key lies in the informational advantage held by the platform. Without recommendations, the seller bases her decision on $\hat{q}_S$, which ignores the seasonal component ξ_t. This leads to systematic errors: in high seasons, the seller understocks and misses sales; in low seasons, she overstocks and incurs excess holding costs. These inefficiencies reduce profits for both parties. When the platform shares recommendations, the seller does not need to adopt them fully for gains to arise. Even partial adoption moves her order closer to the informed quantities $q_t^P(\xi)$ and $q_t^S(\xi)$, mitigating the most severe mismatches between supply and demand. The platform benefits from higher sales, while the seller benefits from fewer costly errors relative to her uninformed benchmark.

The win–win outcome, however, hinges on the conditions in Assumption 3.1. First, the incentive misalignment, captured by the gap $r_P - r_S$, must not be too large; otherwise, the platform's recommendations would always be strictly above the seller's benchmark, and she would rationally disregard them. Second, the platform's informational advantage, measured by the signal-to-noise ratio ξ/ϵ, must be sufficiently strong to ensure that recommendations are informative. Together, these conditions guarantee that information sharing is both credible and mutually valuable.

Figure 3.3 illustrates how the platform's and the seller's long-run average profits vary with η.[3] In both cases, information sharing strictly dominates the no-information benchmark, as

[3]Numerical setup is given in Section 3.6.

indicated by the dashed lines.

Figure 3.3: Platform and seller long-run average profits as a function of the platform's DoS

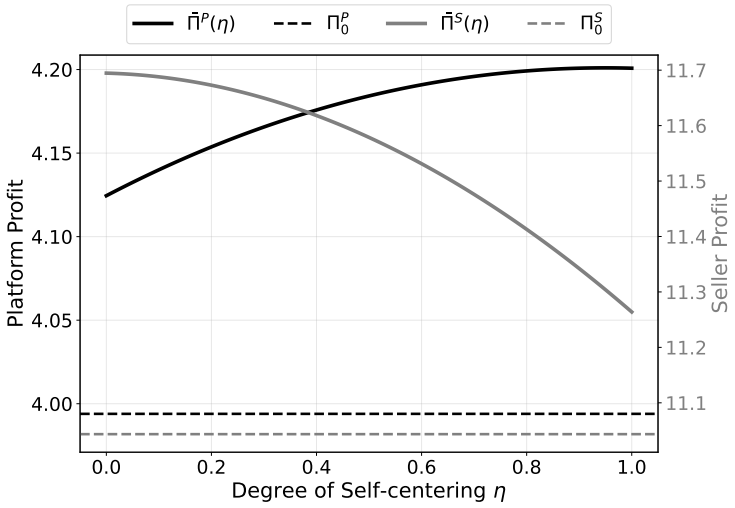

Profit-maximizing Recommendation. The platform determines the profit-maximizing static recommendation by selecting the optimal η, denoted by η_P^*:

$$\eta_P^* = \arg\max_{\eta \in [0,1]} \bar{\Pi}^P(\eta), \tag{3.17}$$

which is characterized as follows.

Proposition 3.4. *Under Assumption 3.1, the platform's profit-maximizing DoS is*

$$\eta_P^* = \begin{cases} 1, & \text{if } r_S \geq \frac{1}{2}, \\ 0, & \text{if } r_S < \frac{1}{2} \text{ and } r_P - r_S \leq \frac{1-2r_S}{4(\epsilon/\xi)+8(\epsilon/\xi)^2}, \\ \tilde{\eta}, & \text{otherwise,} \end{cases} \tag{3.18}$$

where $\tilde{\eta} \in (0,1)$ decreases in $\frac{\xi}{\epsilon}$ and $\frac{p}{c}$. Moreover, η_P^* is weakly increasing in $\frac{\xi}{\epsilon}$ and $\frac{p}{c}$.

Proposition 3.4 shows that the platform's profit-maximizing stance can range from fully self-centered ($\eta = 1$) to fully seller-centered ($\eta = 0$), with intermediate values ($0 < \eta < 1$) balancing these extremes. Which outcome arises depends on the seller's stocking posture and on how closely the platform's preferences align with hers. When the seller is aggressive ($r_S \geq \frac{1}{2}$), she already tends to stock deeply. In this case, the platform can profitably recommend its own preferred order quantity without compromise ($\eta_P^* = 1$). Aggressive sellers are less

concerned about overstocking and more receptive to assertive recommendations, so full self-centering maximizes profit. When the seller is cautious ($r_S < \frac{1}{2}$), she prioritizes avoiding overstock. Here the platform faces a trade-off: more self-centered recommendations increase profit when adopted but also raise the risk of eroding trust and discouraging adoption. The optimal stance depends on the size of the gap $r_P - r_S$. If the gap is small, the platform is better off accommodating the seller by setting $\eta_P^* = 0$, thereby sustaining trust and ensuring adoption. If the gap is large, the platform chooses an interior solution $\eta_P^* = \tilde{\eta} \in (0, 1)$, balancing the incremental profit from more assertive recommendations against the risk of discouraging the seller.

Comparative statics further refine this picture. A stronger informational advantage (higher SNR) makes self-centered recommendations more credible, allowing the platform to lean further toward its own preferred stocking level without losing adoption.

By contrast, higher profit margins (p/c) make the platform more accommodating: the larger profit "pie" reduces the incentive to extract additional gains through aggressive recommendations, and also narrows the misalignment $r_P - r_S$ (see (3.1)), making the seller's and platform's preferences naturally closer. Proposition 3.4 formalizes these monotonicities: η_P^* is increasing in SNR and decreasing in p/c.

In summary, the platform recommends fully self-centered quantities when facing aggressive sellers, fully seller-centered quantities when facing cautious sellers with aligned preferences, and an interior stance otherwise. The direction of adjustment is shaped by the environment: stronger information advantages push the platform toward its own preferred orders, while higher margins make it more willing to accommodate the seller.

Having characterized the platform's optimal recommendation stance η_P^*, we next examine how its profitability is affected by increased demand volatility. Two sources of uncertainty are relevant. The seasonal component ξ governs the strength of the platform's informational advantage, while the noise term ϵ captures idiosyncratic demand shocks that neither party can anticipate. The following result shows that greater variability along either dimension reduces the platform's long-run profit, even when it optimally chooses its recommendation stance.

Proposition 3.5. *With the profit-maximizing η_P^* as defined in* (3.18)*, the platform's long-run average profit decreases monotonically in both the variance of the seasonal component and the market randomness. Formally,*

$$\frac{\mathrm{d}\bar{\Pi}^P(\eta_P^*)}{\mathrm{d}\xi} < 0 \quad \textit{and} \quad \frac{\mathrm{d}\bar{\Pi}^P(\eta_P^*)}{\mathrm{d}\epsilon} < 0.$$

The decline in profit with increased market noise (ϵ) is intuitive: greater randomness

reduces the accuracy of both the platform's and the seller's decisions, thereby lowering expected profit regardless of the recommendation policy. More surprising is the effect of seasonal variance (ξ). A higher ξ raises the signal-to-noise ratio and strengthens the platform's informational advantage, yet the platform's long-run profit still falls.

The reason lies in incomplete adoption. Under Assumption 3.1, (3.10) together with (3.15) implies $\rho(\eta) \in (0,1)$ for all $\eta \in [0,1]$, so even at the profit-maximizing stance the seller ignores the recommendation with positive probability in steady state. The platform's steady-state profit can be written as a mixture of adoption and non-adoption payoffs:

$$\bar{\Pi}^P(\eta) = \rho(\eta)\,\mathbb{E}\left[\Pi^P\left(\bar{q}^R(\eta), d_t\right)\right] + \left(1-\rho(\eta)\right)\mathbb{E}\left[\Pi^P\left(\hat{q}^S, d_t\right)\right].$$

An increase in ξ affects the two components in the mixture asymmetrically. Conditional on adoption, the platform's recommendation is season-contingent: both demand $d_t = \mu + \xi_t + \epsilon_t$ and the adopted order $\bar{q}^R(\eta)$ shift one-for-one with ξ_t. Because the newsvendor payoff depends on the mismatch between q_t and d_t, this common shift cancels, so the adopted payoff $\mathbb{E}\left[\Pi^P(\bar{q}^R(\eta), d_t)\right]$ does not depend on ξ. In contrast, when the seller does not adopt, she orders the season-invariant benchmark $\hat{q}^S$, and larger seasonal swings make this benchmark increasingly "off-target" (too low in high seasons and too high in low seasons) thereby reducing $\mathbb{E}\left[\Pi^P(\hat{q}^S, d_t)\right]$; moreover, because the newsvendor payoff is concave, the expected loss from this seasonal mismatch grows more than proportionally in ξ. While a larger ξ improves the credibility of the recommendation and thus raises the steady-state adoption rate $\rho(\eta)$ as defined in (3.15), incomplete adoption persists for all η under Assumption 3.1. Consequently, the amplified losses in non-adoption periods dominate the trust-driven reweighting toward adoption, which explains why the platform's long-run profit decreases in ξ even at the profit-maximizing stance η_P^*.

3.5 The Case with Demand Censoring

In the baseline model, we assumed that the platform and the seller observe the full demand realization, $d_t = \mu + \xi_t + \epsilon_t$, at the end of each period. In practice, however, demand is often censored for both the platform and sellers. In our model, this censoring is immaterial for the platform's behavior, because the platform is assumed to know the demand environment and commits to a static recommendation rule that does not depend on ex post learning. By contrast, censoring directly affects the seller's ability to evaluate the performance of recommendations. In particular, when the seller's order quantity falls short of realized demand, she observes sales capped at her order quantity and only learns that $q_t \leq d_t$,

thereby making the performance of the recommendation, as modeled in (3.12) and (3.13), potentially ambiguous.

It is important to note that demand censoring does not always make the performance of a recommendation ambiguous. In some cases, the seller can still determine the sign of $\Delta_t \Pi^S$ in (3.12) even when demand is censored ($q_t \leq d_t$). For example, if both the platform's recommendation and the seller's own benchmark fall short of realized demand, and the recommendation is strictly smaller (i.e., $q_t^R < \hat{q}_S \leq d_t$), then the recommendation necessarily induces a larger stockout than the seller's decision. In this case, the recommendation is strictly worse regardless of how high the true demand is. More generally, even under censored observations, there are situations in which the seller can infer the recommendation's performance—specifically, the sign of $\Delta_t \Pi^S$—with certainty. We refer to such cases as unambiguous censored demand. The formal definition follows.

Definition 3.6. *When demand d_t is censored due to understock (i.e., $q_t \leq d_t$), the censored observation is ambiguous if $q_t = \min\{q_t^R, \hat{q}_S\}$, and unambiguous if $q_t = \max\{q_t^R, \hat{q}_S\}$.*

In other words, censored demand is ambiguous precisely when the seller chooses the smaller of the recommended order quantity and her own benchmark, so that she cannot infer whether the larger alternative would have improved her profit.

3.5.1 Bayesian-updating Sellers

In the unambiguous case, the seller can determine the sign of the recommendation's performance, and therefore updates trust in exactly the same way as under fully observed demand. For ambiguous censored demand, we consider a Bayesian seller who updates her belief from the prior distribution $F_{\xi+\epsilon}$ (as in the baseline model) to a posterior distribution $\hat{F}_{\xi+\epsilon}$ at the end of each period. By Bayes' rule, we have

$$\hat{F}_{\xi+\epsilon}(x) = \Pr\left[\xi_t + \epsilon_t < x \,|\, \xi_t + \epsilon_t \geq q_t - \mu\right] = \frac{F_{\xi+\epsilon}(x) - F_{\xi+\epsilon}(q_t - \mu)}{1 - F_{\xi+\epsilon}(q_t - \mu)}, \tag{3.19}$$

with corresponding density

$$\hat{f}_{\xi+\epsilon}(x) = \hat{F}'_{\xi+\epsilon}(x) = \frac{f_{\xi+\epsilon}(x)}{1 - F_{\xi+\epsilon}(q_t - \mu)}. \tag{3.20}$$

This Bayesian updating allows the seller to evaluate the expected profit difference between adopting the recommendation and following her own benchmark using the posterior belief

$\hat{F}_{\xi+\epsilon}$. Accordingly, the performance measure originally defined in (3.12) is updated to

$$\Delta_t\hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) := \begin{cases} \Pi^S(q_t^R, d_t) - \Pi^S(\hat{q}_S, d_t), & \text{if } q_t > d_t, \\ \mathbb{E}_{\xi_t+\epsilon_t \sim \hat{F}_{\xi+\epsilon}} \left[\Pi^S(q_t^R, \mu + \xi_t + \epsilon_t) - \Pi^S(\hat{q}_S, \mu + \xi_t + \epsilon_t)\right], & \text{if } q_t \leq d_t. \end{cases} \tag{3.21}$$

The first case in (3.21) corresponds to uncensored demand and coincides with (3.12). The second case applies when demand is censored. In unambiguous censored cases (Definition 3.6), the seller can still infer the recommendation's performance directly, whereas in ambiguous cases she must rely on posterior expectations. We establish the following result.

Lemma 3.7. *If censored demand is unambiguous, a Bayesian-updating seller correctly identifies the performance of the recommendation. Formally,*

$$\Delta_t\hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) \begin{cases} < 0, & \text{if } q_t^R < \hat{q}^S, \\ \geq 0, & \text{if } \hat{q}^S < q_t^R. \end{cases}$$

Lemma 3.7 confirms that in unambiguous cases, Bayesian evaluation coincides with the seller's intuitive judgment: the seller can determine with certainty whether the recommendation underperformed or not. Thus, the formulation in (3.21) provides a unified treatment of recommendation performance that accommodates both ambiguous and unambiguous censored observations.

With this updated performance measure under censored demand, the seller continues to update her trust according to the same exponential-smoothing rule as in the baseline model. We denote the trust level under censoring by α_t^c to distinguish it from α_t in the fully observed case. Specifically,

$$\alpha_{t+1}^c = \begin{cases} \delta \cdot 1 + (1-\delta) \cdot \alpha_t^c, & \text{if } \Delta_t\hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) > 0, \\ \alpha_t^c, & \text{if } \Delta_t\hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) = 0, \\ \delta \cdot 0 + (1-\delta) \cdot \alpha_t^c, & \text{if } \Delta_t\hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) < 0. \end{cases} \tag{3.22}$$

Thus, censoring affects trust dynamics only through the seller's evaluation of recommendation performance, while the trust-update mechanism itself remains unchanged.

We now examine how a Bayesian seller evaluates a recommendation under ambiguous censored demand, a phenomenon that does not arise when demand is fully observed.

Lemma 3.8. *When censored demand is ambiguous, there exist two thresholds, $\underline{r}_S$ and $\bar{r}_P$, such that a Bayesian-updating seller increases her trust, i.e., $\Delta_t\hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) > 0$, if and*

only if the realized demand state is high season ($\xi_t = \xi$) and at least one of the following conditions holds:

1. $r_S > \underline{r}_S$ *or* $r_P < \bar{r}_P$*; or*

2. $\eta < \bar{\eta}^c$*, where* $\bar{\eta}^c$ *is defined implicitly by*

$$\mathbb{E}_{\xi_t+\epsilon_t \sim \hat{F}_{\xi+\epsilon}}\Big[\Pi^S(\bar{q}_t^R(\bar{\eta}^c), \mu + \xi_t + \epsilon_t) - \Pi^S(\hat{q}_S, \mu + \xi_t + \epsilon_t)\Big] = 0.$$

In all other cases, the seller decreases her trust, i.e., $\Delta_t \hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) < 0$.

Lemma 3.8 highlights a sharp departure from the baseline model. Without censoring, the seller observes realized demand d_t and can evaluate ex post whether the recommendation outperformed her own benchmark. Under ambiguous censored demand (Definition 3.6), however, the seller only knows that the order quantity she chose—the smaller of the recommendation and her own estimate, $q_t = \min\{q_t^R, \hat{q}_S\}$—lies below true demand. She cannot directly infer whether the larger alternative $\max\{q_t^R, \hat{q}_S\}$ would have improved her profit, since increasing the order could have captured additional unmet demand but might also have generated costly overstock.

Whether the recommendation appears beneficial therefore depends on the magnitude of the gap between the two candidate order quantities,

$$\max\{q_t^R, \hat{q}_S\} - \min\{q_t^R, \hat{q}_S\} = |q_t^R - \hat{q}_S| = \Big|2\epsilon\eta(r_P - r_S) + \xi_t\Big|, \tag{3.23}$$

which captures the combined effect of interest misalignment and seasonality. When this gap is "small enough," the larger alternative appears safe and is evaluated favorably; when it is large, the same alternative is perceived as overly aggressive and likely to backfire through overstock.

Lemma 3.8 characterizes when the gap in (3.23) is "small enough" for the recommendation to be evaluated favorably. The key distinction is between low- and high-season realizations.

1. Low season ($\xi_t = -\xi$). In a low season, the gap in (3.23) becomes

$$|q_t^R - \hat{q}_S| = \Big|2\epsilon\eta(r_P - r_S) - \xi\Big|.$$

   This gap is always "small enough." Intuitively, the negative seasonal shock lowers the platform's recommendation relative to the seller's benchmark, reflecting demand weakness that the seller does not observe. This downward adjustment partially offsets the

upward force created by interest misalignment ($r_P > r_S$), which pushes the recommendation above $\hat{q}_S$ through the term $2\epsilon\eta(r_P - r_S)$. As a result, these opposing forces counterbalance each other in low seasons, keeping the gap bounded and making the larger alternative order appear safe.

2. High season ($\xi_t = \xi$). In a high season, the gap becomes

$$|q_t^R - \hat{q}_S| = \left|2\epsilon\eta(r_P - r_S) + \xi\right|,$$

which need not be small. Here, the seasonal effect and interest misalignment reinforce each other, both pushing the recommendation upward. Whether the larger alternative order appears safe therefore depends on the extent of this reinforcement. As formalized in Lemma 3.8, the gap remains acceptable only when the interest misalignment $r_P - r_S$ is sufficiently small (condition (i)), or when the platform adopts a sufficiently benevolent stance (condition (ii), $\eta < \bar{\eta}^c$). Otherwise, the recommendation is perceived as excessively aggressive and evaluated unfavorably.

Fix a trust level α_t^c and a degree of self-centering η. With probability α_t^c, the seller adopts the recommendation and sets $q_t = q_t^R(\eta)$; with probability $1 - \alpha_t^c$, she orders her uninformed benchmark $q_t = \hat{q}^S$. The probability that the seller increases her trust by the end of period t can therefore be written as an ex ante probability, taken with respect to both the adoption decision and the demand realization, that the posterior performance signal as defined in (3.21) is positive:

$$\rho^c(\alpha_t^c, \eta) := \Pr\left(\Delta_t \hat{\Pi}^S\left(q_t, q_t^R(\eta), \hat{q}^S\right) > 0\right). \tag{3.24}$$

This probability differs fundamentally from its counterpart in the baseline model. When demand is fully observed, the trust-increase probability $\rho(\eta)$ (Proposition 3.2) equals the probability that the recommendation outperforms the seller's benchmark, an event that is independent of whether the recommendation is actually adopted. Consequently, $\rho(\eta)$ is independent of the trust level α_t and stationary over time. Under censored demand, by contrast, adoption and evaluation are endogenously linked. Whether the seller follows the recommendation affects whether demand is censored and, if so, whether the resulting observation is ambiguous. These features in turn shape the seller's posterior assessment of recommendation performance. As a result, the probability of a trust increase depends explicitly on the current adoption probability α_t^c, which therefore enters directly into $\rho^c(\alpha_t^c, \eta)$.

Lemma 3.9. *For any fixed η, the trust-increase probability $\rho^c(\alpha_t^c, \eta)$ is decreasing in α_t^c when $\eta < \bar{\eta}^c$, and increasing in α_t^c when $\eta \geq \bar{\eta}^c$. Moreover, for any fixed α_t^c, $\rho^c(\alpha_t^c, \eta)$ is piecewise increasing in η, increasing on each of the intervals $[0, \bar{\eta}^c]$ and $(\bar{\eta}^c, 1]$.*

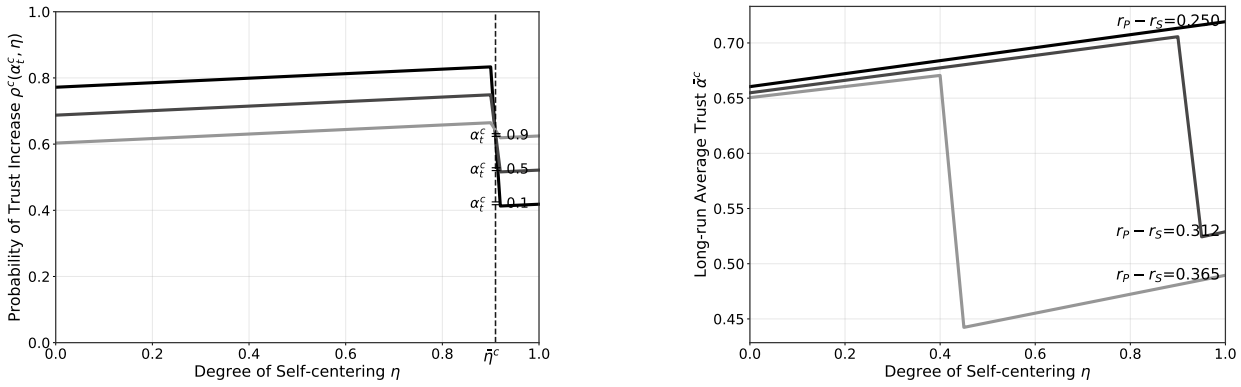

(a) Probability of trust increase $\rho^c(\alpha_t^c, \eta)$

(b) Long-run average trust $\bar{\alpha}_c$

Figure 3.4: Probability that the seller increases trust, $\rho^c(\alpha, \eta)$, and the seller's long-run average trust, $\bar{\alpha}_c$, as functions of the DoS η.

Recall that $\bar{\eta}^c$ denotes the DoS threshold at which a Bayesian seller, facing ambiguous high-season demand, is indifferent between following the platform's recommendation and relying on her own benchmark (Lemma 3.8). The lemma shows that ambiguous high-season demand is interpreted as favorable to the recommendation in two distinct situations: (i) when interest misalignment is sufficiently small, or (ii) when misalignment is larger but the platform is not too self-centering ($\eta < \bar{\eta}^c$). In the first case, ambiguous high-season demand is always interpreted as supporting the platform, regardless of η, so the threshold $\bar{\eta}^c$ is effectively inactive. In the second case, $\bar{\eta}^c$ becomes a genuine tipping point: when $\eta < \bar{\eta}^c$, the seller views ambiguous high-season realizations as evidence in favor of the recommendation, whereas for $\eta \geq \bar{\eta}^c$ the same observations are interpreted as evidence against it. Lemma 3.9 shows that the trust-increase probability $\rho^c(\alpha_t^c, \eta)$ inherits this threshold structure exactly. The left panel of Figure 3.4 visualizes the resulting pattern.

Two features of this relationship are particularly noteworthy.

1. Threshold-induced discontinuity. Under ambiguous high-season demand, the seller cannot observe the exact demand realization and must rely on Bayesian inference to compare the larger and smaller candidate order quantities, $\max\{q_t^R, \hat{q}^S\}$ and $\min\{q_t^R, \hat{q}^S\}$. When the DoS crosses $\eta = \bar{\eta}^c$, this Bayesian evaluation switches sign: all ambiguous high-season realizations that were previously treated as good news (trust-increasing) are reclassified as bad news (trust-decreasing). Because this reclassification applies to the entire set of ambiguous high-season realizations at once, $\rho^c(\alpha_t^c, \eta)$ exhibits a discrete downward jump at $\eta = \bar{\eta}^c$. This is in sharp contrast to the smooth response under the fully observed demand case.

2. Piecewise monotonicity and the informational shield. Aside from the discontinuity at $\eta = \bar{\eta}^c$, the function $\rho^c(\alpha_t^c, \eta)$ is increasing in η on both sides of the threshold. A higher η makes the platform more self-centering and thus more aggressive in its recommendations. Under full demand observability, Proposition 3.2 highlights a fundamental trade-off: greater aggressiveness mitigates understock in high seasons (raising trust) but exacerbates overstock in low seasons (lowering trust), so the net effect depends on the seller's service level. Under demand censoring, however, the seller can always detect stockouts (understock), while in ambiguous states she may not be able to infer whether a more aggressive recommendation would have overshot demand. Thus, the downside overstock channel is partly masked, whereas the upside from avoiding stockouts remains salient. This asymmetric visibility creates an informational shield that makes more aggressive recommendations appear safer than they truly are, generating the piecewise increasing pattern shown in Figure 3.4.

Lemma 3.9 also characterizes how the trust-increase probability responds to the current trust level α_t^c. The key driver is again the seller's behavior under ambiguous high-season demand (in low seasons, ambiguous demand always leads to trust decreases and therefore does not affect the comparative statics). Such ambiguous high-season events arise only when the seller rejects the recommendation and subsequently stocks out. When $\eta \geq \bar{\eta}^c$, these ambiguous realizations are interpreted unfavorably and reduce trust. A higher adoption rate α_t^c therefore reduces the frequency of these negative events, increasing the overall probability of a trust increase. Consequently, $\rho^c(\alpha_t^c, \eta)$ is increasing in α_t^c for $\eta \geq \bar{\eta}^c$. In contrast, when $\eta < \bar{\eta}^c$, ambiguous high-season demand is interpreted favorably and leads to trust increases. A higher adoption rate then reduces the number of such favorable rejections, lowering the likelihood of positive feedback. As a result, $\rho^c(\alpha_t^c, \eta)$ is decreasing in α_t^c for $\eta < \bar{\eta}^c$. The left panel of Figure 3.4 illustrates these opposing effects of trust on the probability of trust increase under censored demand.

3.5.2 Steady-state Trust Level

Analogous to the baseline model, we define the steady-state trust under demand censoring as

$$\bar{\alpha}^c = \lim_{t \to \infty} \mathbb{E}[\alpha_t^c]. \tag{3.25}$$

The next result characterizes how demand censoring modifies the relationship between the platform's degree of self-centering (DoS) and steady-state trust.

Proposition 3.10. *With censored demand, the steady-state trust in* (3.25) *satisfies the fixed-point condition* $\bar{\alpha}^c = \rho^c(\bar{\alpha}^c, \eta)$. *If* $r_S < \underline{r}_S$ *and* $r_P > \bar{r}_P$, *then*

$$\bar{\alpha}^c = \begin{cases} \dfrac{2(\xi+\epsilon) - r_S\Big(2\eta\epsilon(r_S - r_P) + \xi\Big)}{2(\xi - \xi r_S + 2\epsilon)}, & \text{if } \eta < \bar{\eta}^c, \\ \dfrac{\xi - r_S\Big(\xi + 2\eta\epsilon(r_S - r_P) - 2\epsilon\Big)}{\xi + 2r_S(\epsilon - \xi) + 2\epsilon}, & \text{if } \eta \geq \bar{\eta}^c, \end{cases} \tag{3.26}$$

which is piecewise increasing in η. *Otherwise, if* $r_S > \underline{r}_S$ *or* $r_P < \bar{r}_P$, *then*

$$\bar{\alpha}^c = \frac{2(\xi+\epsilon) - r_S\Big(2\eta\epsilon(r_S - r_P) + \xi\Big)}{2(\xi - \xi r_S + 2\epsilon)}, \tag{3.27}$$

which increases smoothly in η. *In both cases,* $\bar{\alpha}^c$ *increases in the signal-to-noise ratio (SNR),* ξ/ϵ.

Proposition 3.10 highlights a sharp contrast between censored and fully observable demand. In the baseline model (Proposition 3.2), the effect of the platform's DoS on trust depends on the seller's service level and may even be negative when the seller is highly averse to overstocking. Under demand censoring, by contrast, steady-state trust $\bar{\alpha}^c$ is always increasing in η, either smoothly or in a piecewise manner. This follows directly from the fixed-point condition $\bar{\alpha}^c = \rho^c(\bar{\alpha}^c, \eta)$ and the monotonicity properties of ρ^c in η established in Lemma 3.9. The underlying mechanism is the informational shield: censoring masks negative evidence from overstock, which would otherwise erode trust, while preserving positive evidence from avoided stockouts. As a result, more self-centered recommendations systematically raise steady-state trust in the censored-demand environment.

The precise shape of $\bar{\alpha}^c$—piecewise as in (3.26) or globally monotone as in (3.27)—is inherited from the behavior of ρ^c and depends on the magnitude of interest misalignment, which governs how a Bayesian seller interprets ambiguous censored demand (Lemma 3.8). When misalignment is small ($r_S > \underline{r}_S$ or $r_P < \bar{r}_P$), the threshold $\bar{\eta}^c$ is never reached. Ambiguous demand is then interpreted favorably for all η, and the trust-increase probability ρ^c rises smoothly with η, yielding the globally monotone expression in (3.27).

When misalignment is large ($r_S < \underline{r}_S$ and $r_P > \bar{r}_P$), the threshold $\bar{\eta}^c$ becomes operative. As shown in Lemma 3.9, ρ^c then exhibits a threshold-induced discontinuity: a discrete downward jump at $\eta = \bar{\eta}^c$ reflecting a sudden reversal in how the seller interprets ambiguous demand. The steady-state trust $\bar{\alpha}^c$ inherits this feature through the fixed-point condition, resulting in the piecewise increasing structure in (3.26). The right panel of Figure 3.4 compares the behavior of $\bar{\alpha}^c$ under different interest misalignment.

Finally, in both (3.26) and (3.27), steady-state trust increases with the SNR ξ/ϵ. A higher SNR strengthens the platform's informational advantage, making its recommendations more likely to outperform the seller's uninformed benchmark, thereby supporting stronger long-run trust.

3.5.3 Platform's Optimal Strategy

Before turning to the platform's optimization problem, we note that the win–win property in Proposition 3.3 carries over to the censored-demand setting. In this extension, demand censoring affects the seller's end-of-period information and thus her trust updates, but it does not alter realized payoffs conditional on an order quantity. Consequently, for any fixed η, the expected profit comparison between following the recommendation $q_t^R(\eta)$ and ordering the uninformed benchmark $\hat{q}^S$ is identical to that in the baseline model.

Let $\bar{\alpha}^{\,c}(\eta)$ denote the steady-state adoption rate under censored demand. Because demand is i.i.d. over time and independent of the trust state, the seller's ($i = S$) and platform's ($i = P$) long-run average profit can be written as

$$\bar{\Pi}^{i,c}(\eta) = \mathbb{E}\left[\Pi^i(\hat{q}^S, d_t)\right] + \bar{\alpha}^{\,c}(\eta)\Big(\mathbb{E}\left[\Pi^i(q_t^R(\eta), d_t)\right] - \mathbb{E}\left[\Pi^i(\hat{q}^S, d_t)\right]\Big). \tag{3.28}$$

Proposition 3.3 implies that the bracketed term is weakly positive for both parties for all $\eta \in [0, 1]$. Therefore, even when demand is censored, information sharing through inventory recommendations continues to increase profits for both the seller and the platform, yielding a win-win outcome relative to the benchmark without recommendations; that is, $\bar{\Pi}^{S,c}(\eta) > \Pi_0^S$ and $\bar{\Pi}^{P,c}(\eta) > \Pi_0^P$ (recall that Π_0^S and Π_0^P are defined in Proposition 3.3).

We now examine the platform's optimal strategy under censored demand. The profit-maximizing degree of self-centering (DoS), denoted by η_c^*, is characterized as follows.

Proposition 3.11. *Under censored demand, the platform's profit-maximizing DoS satisfies* $\eta_c^* \in \{\bar{\eta}^c, 1\}$, *and* $\eta_c^* = \bar{\eta}^c$ *only if* $r_S < \underline{r}_S$ *and* $r_P > \bar{r}_P$. *Moreover,* $\bar{\eta}^c \in (0, 1)$ *is increasing in* r_S *and decreasing in* r_P.

Proposition 3.11 highlights a sharp departure from the uncensored setting (Proposition 3.4). Under censored demand, the platform's profit-maximizing strategy is almost always highly aggressive and often coincides with full self-centering ($\eta_c^* = 1$). This outcome follows directly from the trust dynamics characterized in Proposition 3.10. The platform's long-run profit has two components: (i) profit conditional on adoption and (ii) the adoption probability, captured by steady-state trust. A higher DoS increases profit conditional on adoption by shifting recommendations closer to the platform's own optimal quantity q_t^P. At

the same time, due to the informational shield created by censoring, a higher DoS generally raises steady-state trust $\bar{\alpha}^c$. With both components increasing in η, the platform's incentives are aligned, pushing it toward a highly self-centered stance.

This stands in stark contrast to the baseline model, where the platform often must compromise substantially—sometimes choosing $\eta^* = 0$—to build trust, especially when facing cautious sellers. Under censored demand, such compromise is largely unnecessary. The only exception arises under substantial interest misalignment ($r_S < \underline{r}_S$ and $r_P > \bar{r}_P$), where steady-state trust exhibits a discrete downward jump at the threshold $\bar{\eta}^c$ (Proposition 3.10). This discontinuity creates a sharp trade-off for the platform: increasing η beyond $\bar{\eta}^c$ raises profit conditional on adoption but simultaneously causes a sudden drop in the adoption rate. When this trust loss outweighs the marginal gain in per-adoption profit, the platform optimally moderates its stance and selects $\eta_c^* = \bar{\eta}^c$, attaining the highest profit just before the discontinuity rather than at $\eta = 1$.

Finally, note that the threshold $\bar{\eta}^c$ itself increases as interest misalignment diminishes—that is, as r_S rises or r_P falls. As the platform's and seller's objectives become more closely aligned, the platform can adopt a more aggressive recommendation stance without triggering a loss of trust, and thus needs to compromise less.

3.5.4 Comparison with the Baseline Optimal Profit

We now compare the platform's maximal attainable profit under censored demand with that under full demand observability in the baseline model.

Proposition 3.12. *Suppose that $r_S > \underline{r}_S$ or $r_P < \bar{r}_P$. There exists a threshold $\bar{\sigma}$ such that, when the signal-to-noise ratio satisfies $\xi/\epsilon < \bar{\sigma}$, the platform's optimal profit under censored demand strictly exceeds its optimal profit in the baseline model with fully observed demand. Moreover, the threshold $\bar{\sigma}$ is increasing in r_P and quasi-concave in r_S.*

Proposition 3.12 delivers a counterintuitive insight: the platform may earn higher long-run profits when demand is censored than when demand is fully observed. In other words, reduced information transparency can sometimes be profit-enhancing. This occurs when interest misalignment is limited ($r_S > \underline{r}_S$ or $r_P < \bar{r}_P$) and the market environment is sufficiently unpredictable for both the platform and the seller (i.e., the signal-to-noise ratio is below the threshold $\bar{\sigma}$). The underlying mechanism reflects the interaction between the platform's optimal recommendation stance and the resulting steady-state trust.

1. More aggressive strategy under limited interest misalignment ($r_S \geq \underline{r}_S$ or $r_P < \bar{r}_P$). Proposition 3.11 shows that, when interest misalignment is limited, demand censoring

allows the platform to adopt a fully self-centered recommendation strategy ($\eta_c^* = 1$). This stance is typically more aggressive than the profit-maximizing strategy in the baseline model ($\eta_c^* \geq \eta^*$), and therefore yields higher profits whenever recommendations are adopted.

2. Higher relative trust in low-SNR environments ($\xi/\epsilon < \bar{\sigma}$). For increased aggressiveness to translate into higher profit, recommendations must be adopted sufficiently often. The informational shield (Proposition 3.10) raises trust under censoring by masking recommendation errors. This effect is strongest when the signal-to-noise ratio is low: in this regime, the platform's signal is noisy and would generate many visible mistakes under full transparency, whereas censoring hides precisely those adverse realizations that would otherwise erode trust. Consequently, adoption probabilities increase much more under censored demand than under full observability. As the SNR rises, recommendations become objectively more accurate and are adopted more frequently, so ambiguous stockouts occur less often and fewer mistakes remain hidden. The incremental trust benefit from censoring therefore diminishes, causing steady-state trust under censored demand to increase concavely with SNR. By contrast, under fully observed demand, steady-state trust increases approximately linearly with SNR. As a result, the range of SNR values for which trust is higher under censored demand is necessarily bounded.

To summarize, when the signal-to-noise ratio is small, demand censoring induces a higher steady-state trust level than full observability ($\bar{\alpha}^c > \bar{\alpha}$). Combined with a more aggressive recommendation strategy under limited interest misalignment ($\eta_c^* \geq \eta^*$), this elevated trust translates into higher long-run profits for the platform in unpredictable market environments.

The threshold $\bar{\sigma}$ delineates the range of market conditions under which demand censoring is beneficial. It increases in r_P, since higher platform service levels lead to more aggressive recommendations and amplify the value of the informational shield by masking potential overstock losses. By contrast, its relationship with r_S is quasi-concave: a higher seller service level raises the recommended order quantity (strengthening the shield), but also increases the seller's default order $\hat{q}_S$, which reduces the frequency of stockouts and limits the scope for censoring to operate.

Finally, the potential benefits of demand censoring are not confined to the platform alone. Using numerical analysis over a wide range of parameter combinations (Section 3.6.2), we show that, despite the platform adopting a more aggressive recommendation stance, the seller experiences a larger average increase in long-run profit than the platform. Moreover, in the majority of cases, demand censoring generates a win–win outcome for both the platform and

the seller.

3.6 Numerical Analysis

In this numerical section, we complement the analytical results along three dimensions. First, we quantify the coordination surplus generated by inventory recommendations and characterize its distribution and allocation between the platform and the seller. Second, we conduct comparative statics over a broad parameter space and identify when demand censoring, beyond increasing platform profit as established in Proposition 3.12, also benefits the seller, thereby generating win–win outcomes. Third, we examine robustness under more general trust evolution rules that capture conservative and cautious seller learning, and study how such behaviors discipline the platform's optimal degree of self-centering.

3.6.1 Numerical Setup

Parameter generation and filtering. We generate model instances at random to span a broad range of economically relevant environments. Each instance is characterized by six primitive parameters: the price $p \in [1, 30]$, the cost-to-price ratio $c/p \in [0, 1]$, the platform cost share $\beta \in [0.01, 0.4]$, the platform revenue share $\gamma \in [0.01, 0.4]$, the noise parameter $\epsilon \in [0, 1]$, and the signal-to-noise ratio $\xi/\epsilon \in [0.01, 0.98]$. We fix the baseline demand at $\mu = 2.5$. All parameters are drawn independently and uniformly from their respective ranges, subject to the structural constraint $\beta < \gamma$. We then impose Assumption 3.1 and retain only those instances that satisfy all structural and economic feasibility conditions. Out of 10,000 random draws, 2,408 instances remain after filtering.

Computation and reported outcomes. For each retained instance, we compute equilibrium outcomes under both information regimes: full demand observability (the baseline model) and censored demand. We record the platform's optimal degree of self-centering (DoS), computed over a grid $\eta \in \{0, 0.01, \ldots, 1\}$, along with the corresponding long-run average trust level and the long-run average profits of both the platform and the seller. We also report the changes in these outcomes when moving from full observability to censored demand. Summary statistics for all outcome variables are presented in Table 3.2.

To reduce simulation noise and improve computational efficiency, we compute long-run averages directly from the theoretical steady-state expressions whenever closed-form formulas are available. Specifically, long-run average trust in the baseline regime is computed using equation (3.15), long-run average trust under censored demand is obtained from Proposition 3.10, and long-run average profits are computed using equation (B.12). We validate

these analytical steady-state values against Monte Carlo simulations (Figure 3.5) and rely on simulation-based estimates only in Section 6.4, where alternative trust-update rules preclude closed-form solutions.

Figure 3.5: Convergence of long-run average trust, the platform's profit, and the seller's profit in the baseline demand model.

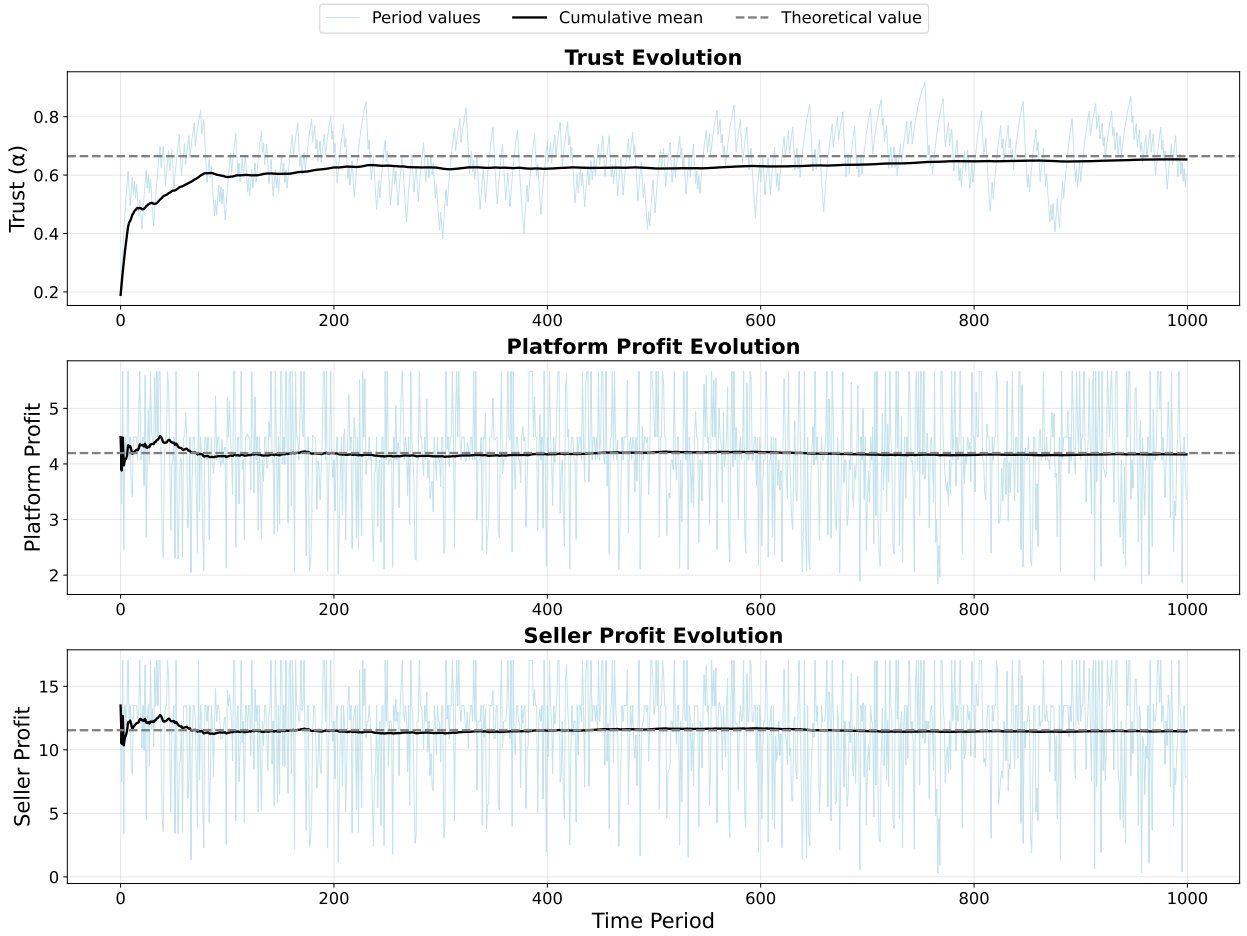

3.6.2 Supply Chain Coordination and the Value of Information

To quantify the value of inventory recommendations, we compare equilibrium profits under the optimal degree of self-centering (DoS) with a benchmark in which no recommendations are provided. In the benchmark, the seller orders her perceived optimal quantity $\hat{q}_S$, yielding profits $\Pi_0^P = \bar{\Pi}^P(\hat{q}_S)$ for the platform and $\Pi_0^S = \bar{\Pi}^S(\hat{q}_S)$ for the seller. We report absolute and percentage profit gains for each party and for the overall supply chain, as well as the platform's share of the total supply-chain gain, in Table 3.3.

Table 3.2: Summary statistics of outcome variables (2,408 parameter combinations)

Variable	Min	Mean	Max	Std	# Positive	#Negative
η^*	0.000	0.951	1.000	0.140		
η_c^*	0.030	0.967	1.000	0.098		
$\bar{\alpha}$	0.503	0.612	0.747	0.059		
$\bar{\alpha}^c$	0.491	0.618	0.783	0.056		
$\bar{\Pi}^{P^*}$	0.115	7.310	28.001	5.229		
$\bar{\Pi}_c^{P^*}$	0.116	7.311	28.001	5.230		
$\bar{\Pi}^{S^*}$	0.294	13.701	56.954	9.795		
$\bar{\Pi}_c^{S^*}$	0.294	13.702	56.954	9.795		
$\Delta\eta^*$	-0.970	0.020	1.000	0.147	686	498
$\Delta\bar{\alpha}$	-0.042	0.007	0.071	0.019	1407	1001
$\Delta\bar{\Pi}^{P^*}$	-0.046	0.001	0.098	0.008	1395	1013
$\Delta\bar{\Pi}^{S^*}$	-0.137	0.001	0.436	0.024	1266	1142

Observation 1. *(1) The optimal recommendation can increase total supply-chain profit by up to 20%, which we refer to as the benefit of optimal recommendation. (2) On average, the seller captures 56% of this benefit, while the platform captures the remaining 44%. (3) In extreme cases, either the seller or the platform can capture nearly all of the coordination benefit, with one party's share as high as 99%.*

Figure 3.6: Distribution of percentage gains from adopting the optimal recommendation for the platform, the seller, and the supply chain

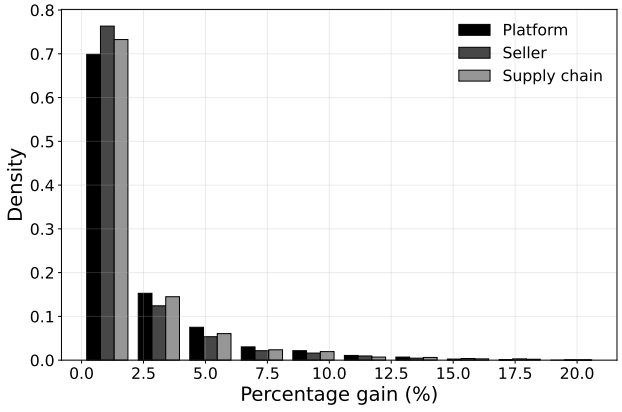

Overall, the optimal recommendation increases total supply-chain profit by an average of 1.76% in the baseline regime (standard deviation 2.61%), with outcomes ranging from essentially zero to 19.52% (Table 3.3). The corresponding mean percentage profit increases are 1.95% for the platform and 1.63% for the seller. These summary statistics are nearly unchanged under demand censoring.

Importantly, the gains from recommendations exhibit substantial heterogeneity across environments. Figure 3.6 shows that the distribution of percentage profit improvements is right-skewed, with a long upper tail of high-impact instances. While many environments exhibit relatively small gains, a nontrivial fraction experience economically large improvements, consistent with the comparative statics identified in the analytical and numerical results.

While the platform, the seller, and the supply chain experience similar average percentage profit increases under the optimal recommendation, the allocation of the resulting surplus varies substantially across parameter combinations. In some environments, nearly all of the coordination benefit accrues to a single party. We examine this heterogeneity in surplus allocation in the next observation.

Table 3.3: Benefits of the optimal recommendation for the platform, the seller, and the supply chain (baseline vs. censored demand)

Variable	Min	Mean	Max	Std
$\Delta\Pi^P = \bar{\Pi}^P(\eta^*) - \bar{\Pi}_0^P$	≈ 0	0.120/0.121	1.824/1.911	0.193/0.195
$\Delta\Pi^P/\bar{\Pi}_0^P$	≈ 0	1.948%/1.957%	17.813%/17.055%	2.639/2.624
$\Delta\Pi^S = \bar{\Pi}^S(\eta^*) - \bar{\Pi}_0^S$	≈ 0	0.164/0.165	3.377/3.249	0.278/0.273
$\Delta\Pi^S/\bar{\Pi}_0^S$	≈ 0	1.630%/1.643%	20.800%/19.886%	2.637/2.602
$\Delta\Pi^{SC} = \Delta\Pi^P + \Delta\Pi^S$	≈ 0	0.284/0.285	4.360/4.195	0.438/0.437
$\Delta\Pi^{SC}/(\Pi_0^P + \Pi_0^S)$	≈ 0	1.758%/1.767%	19.522%/18.675%	2.606/2.579
$\Delta\Pi^P/\Delta\Pi^{sc}$	1.877%	44.456%/44.339%	99.635%	17.568/17.444
$\Delta\Pi^S/\Delta\Pi^{sc}$	0.365%	55.544%/55.661%	98.123%	17.568/17.444

Observation 2. *The total supply-chain benefit from the optimal recommendation is larger when interest misalignment is greater or when the signal-to-noise ratio is higher. By contrast, the platform captures a larger share of this benefit when the signal-to-noise ratio is lower.*

Figure 3.7 plots the percentage increase in total supply-chain profit under the optimal recommendation and the platform's share of this gain as functions of interest misalignment, the seasonal component, and market noise.

Greater interest misalignment increases the value of coordination, while a higher signal-to-noise ratio (SNR) raises the value of the platform's recommendation. Consequently, both factors amplify the total supply-chain benefit from the optimal recommendation. At the same time, a higher SNR makes the recommendation more informative to the seller. By improving the seller's stocking decision, this increased informativeness shifts a larger share of the incremental supply-chain profit toward the seller and correspondingly reduces the platform's share.

Figure 3.7: Sensitivity analysis of (i) the percentage increase in supply-chain profit under the optimal recommendation (benefit of optimal recommendation) and (ii) the platform's share of that benefit, as functions of interest misalignment, the seasonal component, and the market-noise parameter

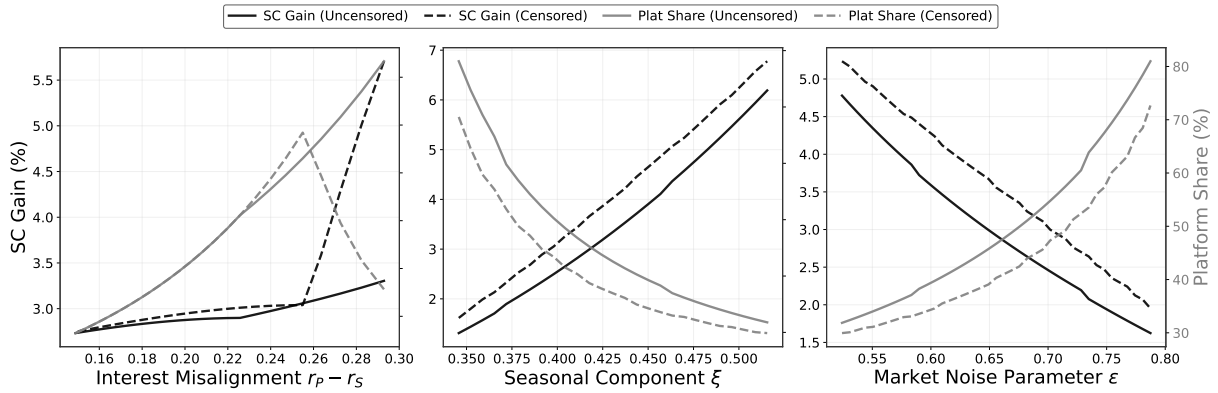

Despite capturing a smaller share of the coordination benefit, a higher signal-to-noise ratio still leads to a larger absolute (and percentage) increase in the platform's long-run average profit. In other words, the positive effect of SNR on total supply-chain surplus dominates the redistribution of that surplus toward the seller. This pattern is illustrated by the high-gain region in the heatmap of the platform's percentage profit increase as a function of interest misalignment and SNR shown in the middle panel of Figure 3.8. (In Figure 3.8, we compute $r_P - r_S$ and the signal-to-noise ratio ξ/ϵ for each parameter combination and discretize the resulting space into n^2 evenly sized bins. Within each bin, we report the average value of the outcome of interest.)

Figure 3.8: Average percentage increase in supply-chain profit, average percentage increase in the platform's profit, and the platform's share of the gain under different interest misalignment levels and signal-to-noise ratios (SNR) ($n = 100$).

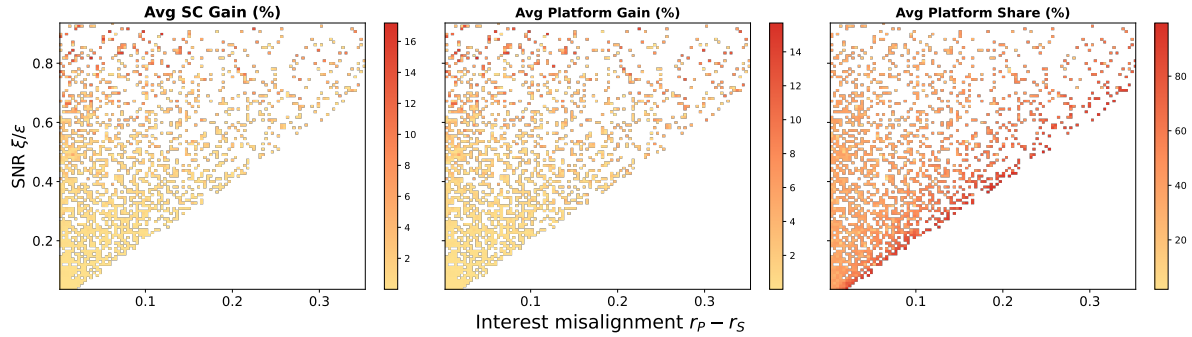

3.6.3 Impact of Demand Censoring and Win–Win Scenarios

We next compare the baseline regime with full demand observability to the censored-demand regime. Our focus is twofold: (i) whether demand censoring alters the value of the platform's optimal recommendation policy, and (ii) when censoring gives rise to win–win outcomes for both the platform and the seller.

Observation 3. *When the platform optimally chooses the degree of self-centering, demand censoring does not materially reduce the benefits of inventory recommendations (relative to the no-recommendation benchmark).*

This robustness result is reassuring. Although demand censoring reduces the seller's ex post observability of realized demand following stockouts, it does not, by itself, impair the platform's ability to translate its informational advantage into coordination value through inventory recommendations, provided the platform adjusts its recommendation stance optimally.

Recall that the benefit of a recommendation is defined as the increase in long-run average profit for the platform and the seller under the optimal recommendation policy. As reported in Table 3.2 and Table 3.3, when the platform sets the DoS optimally, both parties' long-run average profits remain nearly unchanged after demand is censored, with an average change of approximately 0.001 for each party. This pattern persists even when the parameter ranges are expanded. In addition to the fact that both the platform and the seller are sophisticated decision makers (as discussed in Observation 1), two additional forces help explain why censoring has only a limited effect on the benefits of the optimal recommendation.

First, even under demand censoring, the seller can infer latent demand using Bayesian updating, which provides a reasonably accurate assessment of recommendation performance. Moreover, demand is not always censored: when the seller does not stock out, she directly observes realized demand and can evaluate the recommendation without relying on posterior inference.

Second, the platform's optimization of the DoS mitigates potential losses from censoring. As shown in Table 3.2, the platform responds to censoring by adopting a more aggressive recommendation stance (the average $\Delta\eta^*$ is positive), reflecting the masking effect discussed in Section 3.5. This increased aggressiveness raises seller trust in 58% of the parameter combinations and, in turn, increases the likelihood that the recommendation is adopted. On average, demand censoring therefore slightly increases—rather than decreases—the benefits of recommendations.

Finally, although demand censoring has little effect on the overall benefits of the optimal recommendation, its impact on the level of the seller's long-run average profit is more pro-

nounced than its impact on the platform's. This difference is reflected in the wider range and larger standard deviation of $\Delta\bar{\Pi}^{S^*}$ relative to $\Delta\bar{\Pi}^{P^*}$ reported in Table 3.2.

Observation 4. *(1) The increase in the platform's long-run average profit induced by demand censoring is larger when interest misalignment ($r_P - r_S$) is small or moderate and when the signal-to-noise ratio (ξ/ϵ) is low. (2) The increase in the seller's long-run average profit induced by demand censoring is larger when interest misalignment is either small or large and when the signal-to-noise ratio is moderate.*

Across the 2,408 parameter combinations considered, the platform's equilibrium long-run average profit increases in 1,395 cases (57.9%) after demand is censored, while the seller's equilibrium long-run average profit increases in 1,266 cases (52.5%). Classifying outcomes by platform–seller effects, win–win outcomes account for 50.4% (1,213) of the samples, win–lose outcomes for 7.6% (182), lose–win outcomes for 2.2% (53), and lose–lose outcomes for 40.0% (960).

Figure 3.9: Changes in the platform's and seller's long-run average profits under demand censoring, as functions of interest misalignment, the seasonal component, and the market-noise parameter

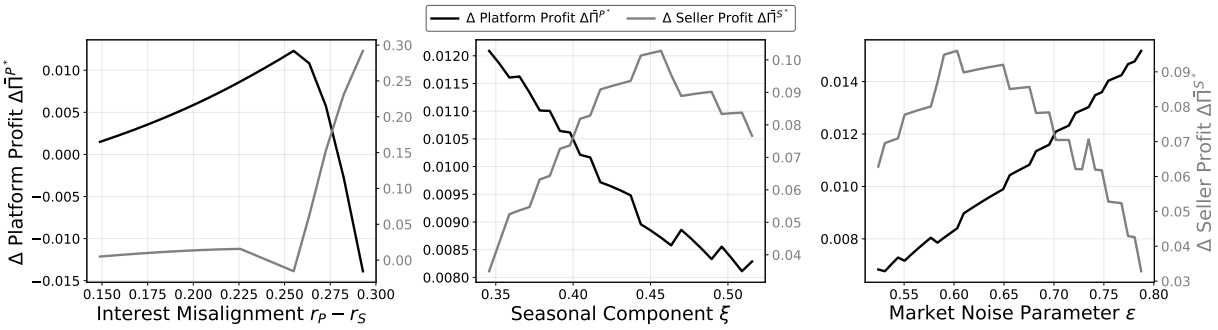

Figure 3.9 summarizes how profit changes under demand censoring vary with interest misalignment, the seasonal component, and market noise. Figure 3.10 reports, for each bin of $(r_P - r_S, \xi/\epsilon)$, the empirical probability that censoring increases the platform's profit, the seller's profit, and both profits (win–win), using a resolution of $n = 100$. We now explain the patterns summarized in Observation 4.

- Interest misalignment $(r_P - r_S)$. Figure 3.9 shows that the platform's profit gain from demand censoring initially increases with interest misalignment and then declines sharply. The decline occurs once the conditions $r_S > \underline{r}_S$ and $r_P < \bar{r}_P$ in Proposition 3.10 cease to hold. When misalignment becomes large, overly aggressive recommendations weaken the masking benefit of censoring and can cause the seller to

Figure 3.10: Probability that equilibrium long-run average profit increases under demand censoring. Resolution: $n = 100$

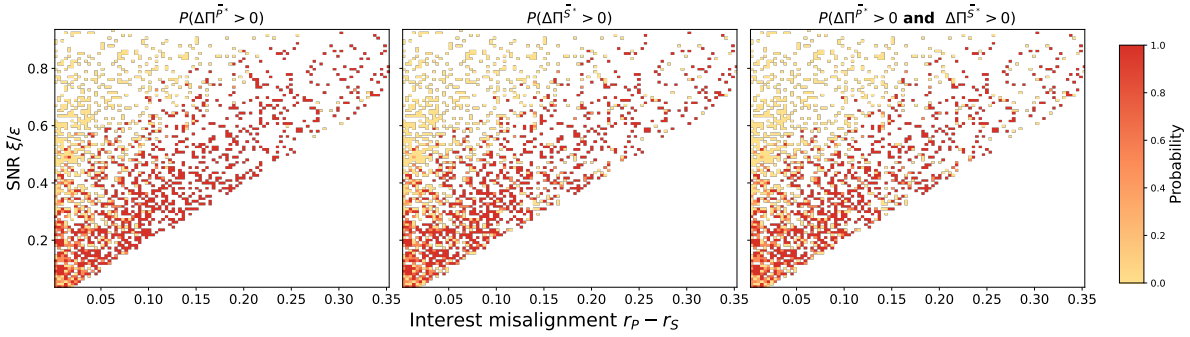

systematically distrust the platform following ambiguous demand realizations. Anticipating this response, the platform optimally reduces its degree of self-centering (η_c^*) as misalignment increases. As the platform becomes more benevolent, its incremental gain from censoring diminishes, while a larger share of the resulting profit improvement accrues to the seller.

- Signal-to-noise ratio (ξ/ϵ). Figure 3.9 indicates that the platform's profit gain from censoring decreases monotonically with the signal-to-noise ratio, and Figure 3.10 confirms that the platform is most likely to benefit from censoring when SNR is low. This pattern extends Proposition 3.12 to parameter regions in which $r_S > \underline{r}_S$ and $r_P < \bar{r}_P$ need not hold. In contrast, the seller's profit response to censoring is non-monotone in SNR, as shown in Figure 3.9. The underlying driver is that the seller's posterior assessment of recommendation performance is itself non-monotone in the seasonal component and market noise (see equations 3.21, B.40, and B.42). This non-monotonicity propagates to the platform's optimal DoS (η_c^*) and, in turn, to the seller's profit change under demand censoring.

3.6.4 Extension to General Trust Evolution

Our baseline and censored-demand analyses assume a relatively "quick-to-trust" seller: trust increases whenever the (posterior) performance of a recommendation is positive and decreases whenever it is negative. In practice, sellers may be more skeptical. They may (i) require a material improvement in realized profit before increasing trust (conservatism), and (ii) react more strongly to losses than to gains (cautious or asymmetric learning). This extension examines whether such more critical seller behavior can discipline the platform's incentives and attenuate the extreme self-centering outcomes observed under censored demand.

To capture these behaviors, we consider the following alternative trust update rule:

$$\alpha_{t+1} = \begin{cases} \delta_u \cdot 1 + (1 - \delta_u) \cdot \alpha_t, & \text{if } \Delta_t \Pi^S > k\, \Pi^S(\hat{q}_S, d_t), \\ \alpha_t, & \text{if } \Delta_t \Pi^S = 0, \\ \delta_d \cdot 0 + (1 - \delta_d) \cdot \alpha_t, & \text{if } \Delta_t \Pi^S < 0, \end{cases} \tag{3.29}$$

where $k \in (0, 1)$ captures seller conservatism—trust increases only when the improvement relative to the benchmark profit is sufficiently large— and $0 < \delta_u < \delta_d < 1$ captures cautious or asymmetric learning, so that trust decreases more rapidly than it increases.

In the numerical experiments below, we separate the effects of asymmetric learning and seller conservatism by studying them in isolation. To examine asymmetric learning, we set $k = 0$ and $(\delta_u, \delta_d) = (0.1, 0.2)$, so that trust decreases more rapidly after losses than it increases after gains. To examine seller conservatism, we instead set $\delta_u = \delta_d = 0.2$ and $k = 0.2$, so that trust increases only when the realized profit improvement exceeds 20% of the benchmark profit.

Because these alternative trust dynamics preclude closed-form steady-state solutions, all optimal values in this section are obtained via simulation. For each $\eta \in \{0, 0.01, 0.02, \ldots, 1\}$, we run 15 independent simulations with a horizon of 1,000 periods and select the degree of self-centering that maximizes the platform's long-run average profit.

Observation 5. *When sellers are more conservative or exhibit asymmetric learning, the seller's long-run average trust is lower, and both the platform's and the seller's long-run average profits decline.*

Observation 5 follows directly from how conservatism and cautiousness alter trust dynamics relative to the baseline update rule. With conservatism, many "small wins" from following a recommendation no longer generate trust increases. With asymmetric learning, negative outcomes reduce trust more sharply than positive outcomes rebuild it. As a result, the expected drift of trust toward higher levels is reduced, and long-run average trust converges to a lower value. The middle panels of Figures 3.12 and 3.13 illustrate this effect by plotting steady-state trust as a function of DoS under cautious and conservative sellers, respectively. Relative to the standard model (black curves), trust curves shift downward throughout.

Lower trust weakens the coordinating role of recommendations, leading to lower long-run average profits for both parties. Figure 3.11 plots the platform's long-run average profit as a function of DoS under the standard model, asymmetric learning, and seller conservatism. In all cases, profits under conservative or cautious trust dynamics are uniformly lower than

in the baseline, regardless of whether demand is censored.

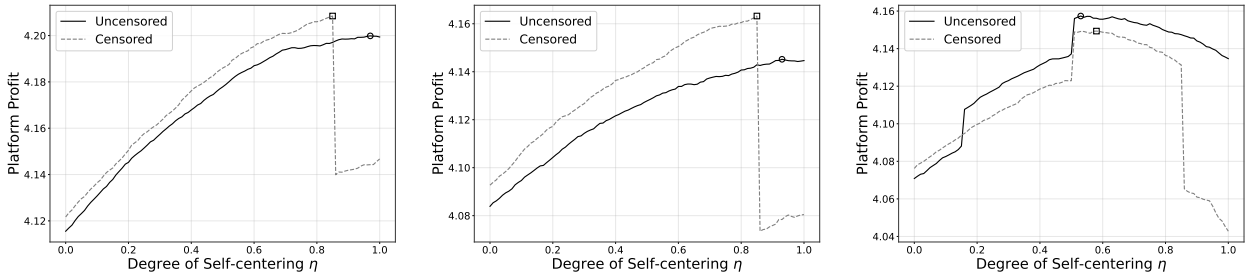

(a) baseline model (b) asymmetric learning (c) conservative seller

Figure 3.11: Platform's long-run average profit as a function of the DoS under alternative seller trust behaviors (marker indicates the maximum).

Observation 6. *Under asymmetric learning with $k = 0$, the platform optimally reduces its degree of self-centering when demand is fully observed, but does not adjust its optimal degree of self-centering when demand is censored.*

Relative to the baseline model, asymmetric learning alters the speed of trust adjustment but not the classification of outcomes: the seller continues to increase trust after positive recommendation performance and to decrease trust after negative performance. As a result, the qualitative relationship between long-run average trust ($\bar{\alpha}$ and $\bar{\alpha}^c$) and the degree of self-centering remains unchanged, but steady-state trust is uniformly lower for any given recommendation stance because negative realizations carry greater weight. In the baseline regime with fully observed demand, this lower trust reduces the effectiveness of recommendations. Anticipating weaker adoption, the platform optimally moderates its recommendation stance by lowering the degree of self-centering in order to improve perceived recommendation quality and sustain coordination.

By contrast, in the censored-demand regime with $k = 0$, the seller's trust update criterion depends only on the sign of the posterior profit difference $\Delta\hat{\Pi}$. Asymmetric learning affects the magnitude of trust adjustments conditional on this sign, but does not shift the threshold $\bar{\eta}^c$ at which posterior evaluation switches from favorable to unfavorable. Consequently, Proposition 3.11 continues to apply, and the platform's optimal degree of self-centering η_c^* coincides with that in the standard model.

When $k > 0$, seller conservatism alters the update criterion itself by raising the bar for trust increases. We analyze the resulting moderation of the platform's recommendation stance and the associated weakening of the informational masking effect in Observation 7.

This pattern is illustrated in the left panel of Figure 3.12, where the optimal DoS under asymmetric learning (gray curves) shifts in the baseline regime but closely overlaps with the standard model in the censored regime.

The right panel further shows that the profit changes induced by demand censoring are only modestly affected by asymmetric learning.

Figure 3.12: The optimal DoS η^*, long-run average trust level $\bar{\alpha}$, and change in the platform and seller's long-run average profit when demand is censored $\Delta\bar{\Pi}$ with asymmetric learning rate.

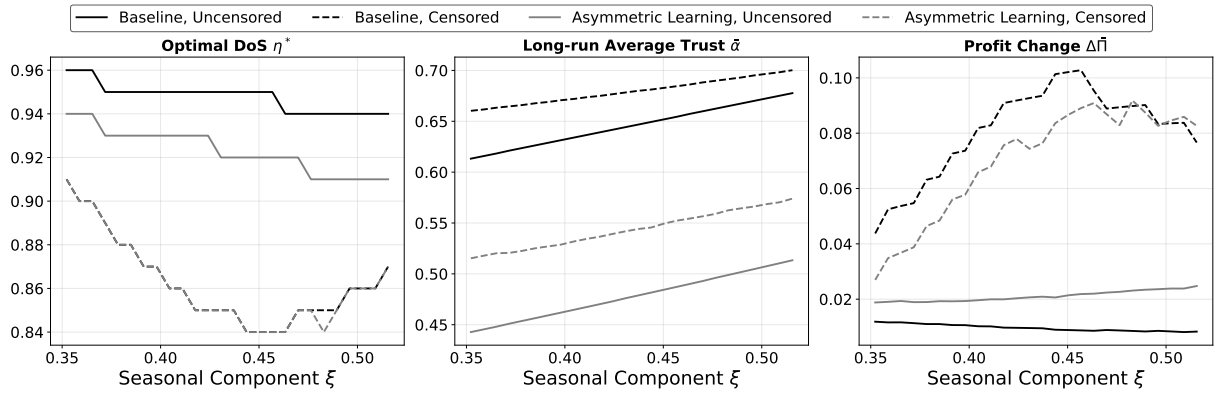

Observation 7. *Seller conservatism induces the platform to adopt a significantly more benevolent recommendation stance in both the baseline and censored-demand regimes, and it weakens the "masking effect" created by demand censoring.*

Unlike asymmetric learning rates, seller conservatism alters not only the speed of trust adjustment but also the criterion under which trust increases. When a recommendation improves realized profit but the improvement is insufficiently large, a conservative seller does not reward the platform with higher trust. This feature has two important implications for the platform's optimal degree of self-centering (DoS).

1. Discretization of trust responses and profit discontinuities. Introducing a "no-update" region—beyond trust increases and decreases—creates multiple possible long-run trust levels and renders the mapping from DoS to steady-state outcomes less smooth. For certain DoS values, the recommended quantity becomes sufficiently close to the seller's perceived benchmark in some demand states that the resulting profit improvement fails to trigger a trust increase. The seller may then become effectively "stuck" at a lower trust level in those states. This mechanism generates discontinuities in long-run outcomes, which appear as upward jumps in the profit curves (right panel of Figure 3.11).

2. Stronger trust-building incentives and a weaker masking effect. In the standard model with fully observed demand, trust updates depend solely on whether the recommendation outperforms the seller's benchmark. With a conservative seller, the platform must deliver meaningful realized profit improvements to build trust, making long-run trust more sensitive to the choice of DoS, as reflected in the pronounced curvature of the middle panel in Figure 3.13. Consequently, the platform faces a sharper trade-off between the short-run gains from self-centered recommendations and the long-run need to sustain trust. This strengthened long-run incentive can dominate the informational masking created by demand censoring, leading the platform to choose a more interior DoS and reducing the gap between the censored and uncensored regimes. This pattern is consistent with the left and right panels of Figure 3.13, which show a smaller shift in the optimal DoS and a muted profit impact of censoring under seller conservatism.

Observation 7 yields a key managerial implication. When sellers are more critical—requiring a noticeable economic gain before revising beliefs upward—the platform can no longer rely on small, frequent "wins" to build trust. Extreme self-centering therefore becomes difficult to sustain. In this sense, seller conservatism acts as a disciplining force that induces more benevolent platform behavior and limits the extent to which demand censoring can "fool" the seller.

Figure 3.13: The optimal DoS η^*, long-run average trust level $\bar{\alpha}$, and change in the platform and seller's long-run average profit when demand is censored $\Delta\bar{\Pi}$ with a conservative seller.

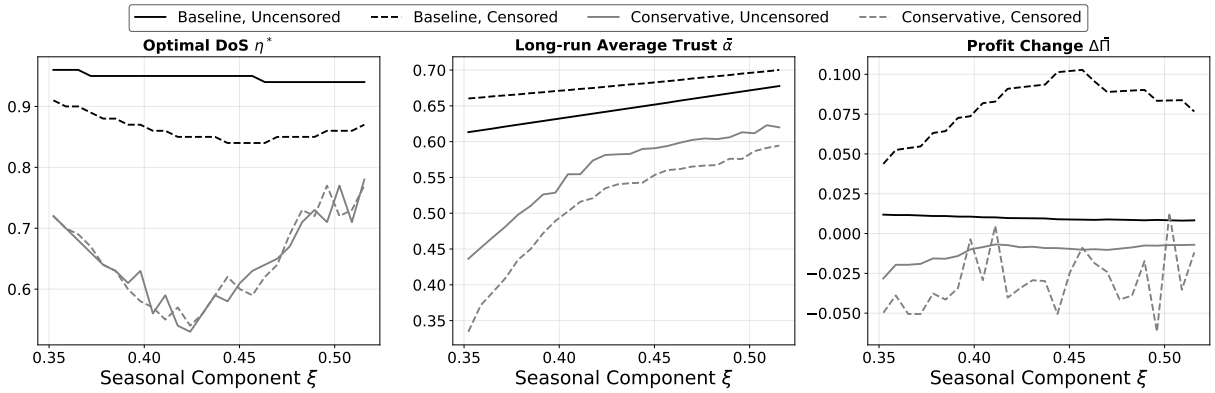

3.6.5 Summary of Numerical Study

Our numerical analysis complements the analytical results along the three dimensions described at the beginning of this section. We summarize the key findings below.

1. Coordination Value and Surplus Allocation. Inventory recommendations generate a meaningful coordination surplus by bridging the gap between the platform's superior demand information and the seller's inventory control. Across our sample, the optimal recommendation increases total supply-chain profit by up to 20%, with the seller capturing, on average, approximately 56% of the resulting surplus. Both the magnitude and the allocation of this surplus vary systematically with market conditions: the total coordination benefit is larger when interest misalignment is greater and when the signal-to-noise ratio is higher, whereas the platform captures a larger share of the surplus when the signal-to-noise ratio is lower.

2. Demand Censoring, the Informational Shield, and Win–Win Outcomes. Demand censoring creates an informational shield during stockouts that can sustain more aggressive (more self-centered) recommendations without triggering the same degree of trust erosion as under full demand observability. When the platform optimally adjusts its degree of self-centering, censoring does not materially alter the incremental benefits of recommendations relative to the no-recommendation benchmark. At the same time, censoring can improve long-run outcomes for both parties. In our simulations, censoring increases the platform's long-run average profit in 57.9% of cases and the seller's in 52.5%, yielding win–win outcomes in 50.4% of environments. The incidence and magnitude of these effects depend systematically on interest misalignment and the signal-to-noise ratio.

3. The Disciplining Effect of Seller Skepticism. Finally, by relaxing the neutral trust-update protocol, we identify an important boundary to the platform's strategic advantage. When sellers are conservative—requiring material profit improvements to build trust—or cautious—penalizing losses more heavily than gains— recommendations become less effective as a coordination mechanism, and long-run profits decline for both parties. In particular, seller conservatism substantially weakens the informational shield, forcing the platform to moderate its recommendation stance and offer more seller-aligned guidance to sustain adoption. These results indicate that seller skepticism can act as a natural disciplining force that curbs platform opportunism and limits the extent to which demand censoring can sustain extreme self-centering.

3.7 Conclusion

This paper studies how a platform should design inventory recommendations in platform-operated fulfillment (POF) services when adoption is voluntary and mediated by seller trust.

We model the interaction as a repeated game in which the platform commits to a static recommendation stance—captured by its degree of self-centering (DoS)—that trades off short-run profit extraction against long-run trust formation and adoption.

Our analysis highlights three central insights. First, inventory recommendations can generate a sustained coordination surplus for both the platform and the seller by transmitting the platform's superior demand information through trust. However, the platform's optimal recommendation stance depends critically on the seller's stocking posture. Under full demand observability, aggressive recommendations erode trust for inventory-conservative sellers but strengthen trust for aggressive ones, forcing the platform to moderate its stance in some environments. Moreover, greater information precision does not necessarily translate into higher long-run profits when adoption is incomplete.

Second, when demand is censored following stockouts, trust dynamics change fundamentally. We identify an informational shield: sellers observe lost sales from understocking but cannot verify the potential overstock costs of larger orders. This asymmetry systematically biases trust formation in favor of more aggressive recommendations. As a result, the platform's profit-maximizing strategy under censored demand is often highly self-centered and, in volatile environments, can yield higher long-run profits than full transparency.

Third, seller trust behavior places important limits on platform opportunism. Numerical results show that demand censoring does not materially reduce the coordination value of optimal recommendations and frequently produces win–win outcomes. At the same time, seller conservatism—requiring material profit improvements to build trust—weakens the informational shield and disciplines the platform's recommendation stance. Thus, trust dynamics act as an endogenous governance mechanism that shapes how informational advantages are ultimately monetized.

Taken together, our findings emphasize that the effectiveness of platform recommendations depends not only on forecast accuracy but also on how recommendations are designed and evaluated. Trust and transparency are operational design variables rather than purely informational features. Platforms that ignore seller trust risk undermining adoption, while sellers that adopt skeptical learning rules can constrain platform self-centering and improve alignment.

Our analysis focuses on a single seller–platform relationship with a static recommendation policy. Future work could examine multi-seller environments, dynamic recommendation stances that respond to evolving trust, and intertemporal inventory dynamics. These extensions would further clarify the role of trust in the design of algorithmic decision support in platform ecosystems.

APPENDIX A

Proofs and Supplementary Materials for Chapter 2

The Appendices for Chapter 2 are organized as follows:

- **Appendix A.1** presents the formulations of the referenced lemmas from Sections 2.4 and 2.5 that are not formally stated in the main body of the paper. These are listed in the order of their citation. This section also includes auxiliary lemmas that support the proofs of lemmas, propositions, and theorems from Sections 2.4 and 2.5.
- **Appendix A.2** contains the proofs for all results in Sections 2.4 and 2.5 (see A.2.1 and A.2.2, respectively). It also includes the proofs of lemmas from Appendix A.1. The details on Assumption 2.1 are provided at the end of A.2.1.
- **Appendix A.3** provides formal analyses and proofs for model extensions in Section 2.6, involving diners with partial demand learning (A.3.1), insatiable diners (A.3.2), environmentally conscious diners (A.3.3), and the impact of the plate size restriction (A.3.4).
- **Appendix A.4** covers additional analyses. Specifically, we (i) endogenize the diner's refill quantity (A.4.1), (ii) explore the impact of demand uncertainty on plate waste (A.4.2), and (iii) prove the equivalence between different monetary instruments (A.4.4).

A.1 Auxiliary Lemmas

Lemma A.1. *The expected total cost incurred by the diner at the buffet is convex in the size of the first plate and has a unique interior minimum. Formally,* $\frac{d^2 g(\varphi,q)}{dq^2} > 0$ *and* $\exists q \in (0, \bar{d})$ *such that* $\frac{dg(\varphi,q)}{dq} = 0$*. Here,* $g(\varphi, q)$ *is defined in Eq. 2.2.*

Lemma A.2. *If* $\iota < \zeta$*, the equilibrium per-diner expected plate waste is convex in the self-service cost. Formally,* $\forall \iota < \zeta$*,* $\frac{d^2\omega^*}{d\iota^2} > 0$*.*

Lemma A.3. *In equilibrium, (i) The buffet operator sets the optimal flat-rate entrance fee, given by:*

$$\varphi^*(q^*) \equiv \varphi^* = \frac{1}{2}\left(\bar{c}_o - \Delta^* + k\tau^*\right). \tag{A.1}$$

(ii) The fraction of customers who choose to join the buffet is given by:

$$\phi^*(\varphi^*, q^*) \equiv \phi^* = \frac{1}{2\bar{c}_o}(\bar{c}_o - \Delta^* - k\tau^*).$$

The expressions for Δ^* and τ^* are given in Appendix A.2.1.

Lemma A.4. *For a given flat-rate entrance fee φ_p and a flat-rate penalty fee v_p: (i) A diner's optimal first plate size and her probability of getting a plate refill are given by:*

$$q_p^*(\varphi_p, v_p) \equiv q_p^*(v_p) = \begin{cases} \frac{(\iota - v_p)^+}{o} & v_p \geq \tilde{v}_p \\ \frac{(u\bar{d} - v_p)^+}{u+o} & v_p < \tilde{v}_p \end{cases}, \quad \mathbb{P}(\textit{refill}) = \frac{1}{\bar{d}}\left(1 - \frac{\iota}{u} - \frac{(\iota - v_p)^+}{o}\right)\mathbf{1}_{v_p \geq \tilde{v}_p}, \tag{A.2}$$

where $\tilde{v}_p$ is defined in Appendix A.2.2.
(ii) The fraction of customers who choose to join the buffet is given by:

$$\phi_p^*(\varphi_p, v_p) = \frac{(\bar{c}_o - \varphi_p - \mu_p(v_p) - \Delta(q_p^*(v_p)))^+}{\bar{c}_o}, \tag{A.3}$$

where $\mu_p(v_p)$ is defined in Appendix A.2.2 and $\Delta(q)$ is defined in Eq. 2.2.

Lemma A.5. *The equilibrium per-diner expected plate waste under the optimal flat-rate penalty fee is strictly lower than that without a penalty. Formally: $\omega_p^* < \omega^*$. Here, $\omega_p^*(v_p^*) = \frac{(q_p^*(v_p^*))^2}{2\bar{d}}$ and $q_p^*(v_p^*)$ is defined in Lemma A.4.*

Lemma A.6. *The per-diner expected plate waste post-education is strictly lower than that pre-education. Formally, $\forall o_e > o$, $\omega^*(o_e) < \omega^*$. Here, $\omega^*(o_e) = \frac{(q^*(o_e))^2}{2\bar{d}}$ and $q^*(o_e)$ is defined in Proposition 2.2.*

Lemma A.7. *In a model where a buffet operator chooses an education level optimally (i.e., if $o_e = o_e^*$): (i) It is optimal to educate if and only if the unit production cost is sufficiently large. Formally, $\forall \iota$, $\exists \underline{k}_e(\iota) > 0$, such that $o_e^* = 0$ if $k \leq \underline{k}_e(\iota)$, otherwise, $o_e^* = \min\{2k, u\} > 0$. The expression for $\underline{k}_e(\iota)$ is given in Appendix A.2.2.*
(ii) Further, for any given unit production cost, there exists a threshold on the self-service cost, such that post-education, a diner has a positive probability of getting a plate refill if and only if the self-service cost is below this threshold. Formally, $\forall k$, $\exists \underline{\iota}_e(k)$ such that $q_e^* = \frac{\iota}{o_e^*}$ and $\mathbb{P}(\textit{refill}) > 0$ if $\iota < \underline{\iota}_e(k)$, otherwise, $q_e^* = \frac{u\bar{d}}{o_e^* + u}$ and $\mathbb{P}(\textit{refill}) = 0$. The expression for $\underline{\iota}_e(k)$ is given in Appendix A.2.2.

Lemma A.8. *There exists a threshold on the post-education unit overage cost, such that the threshold on the unit production cost in Theorem 2.13 weakly decreases in the post-education unit overage cost if the post-education unit overage cost is below the threshold. Formally,* $\exists o \leq o_e^k \leq u$ *such that* $\frac{dk_e(o_e)}{do_e} \leq 0$ *if* $o_e < o_e^k$ *and* $\frac{dk_e(o_e)}{do_e} > 0$ *otherwise. The expression for* o_e^k *is given in Appendix A.2.2.*

A.2 Proofs for Results in Sections 2.4 and 2.5

In what follows, we will use the following notation: $\wedge$ that stands for the logical operation 'and', and $\vee$ that stands for the logical operation 'or'.

A.2.1 Proofs of Results and Auxiliary Appendix Lemmas in Section 2.4

Proof of Auxiliary Appendix Lemma A.1

By definition (Eq. 2.2), $g(\varphi, q)$ is a piece-wise function of q. Specifically, by $u(D-q) \geq \iota \Leftrightarrow D \geq \frac{\iota}{u} + q$, we have

$$g(\varphi, q) = \varphi + \iota + o\int_0^q (q-d)f(d)\mathrm{d}(d) + \begin{cases} u\int_q^{\frac{\iota}{u}+q}(d-q)f(d)\mathrm{d}(d) + \int_{\frac{\iota}{u}+q}^{\bar{d}} cf(d)\mathrm{d}(d) & \text{if } q < \bar{d} - \frac{\iota}{u} \\ u\int_q^{\bar{d}}(d-q)f(d)\mathrm{d}(d) & \text{if } q \geq \bar{d} - \frac{\iota}{u} \end{cases}. \tag{A.4}$$

When $q < \bar{d} - \frac{\iota}{u}$,

$$\begin{aligned} \frac{dg(\varphi,q)}{dq} &= o[F(q) + qf(q) - qf(q)] + u[(\frac{\iota}{u}+q)f(\frac{\iota}{u}+q) - qf(q)] \\ &\quad - u(F(q+\frac{\iota}{u}) - F(q)) - uq(f(q+\frac{\iota}{u}) - f(q)) \\ &\quad - cF(q+\frac{\iota}{u}) = oF(q) - u(F(q+\frac{\iota}{u}) - F(q)) = o\frac{q}{\bar{d}} - \frac{\iota}{\bar{d}}, \end{aligned} \tag{A.5}$$

$$\frac{d^2g(\varphi,q)}{d(q)^2} = of(q) - u(f(q+\frac{\iota}{u}) - f(q)) = \frac{o}{\bar{d}}, \tag{A.6}$$

where we used that D is uniformly distributed with the cdf $F(x) = \frac{x}{\bar{d}}$ and pdf $f(x) = \frac{1}{\bar{d}}$. By Eq. A.6, $g(\varphi, q)$ is convex in q on $[0, \bar{d} - \frac{\iota}{u}]$. Similarly, When $q \geq \bar{d} - \frac{\iota}{u}$,

$$\begin{aligned}\frac{dg(\varphi,q)}{dq} &= o[F(q) + qf(q) - qf(q)] + u[-qf(q) - (1 - F(q)) + qf(q)] = -u + (o+u)F(q) \\ &= -u + \frac{(o+u)q}{\bar{d}}, \\ \frac{d^2g(\varphi,q)}{d(q)^2} &= (o+u)f(q) = \frac{o+u}{\bar{d}},\end{aligned} \tag{A.7}$$

and $g(\varphi, q)$ is convex in q on $[\bar{d} - \frac{\iota}{u}, \bar{d}]$. Furthermore, $\frac{dg(\varphi,q)}{dq}$ is continuous at $q = \bar{d} - \frac{\iota}{u}$ since

$$oF(q) - u(F(q + \frac{\iota}{u}) - F(q)) = oF(q) - u(1 - F(q)) = -u + (o+u)F(q).$$

Hence, $g(\varphi, q)$ is continuously differentiable and convex in q on the entire range $[0, \bar{d}]$. By Equations A.5 and A.7 and since $\iota > 0$, we obtain $\frac{dg(\varphi,q)}{dq}|_{q=0} < 0$ and $\frac{dg(\varphi,q)}{dq}|_{q=\bar{d}} > 0$, and $g(\varphi, q)$ has a unique minimum within $(0, \bar{d})$. □

Proof of Proposition 2.2

The first order condition $\frac{dg(\varphi,q)}{dq} = 0$ in two ranges (Eq. A.5 and A.7) implies

$$o\frac{q}{\bar{d}} - \frac{\iota}{\bar{d}} = 0 \Rightarrow q^* = \frac{\iota}{o}, \tag{A.8}$$

$$-u + \frac{(o+u)q}{\bar{d}} = 0 \Rightarrow q^* = \frac{u}{u+o}\bar{d}. \tag{A.9}$$

Observing the following relationship

$$\iota < \frac{ou\bar{d}}{o+u} \Leftrightarrow \frac{\iota}{o} < \frac{u}{u+o}\bar{d} \Leftrightarrow \frac{\iota}{o} < \bar{d} - \frac{\iota}{u} \Leftrightarrow \frac{u}{u+o}\bar{d} < \bar{d} - \frac{\iota}{u}, \tag{A.10}$$

we conclude that $q^* = \min\{\frac{\iota}{o}, \frac{u}{u+o}\bar{d}\}$.
(i) Define $\zeta = \bar{d}(\frac{1}{u} + \frac{1}{o})^{-1}$, we have

$$q^* = \frac{\min\{\iota, \zeta\}}{o}. \tag{A.11}$$

A diner refills if $d > q^* + \frac{\iota}{u}$ (the inequality can be non-strict, in that case, we add an event with probability zero). When $\iota < \zeta$, it implies a diner refills if $d > \frac{\iota}{o} + \frac{\iota}{u}$. When $\iota \geq \zeta$, since $q^* + \frac{\iota}{u} = \frac{u\bar{d}}{o+u} + \frac{\iota}{u} \geq \frac{u\bar{d}}{o+u} + \frac{o}{u+o}\bar{d} = \bar{d}$, the diner never refills.
(ii) By Eq. A.11, q^* weakly increases in ι. By definition, $\mathbb{P}(\text{refill}) = \mathbb{P}(u(D - q^*) > \iota)$. When

$\iota < \zeta$, $\mathbb{P}(\text{refill}) = \frac{\bar{d}-(q^*+\frac{\iota}{u})}{\bar{d}} = 1 - \frac{\iota}{\zeta}$ and $\mathbb{P}(\text{refill})$ decreases in ι and. When $\iota \geq \zeta$, since $u(\bar{d} - q^*) \leq \iota$, $\mathbb{P}(\text{refill}) = 0$. In sum, ex-ante,

$$\mathbb{P}(\text{refill}) = 1 - \min\{\frac{\iota}{\zeta}, 1\}. \tag{A.12}$$

□

Proof of Theorem 2.3

By definition (Eq. 2.4),

$$\omega^* = \mathbb{E}_D[(q^* - D)^+] = \int_0^{q^*} (q^* - d)\frac{1}{\bar{d}}\mathrm{d}(d) = \frac{(q^*)^2}{2\bar{d}}. \tag{A.13}$$

By Propostion 2.4 and Eq. A.11, Theorem 2.3 holds. Further ω^* weakly increases in u and decreases in o. If $\iota < \zeta$, ω^* decreases in $\bar{d}$. Otherwise, ω^* increases in $\bar{d}$. □

Proof of Auxiliary Appendix Lemma A.2

By Equations A.11 and A.13, when $\iota < \zeta$,

$$\frac{d^2\omega^*}{d\iota^2} = \frac{\partial^2\omega^*}{d(q^*)^2}\frac{dq^*}{d\iota} > 0.$$

□

Proof of Auxiliary Appendix Lemma A.3

A customer will join if $c_o \geq g(\varphi, q^*)$. Therefore,

$$\phi(\varphi, q^*) = \mathbb{P}_{c_o}[(c_o - g(\varphi, q^*))^+] = \frac{(\bar{c}_o - g(\varphi, q^*))^+}{\bar{c}_o} = \frac{(\bar{c}_o - \varphi - \Delta(q^*))^+}{\bar{c}_o}. \tag{A.14}$$

Hence, by definition (Eq. 2.3):

$$\pi(\varphi, q^*) = \phi(\varphi, q^*)(\varphi - k\tau(q^*)) = \frac{(c_o - \varphi - \Delta(q^*))^+}{\bar{c}_o}(\varphi - k\tau(q^*)). \tag{A.15}$$

(i) If $\bar{c}_o > \Delta(q^*) + \tau(q^*) = \Delta^* + \tau^*$, $\pi(\varphi)$ is concave in φ on $(k\tau^*, \bar{c}_o - \Delta^*)$ with a unique maximum at

$$\varphi^* = \frac{\bar{c}_o - \Delta^* + k\tau^*}{2}. \tag{A.16}$$

Otherwise, the buffet operator is unprofitable. (ii) is derived by inserting Eq. A.16 into Eq. A.14. □

Proof of Proposition 2.6

By Eq. A.15 and Appendix Lemma A.3,

$$\pi(\varphi^*, q^*) \equiv \pi^* = \frac{(\bar{c}_o - \Delta^* - k\tau^*)^2}{4\bar{c}_o}. \tag{A.17}$$

When a diner takes q^* on her first plate, let the total amount of food the customer takes in the buffet be denoted by $y^* = y^*(d)$ (which is also a function of the demand realization d). Hence:

$$\begin{aligned}\tau^* &= \mathbb{E}_D[y^*] = \mathbb{E}_D[y^*\mathbf{1}_{y^*=D}] + \mathbb{E}_D[y^*\mathbf{1}_{y^*<D}] + \mathbb{E}_D[y^*\mathbf{1}_{y^*>D}] \\ &= \mathbb{E}_D[D\mathbf{1}_{y^*=D}] + \mathbb{E}_D[D\mathbf{1}_{y^*<D}] - \mathbb{E}_D[(D-y^*)\mathbf{1}_{y^*<D}] + \mathbb{E}_D[D\mathbf{1}_{y^*>D}] + \mathbb{E}_D[(y^*-D)\mathbf{1}_{y^*>D}] \\ &= \mathbb{E}_D[D] + \mathbb{E}_D[(y^*-D)\mathbf{1}_{y^*>D}] - \mathbb{E}_D[(D-y^*)\mathbf{1}_{y^*<D}] = \mathbb{E}_D[D] + w^* - h^*. \end{aligned} \tag{A.18}$$

By definition, the expected amount of food consumed by the diner (that we denote by c^*) is

$$\begin{aligned} c^* &= \mathbb{E}_D[\min\{y^*, D\}] = \mathbb{E}_D[y^* - (y^*-D)^+] = \mathbb{E}_D[y^*] - \mathbb{E}_D[(y^*-D)^+] \\ &= \tau^* - w^* = \mathbb{E}_D[D] - h^*, \end{aligned} \tag{A.19}$$

where the last transition follows from Eq. A.18. □

Proof of Lemma 2.5

By definition (Eq. 2.2): When $\iota \geq \zeta = \frac{ou}{o+u}\bar{d}$,

$$\Delta^* = \iota + \frac{o(q^*)^2 + u(\bar{d} - (q^*))^2}{2\bar{d}} = \iota + \frac{ou\bar{d}}{2(u+o)} = \iota + \frac{\zeta}{2}, \tag{A.20}$$

$$\tau^* = q^* = \frac{\zeta}{o}. \tag{A.21}$$

Since $\frac{d\Delta^*}{d\iota} > 0$ and $\frac{d\iota}{d\iota} = 0$, by Appendix Lemma A.3, $\frac{d\varphi^*}{d\iota} < 0$.
When $\iota < \zeta$,

$$\begin{aligned}\Delta^* &= \iota + \frac{\frac{o}{2}(q^*)^2 + u\int_{q^*}^{q^*+\frac{\iota}{u}}(d-q^*)\mathrm{d}(d) + \iota\int_{q^*+\frac{\iota}{u}}^{\bar{d}}\mathrm{d}(d)}{\bar{d}} \\ &= \iota + \frac{\frac{o}{2}(\frac{\iota}{o})^2 + u\left(\frac{(\frac{\iota}{o}+\frac{\iota}{u})^2-(\frac{\iota}{o})^2}{2} - \frac{\iota}{o}\frac{\iota}{u}\right) + \iota\left(\bar{d} - \frac{\iota}{o} - \frac{\iota}{u}\right)}{\bar{d}} = 2\iota - \frac{\iota^2}{2}\frac{u+o}{uo\bar{d}} = 2\iota - \frac{\iota^2}{2\zeta} \quad \text{(A.22)}\end{aligned}$$

$$h^* = \mathbb{E}_D[(D-y^*)\mathbf{1}_{y^*<D}] = \int_{q^*}^{q^*+\frac{\iota}{u}}(d-q^*)\frac{\mathrm{d}d}{\bar{d}} = \frac{\iota^2}{2u^2\bar{d}}. \quad \text{(A.23)}$$

$$\tau^* = \frac{\bar{d}}{2} + \frac{\iota^2}{2\bar{d}}(\frac{1}{o^2} - \frac{1}{u^2}) = \frac{\bar{d}}{2} + (1 - \frac{o^2}{u^2})\omega^*, \quad \text{(A.24)}$$

where the last line is due to Equations A.13, A.18 and A.23. By Appendix Lemma A.3,

$$\begin{aligned}&\frac{d\varphi^*}{d\iota} = \frac{\iota(o+u)(u(k+o)-ko)}{2\bar{d}o^2u^2} - 1 > 0 \\ &\Leftrightarrow k > \kappa_\varphi = \frac{ou}{u-o} \wedge \zeta > \iota > \iota_\varphi = \frac{2\bar{d}o^2u^2}{(o+u)(u(k+o)-ko)}.\end{aligned}$$

□

Proof of Theorem 2.7

By Eq. A.17, π^* (strictly/weakly) decreases in ι if and only if $\Delta^* + k\tau^*$ (strictly/weakly) increases in ι. When $\iota \geq \zeta$, by Equations A.20 and A.21, $\Delta^* + k\tau^*$ strictly increases in ι. When $\iota < \zeta$, Δ^* strictly increases in ι since Eq. A.22 and $\iota < \zeta$. τ^* strictly increases in ι since Eq. A.24 and $o < u$. Therefore, $\Delta^* + k\tau^*$ strictly increases in ι. □

Derivations for Assumption 2.1

For the buffet to be profitable, there are two requirements. First, the fraction of customers who choose to join the buffer is positive. Second, the per-meal profit is positive. Specifically:

$$\phi(\varphi, q^*) > 0 \wedge \varphi - k\tau^* > 0 \Leftrightarrow \bar{c}_o - \Delta^* > \varphi > k\tau^* \Rightarrow \bar{c}_o > \Delta^* + k\tau^*. \quad \text{(A.25)}$$

By Eq. A.17, π^* strictly decreases in k. Therefore, for all ι, $\exists\bar{k}(\iota)$ such that $\pi^* > 0$ if $k < \tilde{k}(\iota)$. Specifically,

$$\bar{k}(\iota) = \max\{\frac{\bar{c}_o - \Delta^*}{\tau^*}, 0\}. \quad \text{(A.26)}$$

By Theorem 2.7, the firm's profit strictly decreases in ι. Therefore, $\exists \iota > 0, k > 0$ such that $\pi^* > 0$ if and only if $\bar{c}_o > 0$. To identify the support of $\bar{k}(\iota)$, we first let $k = 0$. Define $\bar{\iota} = \min\{\iota, \pi^*|_\iota = 0, k = 0, \iota > 0\}$. By the monotonicity of π^* in ι, if $\pi^*|_{\iota=\zeta} > 0$, then $\bar{\iota} > \zeta$. Otherwise, $\bar{\iota} < \zeta$. Since $\bar{\iota} > \zeta \Leftrightarrow \pi^*|_{\iota=\zeta} > 0 \Leftrightarrow \bar{c}_o > \Delta^* + k\tau^* = \zeta + \frac{\zeta}{2} + k\frac{\zeta}{o} = \frac{3}{2}\zeta$, by Equations A.20 and A.22,

$$\pi^*|_{\iota=\bar{\iota},k=0} = 0 \Leftrightarrow \bar{c}_o = \Delta^* \Leftrightarrow \bar{\iota} = \begin{cases} 2(1-\sqrt{1-\frac{\bar{c}_o}{2\zeta}})\zeta & \bar{c}_o < \frac{3}{2}\zeta \\ \bar{c}_o - \frac{\zeta}{2} & \bar{c}_o \geq \frac{3}{2}\zeta \end{cases}. \tag{A.27}$$

Note: If $\iota > u\bar{d}$, a diner never takes food in the buffet, i.e., $\tau^* = 0$. Therefore, we require $\bar{\iota} < u\bar{d}$ to avoid such an unrealistic scenario. This requirement has no impact on the key results of our paper and can be easily relaxed. □

A.2.2 Proofs of Results and Auxiliary Appendix Lemmas in Section 2.5

Proof of Auxiliary Appendix Lemma A.4

(i) By the expression of $g(\varphi, q)$ (Eq. 2.2), we can rewrite $g_p(q_p)$ (Eq. 2.7) and its first-order derivative as

$$\begin{aligned} &g_p(\varphi_p, v_p, q_p) = g(\varphi_p, q_p) + v_p\mathbb{P}(D < q_p) = g(\varphi_p, q_p) + v_p\frac{q_p}{\bar{d}} \\ &\Rightarrow \frac{dg_p(\varphi_p, v_p, q_p)}{dq_p} = \frac{v_p}{\bar{d}} + \frac{dg(\varphi_p, q_p)}{dq_p}. \end{aligned} \tag{A.28}$$

Moreover, since the second-order derivative of g and that of g_p are equal to each other, by Appendix Lemma A.1, g_p is convex in q_p (on the entire domain $[0, \bar{d}]$) and has a unique minimum. However, this minimum is not necessarily interior. When $q_p < \bar{d} - \frac{\iota}{u}$, by Eq. A.5:

$$\frac{dg_p(\varphi_p, v_p, q_p)}{dq_p} = \frac{v_p}{\bar{d}} + \frac{oq_p - \iota}{\bar{d}} \Rightarrow q_p^*(v_p) = \frac{(\iota - v_p)^+}{o}. \tag{A.29}$$

When $q_p \geq \bar{d} - \frac{\iota}{u}$, by Eq. A.7:

$$\frac{dg_p(\varphi_p, v_p, q_p)}{dq_p} = \frac{v_p}{\bar{d}} - u + (u+o)\frac{q_p}{\bar{d}} \Rightarrow q_p^*(v_p) = \frac{(u\bar{d} - v_p)^+}{u+o}. \tag{A.30}$$

We discuss a diner's refill decisions under the penalty.
(1) When $v_p \geq u\bar{d}$, $q_p^* = 0$ since $\iota < u\bar{d} < v_p$. Moreover, since $\bar{d} - \frac{\iota}{u} > 0 = q_p^*(v_p)$, a diner

has a positive chance of refilling.
(2) When $\iota \leq v_p < u\bar{d}$, since $\frac{(u\bar{d}-v_p)^+}{u+o} < \frac{u\bar{d}-\iota}{u+o} < \bar{d} - \frac{\iota}{u}$ and $\frac{(\iota-v_p)^+}{o} = 0 < \bar{d} - \frac{\iota}{u}$, a diner has a positive chance of refilling and $q_p^* = 0$.
(3) When $v_p < \iota < u\bar{d}$, similar to Eq. A.10, a diner refills if and only if

$$\bar{d} - \frac{\iota}{u} > \frac{u\bar{d} - v_p}{u+o} \Leftrightarrow \bar{d} - \frac{\iota}{u} > \frac{\iota - v_p}{o} \Leftrightarrow v_p > \tilde{v}_p = o(\frac{\iota}{o} + \frac{\iota}{u} - \bar{d}) = o\bar{d}(\frac{\iota}{\zeta} - 1). \tag{A.31}$$

(3-i) If $\iota < \zeta$, $\tilde{v}_p < 0$, a diner has a positive chance of refilling.
(3-ii) If $\iota \geq \zeta$: If $v_p < \tilde{v}_p$, $\mathbb{P}(\text{refill}) = 0$.

In sum,

$$\mathbb{P}(\text{refill}) = \frac{1}{\bar{d}} \left(1 - \frac{\iota}{u} - \frac{(\iota - v_p)^+}{o}\right) \mathbf{1}_{v_p \geq \tilde{v}_p}, \tag{A.32}$$

and this expression is increasing in v_p.
(ii) The proof is similar to the baseline model (Lemma A.3),

$$\phi_p(\varphi_p, v_p, q_p^*(v_p)) = \mathbb{P}_{c_o}[(c_o - g_p(\varphi_p, v_p, q_p^*(v_p)))^+] = \frac{(\bar{c}_o - \varphi_p - \mu_p(v_p) - \Delta(q_p^*(v_p)))^+}{\bar{c}_o},$$

where $\mu_p(v_p) = v_p \mathbb{P}(D < q_p^*(v_p))$. □**Note:** In the non-trivial cases of $\iota < u\bar{d}$, by cases (1) and (2), if $v_p > \iota$, $q_p^*(v_p) = 0$. Whenever $q_p^*(v_p) = 0$, a diner's refill decision and the amount of food taken are independent of the specific v_p. Therefore, in our further analysis, we will treat the discussion of any $v_p > \iota$ as $v_p = \iota$. By this treatment, together with the assumption of $\iota < u\bar{d}$, we can simplify $q_p^*(v_p)$ in Eq. A.29 and A.30 as

$$q_p^*(v_p) = \begin{cases} \frac{\iota - v_p}{o} & v_p \geq \tilde{v}_p \\ \frac{u\bar{d} - v_p}{u+o} & v_p < \tilde{v}_p \end{cases}. \tag{A.33}$$

Proof of Auxiliary Appendix Lemma A.5

Since Eq. A.13 continues to hold when there is a flat-rate penalty fee, $\omega_p^* = (q_p^*(v_p))^2/2/\bar{d}$. Since $dq_p(v_p)^*/dv_p \leq 0$, then $d\omega_p^*/dv_p \leq 0$, and $q_p^*(v_p)$ is continuous in v_p (by Eq. A.2). □

Proof of Lemma 2.8

By definition (Eq. 2.8) and Eq. A.3 (using the shorthand $\Delta(q_p^*(v_p))$ as $\Delta(v_p)$ and $\tau(q_p^*(v_p))$ as $\tau(v_p)$), and with slight abuse of notation:

$$\pi_p \equiv \pi_p(\varphi_p, v_p) \equiv \pi_p(\varphi_p, v_p, q_p^*(v_p)) = \frac{1}{\bar{c}_o}(\bar{c}_o - \varphi_p - \mu_p(v_p) - \Delta(v_p))(\varphi_p + \mu_p(v_p) - k\tau(v_p)). \tag{A.34}$$

The first order conditions with respect to φ_p and v_p are, respectively, given by:

$$-(\varphi_p + \mu_p(v_p) - k\tau(v_p)) + (\bar{c}_o - \varphi_p - \mu_p(v_p) - \Delta_p(v_p)) = 0, \tag{A.35}$$

$$(\varphi_p + \mu_p(v_p) - k\tau(v_p))\frac{d(-\mu_p(v_p) - \Delta_p(v_p))}{dv_p} + (\bar{c}_o - \varphi_p - \mu_p(v_p) - \Delta_p(v_p))\frac{d(\mu_p(v_p) - k\tau(v_p))}{dv_p} = 0. \tag{A.36}$$

If an interior extreme point exists, we must have

$$\varphi_p^* \equiv \varphi_p^*(v_p^*) = \frac{1}{2}\left(\bar{c}_o - \Delta_p^* + k\tau_p^*\right) - \mu_p^*, \tag{A.37}$$

where Δ_p^*, τ_p^* and μ_p^* are shorthand notations for $\Delta(v_p^*), \tau(v_p^*)$ and $\mu_p(v_p^*)$, respectively. Inserting Eq. A.37 into Eq. A.36, we have v_p^* that satisfies

$$\frac{1}{2}\left(\bar{c}_o - \Delta_p^* - k\tau_p^*\right)\frac{d(-\Delta_p^* - k\tau_p^*)}{dv_p} = 0 \Leftrightarrow \frac{d(\Delta_p^* + k\tau_p^*)}{dv_p} = 0. \tag{A.38}$$

To examine the second-order condition, we compute each of the entries of the Hessian Matrix,

$$\frac{d^2\pi_p}{d\varphi_p^2} = -2, \ \frac{d^2\pi_p}{d\varphi_p dv_p} = \frac{d^2\pi_p}{dv_p d\varphi_p} = \frac{d(-\mu_p(v_p) - \Delta(v_p))}{dv_p} - \frac{d(\mu_p(v_p) - k\tau(v_p))}{dv_p}$$

$$\frac{d^2\pi_p}{dv_p^2} = 2\frac{d(-\mu_p(v_p) - \Delta(v_p))}{dv_p}\frac{d(\mu_p(v_p) - k\tau(v_p))}{dv_p} + (\varphi_p + \mu_p(v_p) - k\tau(v_p))\frac{d^2(-\mu_p(v_p) - \Delta(v_p))}{dv_p^2} + (\bar{c}_o - \varphi_p - \mu_p(v_p) - \Delta(v_p))\frac{d^2(\mu_p(v_p) - k\tau(v_p))}{dv_p^2}.$$

By Eq. A.37,

$$\frac{d^2\pi_p}{dv_p^2}|_{\varphi_p=\varphi_p^*} = 2\frac{d(-\mu_p(v_p)-\Delta(v_p))}{dv_p}\frac{d(\mu_p(v_p)-k\tau(v_p))}{dv_p} - \frac{1}{2}\left(\bar{c}_o-\Delta(v_p)-k\tau(v_p)\right)\frac{d^2(\Delta(v_p)+\tau(v_p))}{dv_p^2}.$$

Therefore,

$$|\mathbb{H}| = \left(\frac{d^2\pi_p}{d\varphi_p^2}\frac{d^2\pi_p}{dv_p^2} - \frac{d^2\pi_p}{d\varphi_p dv_p}\frac{d^2\pi_p}{dv_p d\varphi_p}\right)|_{\varphi_p=\varphi_p^*, v_p=v_p^*} = \left(\bar{c}_o-\Delta_p^*-k\tau_p^*\right)\frac{d^2(\Delta_p^*+\tau_p^*)}{dv_p^2}, \quad \text{(A.39)}$$

where we used Eq. A.38 in the last equality.

Hereafter, we examine first-order and second-order conditions (Eq. A.37 to A.39) over different v_p based on Lemma A.4:

(1) If $v_p > \tilde{v}_p$, a diner has a positive probability of refilling. Similar to the derivation of Δ and τ when a diner refills in the base model (Eq. A.18, A.22, A.23 and A.24), when the penalty fee is v_p, we have

$$\Delta(v_p) = \Delta(q_p^*(v_p)) = \iota + o\frac{(q_p^*(v_p))^2}{2\bar{d}} + \frac{\iota^2}{2u\bar{d}} + \iota\frac{\bar{d}-(q_p^*(v_p)+\frac{\iota}{u})}{\bar{d}}$$
$$= \iota + \frac{uv_p^2 + 2ou\bar{d}\iota - (o+u)\iota^2}{2ou\bar{d}}. \quad \text{(A.40)}$$

$$\tau(v_p) = \tau(q_p^*(v_p)) = \frac{\bar{d}}{2} + \frac{1}{2\bar{d}}((\frac{\iota-v_p}{o})^2 - \frac{\iota^2}{u^2}). \quad \text{(A.41)}$$

$$\frac{d(\Delta(v_p)+k\tau(v_p))}{dv_p} = 0 \Leftrightarrow \frac{(k+o)v_p - k\iota}{o^2\bar{d}} = 0 \Leftrightarrow v_p^* = \frac{k}{k+o}\iota. \quad \text{(A.42)}$$

$$\frac{d^2(\Delta(v_p)+k\tau(v_p))}{dv_p^2} = \frac{1}{(o+u)\bar{d}} = \frac{k+o}{o^2\bar{d}} > 0.$$

By the condition of the case (1), we require

$$v_p^* = \frac{k}{k+o}\iota > \tilde{v}_p = o\bar{d}(\frac{\iota}{\zeta}-1) \Leftrightarrow \iota < (\frac{1}{k+o}+\frac{1}{u})^{-1}\bar{d}. \quad \text{(A.43)}$$

In this case,

$$q_p^* = q_p^*(v_p^*) = \frac{\iota - \frac{k}{k+o}\iota}{o} = \frac{\iota}{o+k}. \quad \text{(A.44)}$$

(2) If $v_p \leq \tilde{v}_p$, a diner never refills. Similar to the derivation of Δ and τ when a diner never refills in the base model (Eq. A.20 and A.21), when the penalty fee is v_p, we have:

$$\Delta(v_p) = \iota + o\frac{(q_p^*(v_p))^2}{2\bar{d}} + u\frac{(\bar{d} - q_p^*(v_p))^2}{2\bar{d}} = \iota + \frac{\zeta}{2} + \frac{v_p^2}{2(o+u)\bar{d}}. \tag{A.45}$$

$$\tau(v_p) = q_p^*(v_p) = \frac{u\bar{d} - v_p}{u+o}. \tag{A.46}$$

$$\frac{d(\Delta(v_p) + k\tau(v_p))}{dv_p} = 0 \Leftrightarrow v_p^* = k\bar{d}. \tag{A.47}$$

$$\frac{d^2(\Delta(v_p) + k\tau(v_p))}{dv_p^2} = \frac{1}{(o+u)\bar{d}} > 0.$$

By the condition of the case (2), we require

$$k\bar{d} \leq \tilde{v}_p \Leftrightarrow \iota \geq \frac{(o+k)u\bar{d}}{o+u}. \tag{A.48}$$

In this case,

$$q_p^* = q_p^*(v_p^*) = \frac{u\bar{d} - k\bar{d}}{o+u}. \tag{A.49}$$

(3) When neither of the Equations A.43 and A.48 holds, i.e., $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} \leq \iota < \frac{(o+k)u\bar{d}}{o+u}$, it indicates that there does not exist (φ, v_p) that can simultaneously satisfy both first-order conditions (on φ_p (Eq. A.37) and v_p (Eq A.38)). In this case, the optimal solution must be at the boundary. Moreover, given any v_p, since π_p^* is maximized if Eq. A.37 holds, v_p^* must be at the boundary and minimize $\Delta(v_p) + k\tau(v_p)$. Since $\iota < \frac{(o+k)u\bar{d}}{o+u}$, we have $k\bar{d} > \tilde{v}_p$ and $\Delta(v_p) + k\tau(v_p)$ strictly decreases in $(0, \tilde{v}_p)$; since $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} \leq \iota$, we have $\frac{k}{k+o}\iota \leq \tilde{v}_p$ and $\Delta(v_p) + k\tau(v_p)$ strictly increases in $(\tilde{v}_p, \iota)$; finally, since $\Delta(v_p) + k\tau(v_p)$ is continuous, $v_p^* = \tilde{v}_p$. In this case,

$$q_p^* = q_p^*(v_p^*) = \frac{\iota - \tilde{v}_p}{o} = \frac{\iota - o\bar{d}(\frac{\iota}{\zeta} - 1)}{o} = \bar{d} - \frac{\iota}{u} = \frac{u\bar{d} - \tilde{v}_p}{u+o}. \tag{A.50}$$

Figure A.1a illustrates the optimal flat-rate penalty fee under different values of the self-service cost and unit production cost. Curve $k = \frac{(o+u)\iota - ou\bar{d}}{u\bar{d} - \iota}$ and $k = \frac{u+o}{u\bar{d}}\iota - o$ are derived from Equations A.43 and A.48. □

Proof of Lemma 2.9

By Equations A.11, A.13, A.44, A.49 and A.50, and Appendix Lemma A.5, we have

$$\frac{d(\omega^* - \omega_p^*)}{dk} = -\frac{d\omega_p^*}{dk} = \frac{\partial \omega_p^*}{\partial k}\frac{dq_p^*}{dk} \geq 0.$$

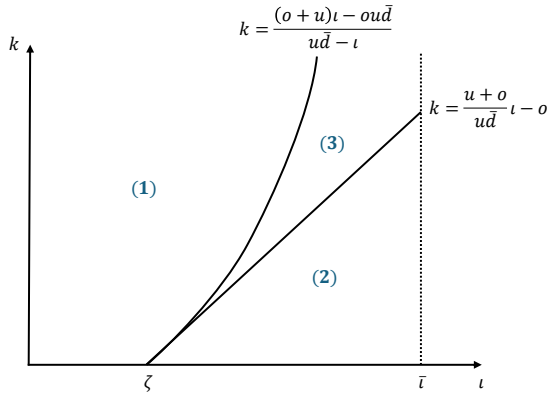

(a) The optimal flat-rate penalty fee v_p^* equals $\frac{k}{o+k}\iota$ in region (1), $k\bar{d}$ in region (2), and $o(\frac{\iota}{\zeta} - 1)$ in region (3).

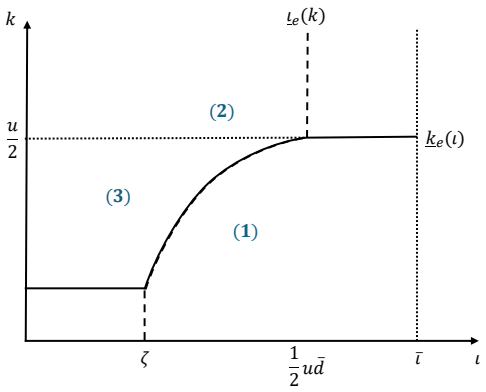

(b) The optimal education level o_e^* equals o in region (1), u in region (2), and $2k$ in region (3).

Figure A.1: The optimal flat-rate penalty fee and the optimal education level across different combinations of self-service costs and unit production costs.

Note: Definitions of the expressions used are provided in the proofs of Lemma 2.8 and Appendix Lemma A.7, respectively.

□

Proof of Theorem 2.10

Inserting Eq. A.37 into Eq. A.34, the expression for the equilibrium profit when the buffet operator sets the optimal penalty is given by:

$$\pi_p^* = \pi_p(\varphi_p^*(v_p^*), v_p^*) = \frac{1}{4\bar{c}_o}(\bar{c}_o - \Delta(v_p^*) - k\tau(v_p^*))^2. \tag{A.51}$$

By Eq. 2.9 and the fact that $v_p^* > 0$ from the proof of Appendix Lemma A.5, Theorem 2.10 holds. □

Proof of Lemma 2.11

By Equations 2.5 and A.51,

$$\pi_p^* - \pi^* = \frac{1}{4\bar{c}_o}\left((\Delta^* + k\tau^*) - (\Delta_p^* + k\tau_p^*)\right)\left(\bar{c}_o - (\Delta^* + k\tau^*) + \bar{c}_o - (\Delta_p^* + k\tau_p^*)\right).$$

When $\iota < \zeta$, by Equations A.22, A.24, A.40, A.41 and A.42,

$$\frac{d(\pi_p^* - \pi^*)}{dk} = \frac{1}{4\bar{c}_o} \frac{\iota^2 k}{2\bar{d}^2 o^4 u^2 (k+o)^3} f_\Delta(k),$$

where

$$\begin{aligned} f_\Delta(k) = {} & \bar{d} o^2 u^2 (k+o)(2\bar{c}_o(k+2o) - \bar{d}k(2k+3o)) - 4\bar{d}\iota o^2 u^2 (k+o)(k+2o) \\ & + \iota^2 \left(-u^2 \left(k^3 + 2k^2 o - 2o^3 \right) + o^2 u (k+o)(k+2o) + k o^2 (k+o)(2k+3o) \right). \end{aligned} \tag{A.52}$$

Observe that, because of $o < u$ and $\iota < \zeta$, Assumption 2.1, and Eq. A.26,

$$\begin{aligned} f_\Delta(0) &= 2o^3 u (2\bar{c}_o \bar{d} o u + \iota(\iota(o+u) - 4\bar{d} o u)) > 0, \\ f_\Delta(\bar{k}(\iota)) &= \bar{k}(\iota) o^2 \left(\iota^2 (\bar{k}(\iota) + o - u)(\bar{k}(\iota) + o + u) - \bar{c}_o^2 u^2 (\bar{k}(\iota) + o)^2 \right) < 0. \end{aligned}$$

Since $f_\Delta(k)$ is a cubic function of k and the coefficient of the cubic term is negative $(\iota^2 (2o^2 - u^2) - 2\bar{d}^2 o^2 u^2) < 0$, there exists exactly one root between 0 and $\bar{k}(\iota)$. Therefore, there exists k_p^π, such that $d(\pi_p^* - \pi^*)/dk < 0$ if $k > k_p^\pi$ and $d(\pi_p^* - \pi^*)/dk > 0$, otherwise. □

Proof of Auxiliary Appendix Lemma A.6

$\omega^*(o_e) = \omega^*|_{o=o_e}$. Since ω^* strictly decreases in o (proved in Theorem 2.3) and $o_e > o$, Appendix Lemma A.6 holds. □

Proof of Auxiliary Appendix Lemma A.7

With slight abuse of notation, given any mathematical term x (e.g., τ, Δ and π) that is a function of the unit overage cost o (except for o itself), in what follows, we will use these three terms interchangeably: $x_e, x(o_e)$ and $x|_{o=o_e}$.

(i) When the buffet operator educates diners, the equilibrium optimal flat-rate entrance fee and the expected profit are given by Equations A.1 and 2.5, where o is substituted with o_e, i.e., $\varphi_e^* = \varphi^*(o_e), \pi_e^* = \pi^*(o_e)$ When jointly maximizing the entrance fee φ_e and the post-education unit overage cost o_e, the first order condition with respect to φ_e is the same as the base model, therefore, $\varphi_e^* = \frac{1}{2}(\bar{c}_o - \Delta^*(o_e) - k\tau^*(o_e))$. The discussion of the second-order conditions is a replication of Equations A.39 to A.50. Similar to Eq. 2.9, o_e^* minimizes $\Delta_e^* + k\tau_e^*$. If an interior profit maximizer exists, o_e^* satisfies

$$\frac{d\pi_e(\varphi_e, o_e)}{do_e}|_{\varphi_e = \varphi_e^*} = \frac{d(\Delta_e^*(o_e) + k\tau_e^*(o_e))}{do_e} = 0. \tag{A.53}$$

Next, we investigate the sensitivity of $\Delta_e^* + k\tau_e^*$ with respect to o_e.
When $\iota < \zeta_e$, by Equations A.22 and A.24,

$$\frac{d}{do_e}(\Delta_e^* + k\tau_e^*) = \frac{d}{do_e}\left(2\iota - \frac{\iota^2}{2\zeta_e} + k(\frac{\bar{d}}{2} + \frac{\iota^2}{2\bar{d}}(\frac{1}{o_e^2} - \frac{1}{u^2}))\right) = \frac{\iota^2(o_e - 2k)}{2o_e^3\bar{d}},$$
$$\frac{d^2}{do_e^2}(\Delta_e^* + k\tau_e^*) = \frac{\iota^2(3k - o_e)}{o_e^4\bar{d}}. \tag{A.54}$$

When $o_e > 2k$, $\Delta_e^* + k\tau_e^*$ strictly increases in o_e. Otherwise ($o_e \leq 2k$), it weakly decreases in o_e.
When $\iota \geq \zeta_e$, by Equations A.20 and A.21,

$$\frac{d}{do_e}(\Delta_e^* + k\tau_e^*) = \frac{d}{do_e}\left(\iota + \frac{\zeta_e}{2} + k\frac{\zeta_e}{o_e}\right) = \frac{u\bar{d}(u - 2k)}{2(o_e + u)^2}.$$

When $u > 2k$, $\Delta_e^* + k\tau_e^*$ strictly increases in o_e, and weakly decreases otherwise ($u \leq 2k$). Observe that ζ_e itself strictly increases in o_e and we can switch form the case $\iota \geq \zeta_e$ to $\iota < \zeta_e$ when o_e increases. Indeed, $\iota < \zeta_e \Leftrightarrow o_e > \frac{u\iota}{u\bar{d} - \iota}$. To obtain the sensitivity analysis, we rely on the discussion, of two separate cases, which we specify depending on the values that ι takes (Figure A.2): compare panel (a) with panel (b).

If $\iota \geq \frac{1}{2}u\bar{d}$ (Figure A.2b): $\frac{u\iota}{u\bar{d} - \iota} \geq u$. In this case, by the assumption $o_e < u$, we always have

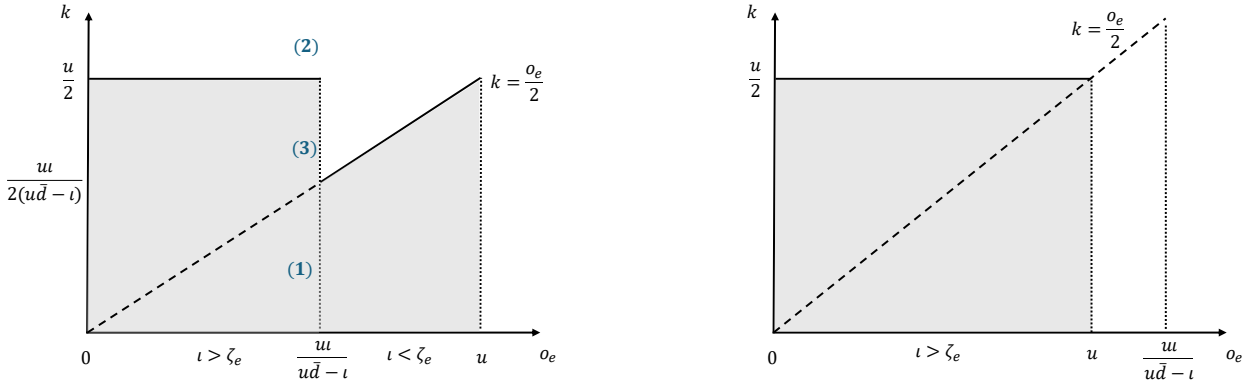

(a) $\iota < \frac{1}{2}u\bar{d}$ (b) $\iota \geq \frac{1}{2}u\bar{d}$

Figure A.2: The marginal impact of consumer education on the buffet operator's expected profit across different values of post-education unit overage costs and unit production costs, for small and large values of the self-service costs.

Note: The gray-shaded area represents combinations of (o_e, k) where the derivative of profit with respect to o_e is negative ($\frac{d\pi_e}{do_e} < 0$), indicating a decline in profit. The non-shaded (white) area corresponds to positive profit impact ($\frac{d\pi_e}{do_e} > 0$). Detailed analyses, along with definitions of terms and curve explanations, are provided in the proof of Appendix Lemma A.7.

$\iota \geq \zeta_e$. Therefore, $\Delta_e^* + k\tau_e^*$ weakly decreases in o_e if $2k \leq u$ and increases otherwise.
If $\iota < \frac{1}{2}u\bar{d}$ (Figure A.2a):
(1) If $2k < \min\{u, \min_{\iota<\zeta_e}\{o_e\}\} \Leftrightarrow k < \min\{\frac{u}{2}, \frac{u\iota}{2(u\bar{d}-\iota)}\} = \frac{u\iota}{2(u\bar{d}-\iota)}$, $\Delta_e^* + k\tau_e^*$ strictly increases in o_e. In this case, $o_e^* = \min_{o\leq o_e\leq u} = o$.
(2) If $2k \geq \max\{u, \max_{\iota<\zeta_e}\{o_e\}\} = u$, $\Delta_e^* + k\tau_e^*$ strictly decreases in o_e. In this case, $o_e^* = \max_{o\leq o_e\leq u} = u$.
(3) If $2k \in [\frac{u\iota}{2(u\bar{d}-\iota)}, u)$: for $o_e < \frac{u\iota}{u\bar{d}-\iota}$, $\Delta_e^* + k\tau_e^*$ strictly increases in o_e since $2k < u$; for $o_e > \frac{u\iota}{u\bar{d}-\iota}$ and $o_e < 2k$, $\Delta_e^* + k\tau_e^*$ strictly decreases in o_e; for $o_e \geq 2k$, $\Delta_e^* + k\tau_e^*$ strictly increases in o_e.
(3-i) If $o \geq 2k$, $\Delta_e^* + k\tau_e^*$ increases in o_e. Similar to case (1), $o_e^* = o$.
(3-ii) If $\frac{u\iota}{2(u\bar{d}-\iota)} \leq o < 2k$, $\Delta_e^* + k\tau_e^*$ decreases in o_e for $o < o_e < 2k$ and increases in o_e for $o_e > 2k$. The first order condition is satisfied at $o_e^* = 2k$. When $o_e^* = 2k$, by Eq. A.54, the Hessian matrix is negative definite and $o_e^* = 2k$ (with φ_e^*) maximizes profit.
(3-iii) If $o < \frac{u\iota}{2(u\bar{d}-\iota)}$, $\Delta_e^* + k\tau_e^*$ increases on $(o, \frac{u\iota}{2(u\bar{d}-\iota)}) \cup [2k, u)$, and decreases on $[\frac{u\iota}{2(u\bar{d}-\iota)}, 2k)$. We compare $\Delta_e^* + k\tau_e^*$ at o and $2k$. By Equations A.22 and A.24,

$$\begin{aligned}
&\Delta_e^* + k\tau_e^*|_{o_e=2k} = -\frac{\iota^2(2k+u)^2}{8\bar{d}ku^2} + \frac{\bar{d}k}{2} + 2\iota, \qquad \text{(A.55)}\\
&(\Delta_e^* + k\tau_e^*)|_{o_e=o} - (\Delta_e^* + k\tau_e^*)|_{o_e=2k}\\
&= \frac{1}{8}\left(\frac{4\bar{d}(u(k+o)-ko)}{o+u} + \frac{\iota^2(2k+u)^2}{\bar{d}ku^2} - 8\iota\right) > 0 \Leftrightarrow k > \underline{k}_e(\iota),
\end{aligned}$$

where

$$\begin{aligned}
\underline{k}_e(\iota) = &\frac{u\left(-ou^2\bar{d}^2 + 2\iota u\bar{d}(o+u) - (\iota^2(o+u))\right)}{2u^2\bar{d}^2(u-o) + 2\iota^2(o+u)}\\
&+ \frac{\sqrt{u\bar{d}\left(ou\bar{d} - \iota(o+u)\right)\left(u\bar{d}\left(ou\bar{d} - 3\iota(o+u)\right) + 4\iota^2(o+u)\right)}}{2u^2\bar{d}^2(u-o) + 2\iota^2(o+u)}.
\end{aligned}$$

Notice that $\underline{k}_e(\iota)$ is only defined for $\iota < \frac{1}{2}u\bar{d}$ and $o < \frac{u\iota}{2(u\bar{d}-\iota)} \Leftrightarrow \iota > \zeta$ (conditions of the case (3-iii)), we extend its definition to the full range of $\iota < \bar{\iota}$. Specifically, we define $\underline{k}_e(\iota) = \frac{o}{2}$ for $\iota \leq \zeta$ and $\underline{k}_e(\iota) = \frac{o}{2}$ for $\iota \geq \frac{1}{2}u\bar{d}$.
(ii):
(1) When $k < \frac{o}{2}$, $o_e^* = 0$. Without education, the threshold on a diner's refilling decision is the same as that of the base model, i.e., $\underline{\iota}_e(k) = \zeta$.
(2) When $k \geq \frac{u}{2}$, $o_e = u$. $\underline{\iota}_e(k) = \zeta|_{o=o_e^*} = \frac{1}{2}u\bar{d}$.
(3) When $\frac{o}{2} \leq k < \frac{u}{2}$, if $k > \underline{k}_e(\iota)$, $o_e^* = 2k$, $o_e^* < \frac{u\iota}{u\bar{d}-\iota} \Leftrightarrow \iota < \zeta_e^*$ and a diner refills. Otherwise,

if $\frac{o}{2} \leq k < \underline{k}_e(\iota)$, we must have $\iota \geq \frac{1}{2}u\bar{d}$. In this case, $o_e^* = o$. Without education, a diner never refills. Therefore, when $\frac{o}{2} \leq k < \frac{u}{2}$, we define $\underline{\iota}_e(k) = \underline{k}_e^{-1}(k)$, where $\underline{k}_e^{-1}$ is the inverse function of $\underline{k}_e$.
Figure A.1b illustrates the optimal post-education unit overage cost for various combinations of self-service and unit production cost values. □

Proof of Theorem 2.13

Since $\pi_e^* = \pi^*|_{o=o_e}$, by Eq. 2.5, $\pi_e^* > \pi^*$ if and only if $\Delta_e^* + k\tau_e^* < \Delta^* + k\tau^*$. Since $\tau_e = \tau|_{o=o_e}$, by Equations A.24 and A.21, if $o_e > o$, we have $\tau_e^* < \tau^*$. By the definition of Δ (Eq. 2.2), and the fact that q^* is the unique minimizer of Δ (Appendix Lemma A.1), and by the Envelope theorem, we have $\frac{\Delta(q^*)}{do} = \frac{\partial \Delta}{\partial o} > 0$ and $\Delta_e^* > \Delta^*$. Therefore,

$$\Delta_e^* + k\tau_e^* < \Delta^* + k\tau^* \Leftrightarrow k > k_e(o_e) = \frac{\Delta_e^* - \Delta^*}{\tau^* - \tau_e^*}. \tag{A.56}$$

□

Proof of Auxiliary Appendix Lemma A.8

By Appendix Lemma A.7, $\Delta_e^* + k\tau_e^*$ strictly increases in o_e if $k > \frac{u}{2}$, and strictly decreases otherwise. Moreover, if $k < \underline{k}_e$, $o_e^* = o$. Therefore, we have

$$\underline{k}_e(\iota) \leq k_e(o_e) \leq \frac{u}{2} < u. \tag{A.57}$$

(1) If $\iota \geq \frac{1}{2}u\bar{d}$, by Appendix Lemma A.7, $\underline{k}_e(\iota) = \frac{u}{2}$. By Eq. A.57, $k_e(o_e) = \frac{u}{2}$ (we can verify this result by Eq. A.56, in the case (1), $\tau_e = \frac{\bar{d}}{2}$ and $\Delta_e = \iota + \frac{u\bar{d}}{2}$). In this case, $o_e^k = u$.
(2) If $\iota < \zeta$, by Eq. A.22, A.24 and A.56, $k_e(o_e) = \frac{oo_e}{o+o_e}$. In this case, $o_e^k = o$.
(3) If $\zeta \leq \iota < \frac{1}{2}u\bar{d}$:
(3-i) When $o_e \leq \frac{u\iota}{u\bar{d}-\iota}$, similar to case (1), $k_e(o_e) = \underline{k}_e(\iota) = \frac{u}{2}$.
(3-ii) When $o_e > \frac{u\iota}{u\bar{d}-\iota}$, a diner does not refill before education (Eq. A.20, A.21) and refills post education (Eq. A.22, A.24), by definition (Eq. A.56),

$$k_e(o_e) = -\frac{uo_e\left(\iota u\left(u\left(\iota - 2\bar{d}o_e\right) + \iota o_e\right) + oo_e\left(\iota - u\bar{d}\right)^2 + \iota^2 ou\right)}{o_e^2\left(u^2\bar{d}^2(u-o) + \iota^2(o+u)\right) - \iota^2u^2(o+u)}.$$

We have $\frac{dk_e(o_e)}{do_e} > 0$ if and only if $o > o_e^k$ and

$$o_e^k = \frac{u\left(u\bar{d}\left(2\iota(o+u) - ou\bar{d}\right) - (\iota^2(o+u))\right)}{u^2\bar{d}^2(u-o) + \iota^2(o+u)} + \frac{u\left(\sqrt{u\bar{d}\left(u\bar{d}\left(ou\bar{d}\left(ou\bar{d} - 4\iota(o+u)\right) + \iota^2(o+u)(7o+3u)\right) - 4\iota^3(o+u)^2\right)}\right)}{u^2\bar{d}^2(u-o) + \iota^2(o+u)}.$$

For each of the three cases ($\iota < \zeta$; $\zeta \leq \iota < \frac{1}{2}u\bar{d}$; and $\iota \geq \frac{1}{2}u\bar{d}$), the comparison of the pre-education equilibrium profit (π^*) and post-education equilibrium profit (π_e^*) under different o_e and k is illustrated in Figures A.3a, A.3b, and A.3c, respectively. □

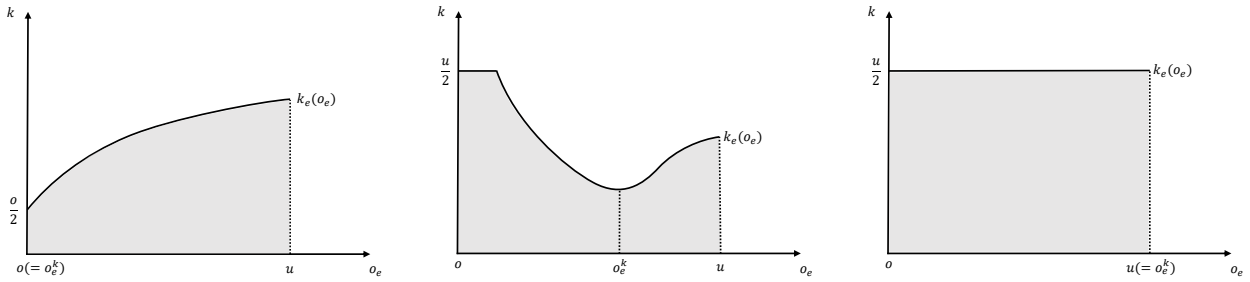

(a) $\iota < \zeta$ (b) $\zeta \leq \iota < \frac{1}{2}u\bar{d}$ (c) $\iota \geq \frac{1}{2}u\bar{d}$

Figure A.3: Equilibrium expected profits before and after education across different combinations of post-education unit overage costs (o_e) and unit production costs (k) when the self-service costs are low, moderate, and high.

Note: Shaded areas indicate reduced profit ($\pi_e^* < \pi^*$), while non-shaded areas show increased profit ($\pi_e^* > \pi^*$). Detailed definitions and expressions can be found in the proof of Appendix Lemma A.8.

Proof of Theorem 2.14

In this proof, for a term x that is a function of o (e.g.,τ), x_e^* stands for $x^*|_{o=o_e^*}$. We begin with the profitability comparison.

Profitability comparison

Take the unit overage cost post education that is set to the optimal level, i.e., $o_e = o_e^*$ in Appendix Lemma A.7. By Lemma A.4, we can always find a flat-rate penalty fee v_p^e such that the equilibrium size of the first plate under this penalty fee is the same as that under the optimal post-education unit overage cost, i.e., $q_p^*(v_p^e) = q_e^* = q_e(o_e^*)$. By Equations A.21, A.24, A.41, and A.46, this further implies $\tau_p(v_p^e) = \tau_e^* = \tau_e(o_e^*)$. Moreover, by

the definition of the expected non-monetary cost (Eq. 2.2), we have

$$\Delta_e^* = \Delta|_{o=o_e^*} = \Delta(q_e^*) + \mathbb{E}_D[(o_e^* - o)\cdot(q_e^* - D)^+] = \Delta(q_p^*(v_p^e)) + \mathbb{E}_D[(o_e^* - o)\cdot(q_e^* - D)^+] > \Delta(q_p^*(v_p^e)).$$

Therefore, $\Delta_e^* + k\tau_e^* > \Delta(q_p^*(v_p^e)) + k\tau(v_p^e)$. By Equations 2.5, A.34, A.37 and A.51,

$$\pi_p^* = \pi_p(\varphi_p^*(v_p^*), v_p^*) > \pi_p(\varphi_p^*(v_p^e), v_p^e) > \pi_e^*.$$

Waste comparison

By Appendix Lemma A.5 and Appendix Lemma A.6, the relationship between q_e^* and

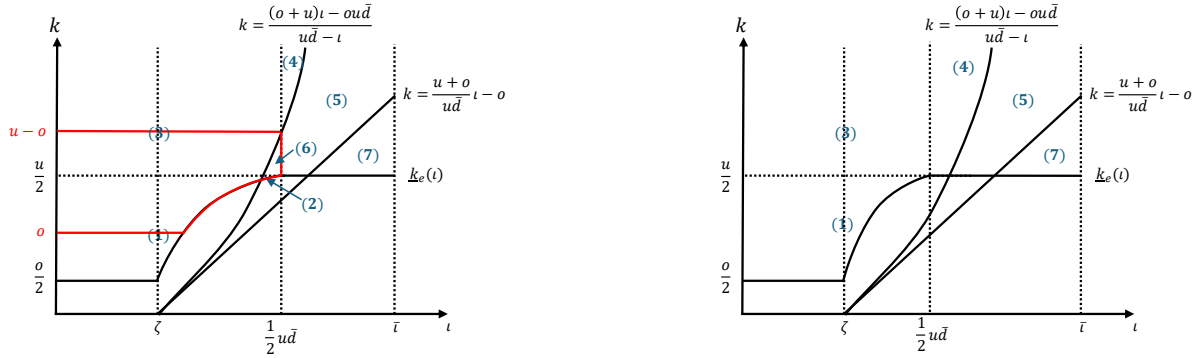

(a) $u > 2o$ (b) $u \leq 2o$

Figure A.4: Per-diner expected plate waste under the optimal flat-rate penalty fee (ω_p^*) and the optimal education level (ω_e^*) across different combinations of self-service costs (ι) and unit production costs (k) when the unit overage costs (o) is small and large.

Note: Consumer education is more effective in reducing plate waste when o and ι are sufficiently low and k is moderate. Specifically, $\omega_p^* > \omega_e^*$ if (ι, k) is within the region bounded by the red boundaries in panel (a); otherwise, $\omega_p^* \leq \omega_e^*$. Detailed definitions and expressions are provided in the proof of Theorem 2.14.

q_p^* is the same as that of ω_e^* and ω_p^*. Therefore, hereafter, we only compare q_p^* (given by Lemma A.4) and q_e^* (given by Appendix Lemma A.6). Moreover, we only look at cases where $o_e^* > o$, i.e., $k > \underline{k}_e^*(\iota)$, otherwise, we immediately have $q_p^* < q_e^*$ (by Theorem 2.13). The comparison is further divided into two cases by considering the point $(\frac{1}{2}u\bar{d}, \underline{k}_e(\frac{1}{2}u\bar{d})) = (\frac{1}{2}u\bar{d}, \frac{u}{2})$ on the graph of (ι, k). The first case is when $(\frac{1}{2}u\bar{d}, \frac{u}{2})$ falls into the region where diners are indifferent between refilling and not refilling under the optimal penalty fee, i.e.,

$$(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d}|_{k=\frac{u}{2}} < \frac{1}{2}u\bar{d} < \frac{(k+o)u}{u+o}\bar{d}|_{k=\frac{u}{2}} \Leftrightarrow u > 2o.$$

Otherwise, in the second case, when $u \leq 2o$, $(\frac{1}{2}u\bar{d}, \frac{u}{2})$ falls into the region where the probability of refilling is non-zero under the optimal penalty fee.

Case: $u > 2o$ (Figure A.4a).

We compare the plate waste under v_p^* and o_e^* in the following sub-regions.

(1) $k \in [\underline{k}_e(\iota), \frac{u}{2})$ and $\iota < (\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d}$. In this case, $q_p^* = \frac{\iota}{o+k}$ and $q_e^* = \frac{\iota}{2k}$. Therefore, if $k \leq o$, $q_p^* \leq q_e^*$. If $k > o$, $q_p^* > q_e^*$.

(2) $k \in [\underline{k}_e(\iota), \frac{u}{2})$ and $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} \leq \iota < \frac{1}{2}u\bar{d}$. In this case, $q_p^* = \bar{d} - \frac{\iota}{u}$ and $q_e^* = \frac{\iota}{2k}$. We have,

$$q_p^* - q_e^* = \bar{d} - \frac{\iota}{u} - \frac{\iota}{2k} > \bar{d} - \frac{\iota}{u} - \frac{\iota}{u} > 0,$$

where we used $k < \frac{u}{2}$ in the first inequality and $\iota < \frac{1}{2}u\bar{d}$ in the second inequality.

(3) $k \geq \frac{u}{2}$ and $\iota < \min\{(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d}, \frac{1}{2}u\bar{d}\}$. In this case, $q_p^* = \frac{\iota}{o+k}$ and $q_e^* = \frac{\iota}{u}$. Therefore, if $k \geq u - o$, $q_p^* \leq q_e^*$. If $\frac{u}{2} < k < u - o$, $q_p^* > q_e^*$.

(4) $k \geq \frac{u}{2}$ (in the case of $u > 2o$, it implies $k > u - o$) and $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} > \iota \geq \frac{1}{2}u\bar{d}$. In this case, $q_p^* = \frac{\iota}{o+k}$ and $q_e^* = \frac{\iota}{\bar{d}}$. We show that $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} < \frac{1}{2}(o+k)\bar{d}$. Rewrite $\iota = (\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d}$ as $k = \frac{\iota(o+u) - \bar{d}ou}{\bar{d}u - \iota}$. Its slope at $\frac{1}{2}u\bar{d}$ is $\frac{d}{d\iota}(\frac{\iota(o+u)-\bar{d}ou}{\bar{d}u-\iota})|_{\iota=\frac{1}{2}u\bar{d}} = \frac{4}{\bar{d}}$. Similarly, rewrite $\iota = \frac{1}{2}(o+k)\bar{d}$ as $k = 2\frac{\iota}{\bar{d}} - o$. Its slope at $\frac{1}{2}u\bar{d}$ is $\frac{2}{\bar{d}} < \frac{4}{\bar{d}}$. Therefore, $k = \frac{\iota(o+u)-\bar{d}ou}{\bar{d}u-\iota}$ and $k = 2\frac{\iota}{\bar{d}} - o$ has no intersection for $\iota > \frac{1}{2}u\bar{d}$. Therefore, $\forall \iota \in [\frac{1}{2}u\bar{d}, (\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d})$, $\iota < \frac{1}{2}(o+k)\bar{d}$ and $q_p^* < q_e^*$.

(5) $k \geq \frac{u}{2}$ and $\max\{(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d}, \frac{1}{2}u\bar{d}\} \leq \iota < \frac{(k+o)u\bar{d}}{u+o}$. In this case, $q_p^* = \bar{d} - \frac{\iota}{u}$ and $q_e^* = \frac{1}{2}\bar{d}$. We have $q_p^* \leq q_e^*$ since $\iota \geq \frac{1}{2}u\bar{d}$.

(6) $k \geq \frac{u}{2}$ and $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} \leq \iota < \frac{1}{2}u\bar{d}$. In this case, $q_p^* = \bar{d} - \frac{\iota}{u}$ and $q_e^* = \frac{\iota}{u}$. We have $q_p^* > q_e^*$ since $\iota < \frac{1}{2}u\bar{d}$.

(7) $k \geq \frac{u}{2}$ and $\frac{(k+o)u\bar{d}}{u+o} \leq \iota < \frac{u}{\bar{d}}$. In this case, $q_p^* = \frac{(u-k)\bar{d}}{u+o}$ and $q_e^* = \frac{1}{2}\bar{d}$. Suppose $\frac{(u-k)\bar{d}}{u+o} > \frac{1}{2}\bar{d}$, we must have $k < \frac{u-o}{2}$. However, since $k \geq \frac{u}{2}$, we arrive at a contradiction. We must have $q_p^* < q_e^*$.

Case: $u \leq 2o$ (Figure A.4b).

In this case, we have all sub-regions that we examined in the case of $u > 2o$ except for (2) and (6). The discussions of (5) and (7) are the same as the respective cases under $u > 2o$. Therefore, in what follows, it is sufficient to only specify the derivations for cases (1), (3), and (4) under $u \leq 2o$.

(1) Since $k \leq \frac{u}{2} < o$, we always have $q_p^* < q_e^*$.

(3) Since $u - o < \frac{u}{2}$ and $k \geq \frac{u}{2}$, we have $k > u - o$ and $q_p^* < q_e^*$.

(4) Since $k > u - o$, $(\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d} - \frac{1}{2}(o+k)\bar{d} = -\frac{d(k+o)(k+o-u)}{2(k+o+u)} < 0$. Therefore, $\forall \iota \in (\frac{1}{2}u\bar{d}, [\frac{1}{k+o} + \frac{1}{u})^{-1}\bar{d})$, $\iota < \frac{1}{2}(o+k)\bar{d}$ and $q_p^* < q_e^*$.

To sum up the above: $q_p^* > q_e^*$ if and only if: (i) $o < 2u$; (ii) $k < \bar{k}^\omega = u - o$; and (iii) $k > \underline{k}^\omega = \max\{o, \underline{k}_e(o)\}$. Figure A.5a provides the comparison of ω_e^* and ω_p^* for different diner's unit overage costs, buffet operator's unit production costs, and a fixed self-service cost. Figure A.5b further compares ω_e^* and ω_p^* by taking a vertical slice of Figure A.5a at o

where $\underline{k}_e(o) < o$. □

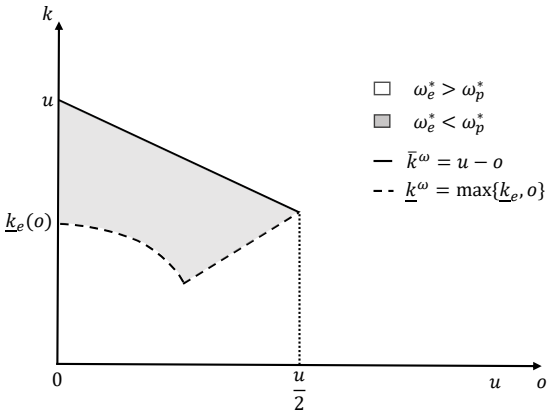

(a) Plate waste comparison across combinations of pre-education unit overage costs (o) (x-axis) and unit production costs (k) (y-axis).

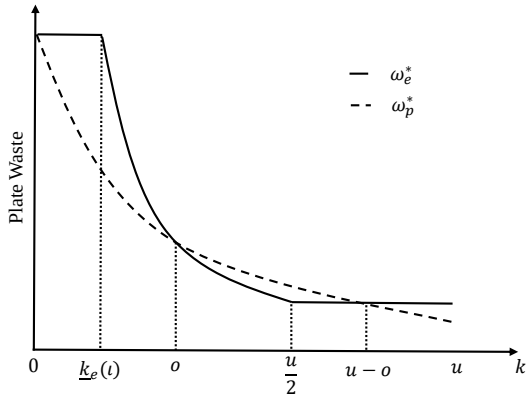

(b) Plate waste comparison as a function of the unit production cost (k) (x-axis), for a fixed pre-education unit overage cost (o).

Figure A.5: Comparison of per-diner expected plate waste under the optimal flat-rate penalty fee (ω_p^*) and the optimal education level (ω_e^*).

Note: In panel (a), the shaded gray area indicates where education is more effective in reducing plate waste ($\omega_e^* < \omega_p^*$). Analyses and definitions of terms and curves are provided in the proof of Theorem 2.14

A.3 Formal Analyses & Proofs for Model Extensions in Section 2.6

For each model extension covered in Section 2.6 of the main paper, this section provides (i) additional nuances of the formal model setup, and (ii) analyses, formal statements, and proofs of the key results under the respective extension.

A.3.1 Partial Demand Learning (Model Extension in Section 2.6.1)

Let the size of the first plate be denoted by q_d. Under partial demand learning: if a diner is still hungry after the first round, i.e., $q_d < D$, the amount of food she has to consume to satisfy her demand is the conditional distribution $D' := D|D > q_d$. Specifically, her residual demand D' is distributed uniformly between 0 and $\bar{d} - q_d$. A diner minimizes the expected

total cost of dining in the buffet given by:

$$g_d(\varphi_d, q_d) = \varphi_d + \iota + \mathbb{E}[o(q_d - D)^+] + \mathbb{P}(D > q) g_{d,2}(q_d), \tag{A.58}$$

where $g_{d,2}(q_d)$ is the expected total cost when the diner is still hungry after the first round of consumption. When a diner is still hungry, she chooses the minimum between the expected underage cost of quitting the buffet directly and the lowest expected cost she can achieve by refilling, i.e.,

$$g_{d,2}(q) = \min\{u\mathbb{E}_{D'}[D'], \min_{0<q'\leq\bar{d}-q} \iota + \mathbb{E}_{D'}[u(D'-q')^+ + o(q'-D')^+]\}. \tag{A.59}$$

If the buffet operator charges a flat-rate penalty fee $v_{d,p}$ on a diner whenever there is plate waste, the expected total cost of dining in the buffet and the expected total cost of being hungry after the first consumption round are given respectively by

$$\begin{aligned} g_{d,p}(\varphi_{d,p}, v_{d,p}, q_{d,p}) &= \varphi_{d.p} + \iota + v_{d,p}\mathbb{P}(D < q_{d,p}) + \mathbb{E}[o(q_{d,p} - D)^+] \\ &\quad + \mathbb{P}(D > q_{d,p}) g_{d,p,2}(v_{d,p}, q_{d,p}). \end{aligned} \tag{A.60}$$

$$\begin{aligned} g_{d,p,2}(v_{d,p}, q) = \min\Bigg\{ & u\mathbb{E}_{D'}[D'], \\ & \min_{0<q'\leq\bar{d}-q} \iota + v_{d,p}\mathbb{P}(D' < q') + \mathbb{E}_{D'}[u(D'-q')^+ + o(q'-D')^+] \Bigg\}. \end{aligned} \tag{A.61}$$

Clearly, Equations A.58 and A.59 are the special cases of Equations A.60 and A.61 with $v_{d,p} = 0$. Therefore, we first derive equilibrium results for the general case, where diners under partial demand learning and the flat-rate penalty. Specifically, in what follows, for any term x, we connect the case without a penalty and the general case with a penalty by $x_d = x_{d,p}|_{v_{d,p}=0}$.

A.3.1.1 Equilibirum Analysis under Partial Demand Learning (with Flat-rate Penalty)

Lemma A.9. *With partial demand learning and a flat-rate penalty, there exists a threshold on the self-service cost such that, when the self-service cost is above the threshold, a diner always refills if she is still hungry after the first consumption round. Otherwise, she never refills. Formally, for equations A.58 and A.59, $\exists\zeta_d$ such that if $\iota < \zeta_d$, $q^*_{d,p,2} > 0$. Otherwise, $q^*_{d,p,2} = 0$.*

Proof: We first solve for the optimal refill decision. Define

$$h_{d,2}(v_{d,p}, q_{d,p}, q_{d,p,2}) = \iota + v_{d,p}\mathbb{P}(D' < q_{d,p,2}) + \mathbb{E}_{D'}[u(D' - q_{d,p,2})^+ + o(q_{d,p,2} - D')^+]\}$$
$$= \iota + v_{d,p}\frac{q_{d,p,2}}{\bar{d} - q_{d,p}} + o\frac{q_{d,p,2}}{\bar{d} - q_{d,p}}\frac{q_{d,p,2}}{2} + u(1 - \frac{q_{d,p,2}}{\bar{d} - q_{d,p}})\frac{\bar{d} - q_{d,p} - q_{d,p,2}}{2},$$

We have

$$q^*_{d,p,2}(v_{d,p}, q_{d,p}) = \arg\min_{0 \le q_{d,p,2} \le \bar{d} - q_{d,p}} h_{d,2}(v_{d,p}, q_{d,p}, q_{d,p,2}) = \frac{u(\bar{d} - q_{d,p}) - v_{d,p}}{o + u}. \tag{A.62}$$

$$h_{d,2}(v_{d,p}, q_{d,p}, q^*_{d,p,2}(v_{d,p}, q_{d,p})) = \iota - \frac{\left(u\left(q_{d,p} - \vec{d}\right) + v_{d,p}\right)^2}{2(o+u)\left(\bar{d} - q_{d,p}\right)} + \frac{1}{2}u\left(\bar{d} - q_{d,p}\right). \tag{A.63}$$

(1) If $v_{d,p} \ge u(\bar{d} - q_{d,p}) \Leftrightarrow q_{d,p} \ge \frac{u\bar{d} - v_{d,p}}{u}$, $h_{d,2}$ strictly increases in $q_{d,p,2}$ and $q^*_{d,p,2}(v_{d,p}, q_{d,p}) = 0$. In this case, a diner prefers not to refill.
(2) If $v_{d,p} < u(\bar{d} - q_{d,p})$, by Eq. A.63,

$$h_{d,2}(v_{d,p}, q_{d,p}, q^*_{d,p,2}(v_{d,p}, q_{d,p})) > u\mathbb{E}[D'] = u\frac{\bar{d} - q_{d,p}}{2} \Leftrightarrow \iota > \iota_d = \frac{\left(u\left(q_{d,p} - \vec{d}\right) + v_{d,p}\right)^2}{2(o+u)\left(\bar{d} - q_{d,p}\right)}. \tag{A.64}$$

Since $\frac{d\iota_d}{dq_{d,p}} = \frac{v^2_{d,p} - u^2(\bar{d} - q_{d,p})^2}{2(\bar{d} - q_{d,p})^2(o+u)} < 0$, ι_d strictly decreases in $q_{d,p}$. This implies:
(2-i) If $\iota \ge \iota_d|_{q_{d,p}=0} = \frac{(u\bar{d} - v_{d,p})^2}{2(o+u)\bar{d}}$, a diner never refills.
(2-ii) If $\iota < \frac{(u\bar{d} - v_{d,p})^2}{2(o+u)\bar{d}}$, there exists a $\hat{q}_{d,p}(v_{d,p}) < \bar{d} - \frac{v_{d,p}}{u}$, such that a diner refills if and only if $q_{d,p} < \hat{q}_{d,p}(v_{d,p})$. Otherwise, she does not refill. $\hat{q}_{d,p}(v_{d,p})$ solves $\iota = \iota_d$.
Next, we solve for the optimal first plate size $q^*_{d,p}$.
(1) If $q_{d,p} \ge \hat{q}_{d,p}(v_{d,p})$ or $\iota \ge \frac{(u\bar{d} - v_{d,p})^2}{2(o+u)\bar{d}}$, the diner never refills. Similar to the baseline model, $q^*_{d,p}(v_{d,p}) = \frac{u\bar{d} - v_{d,p}}{u+o}$ (Lemma A.4). To check whether $q^*_{d,p}(v_{d,p}) \ge \hat{q}_{d,p}(v_{d,p})$ holds, consider

$$\iota - \iota_d(q^*_{d,p}(v_{d,p})) > 0 \Leftrightarrow \iota \ge \iota^1_d = \frac{o^2\left(v_{d,p} - u\vec{d}\right)^2}{2(o+u)^2\left(o\bar{d} + v_{d,p}\right)}. \tag{A.65}$$

Since $\iota^1_d < \frac{(u\bar{d} - v_{d,p})^2}{2(o+u)\bar{d}} < \iota$, hence $q^*_{d,p}(v_{d,p}) \ge \hat{q}_{d,p}(v_{d,p})$, indeed, holds true.
(2) If $\iota \ge \frac{(u\bar{d} - v_{d,p})^2}{2(o+u)\bar{d}}$ and $q_{d,p} < \hat{q}_{d,p}(v_{d,p})$, by Equations A.60 and A.63, and the first order

condition, we have

$$q^*_{d,p}(v_{d,p}) = \arg\min_{q_{d,p}} g_{d,p}(\varphi_{d,p}, v_{d,p}, q_{d,p}) = \frac{(o+u)\iota + ou\bar{d} - ov_{d,p}}{o^2+2ou}, \tag{A.66}$$

$$\begin{aligned} &g_{d,p}(\varphi_{d,p}, v_{d,p}, q^*_{d,p}(v_{d,p})) \\ &= \varphi_d + \frac{o\left(\bar{d}\left(ou\bar{d} + 2\iota(o+u)\right) + 2v_{d,p}\left(2u\bar{d}+\iota\right) - 2v^2_{d,p}\right) - \iota^2(o+u)}{2o\bar{d}(o+2u)} + \iota. \end{aligned} \tag{A.67}$$

For the condition $q^*_{d,p}(v_{d,p}) < \hat{q}_{d,p}(v_{d,p})$ to hold, we must have $\iota - \iota_d(q^*_{d,p}(v_{d,p})) < 0$, which implies

$$\begin{aligned} \iota < \iota^2_d = &\frac{o\left(-(o+2u)\sqrt{o\left(2\bar{d}(o+u)^2 v_{d,p} + o\bar{d}^2(o+u)^2 - (o+2u)v^2_{d,p}\right)}\right)}{(o+u)\left(2o^2+4ou+u^2\right)} \\ &+ \frac{o\left(\bar{d}(o+u)^3 + \left(o^2+ou-u^2\right)v_{d,p}\right)}{(o+u)\left(2o^2+4ou+u^2\right)}. \end{aligned} \tag{A.68}$$

Note: ι^2_d always exists since $\frac{\bar{d}(o+u)\left(\sqrt{2o^2+4ou+u^2}+o+u\right)}{o+2u} > u\bar{d}$ and $v_{d,p} < u\bar{d}$. Moreover, by Eq. A.66, we have $\iota^1_d < \iota^2_d$ (Eq. A.65 and A.68) since

$$\begin{aligned} &\iota - \iota_d(q^*_{d,p}(\iota^2_p)) \\ &= -\frac{\left(2o^2+4ou+u^2\right)\left(v_{d,p}-u\bar{d}\right)^2\left(ou\bar{d}+(o+2u)v_{d,p}\right)^2}{4(o+u)^2(o+2u)\left(o\bar{d}+v_{d,p}\right)} \\ &\cdot \frac{1}{\left(2\bar{d}\left(2o^2+4ou+u^2\right)v_{d,p} + o\bar{d}^2\left(2o^2+4ou+u^2\right) + (o+2u)v^2_{d,p}\right)} < 0. \end{aligned}$$

(3) For $\iota \in (\iota^1_d, \iota^2_d]$, a diner decides whether to refill by comparing $g_{d,p}(\varphi_{d,p}, v_{d,p}, q^*_{d,p}(v_{d,p}))$ (Eq. A.67) and the expected total cost of never refilling, given by Equations 2.7 and A.20. Specifically,

$$\begin{aligned} &\left(\varphi_d + \iota + \frac{ou\bar{d}^2 + 2u\bar{d}v_{d,p} - v^2_{d,p}}{2o\bar{d}+2u\bar{d}}\right) - g_{d,p}(\varphi_{d,p}, v_{d,p}, q^*_{d,p}(v_{d,p})) \\ &= \frac{o^2u^2\bar{d}^2 - 2o\bar{d}\left(ouv_{d,p} + \iota(o+u)^2\right) + \left(\iota(o+u) - ov_{d,p}\right)^2}{2o\bar{d}(o+u)(o+2u)} \le 0 \\ &\Leftrightarrow \iota \ge \zeta_d = o\left(\frac{v_{d,p} - \sqrt{\bar{d}(o+2u)\left(o\bar{d}+2v_{d,p}\right)}}{(o+u)} + \bar{d}\right). \end{aligned} \tag{A.69}$$

In sum, a diner's decision is given by:

$$(q^*_{d,p}(v_{d,p}), q^*_{d,p,2}(v_{d,p})) = \begin{cases} (\frac{u\bar{d}-v_{d,p}}{o+u}, 0) & \text{if } \iota \geq \zeta_d \\ (\frac{(o+u)\iota+ou\bar{d}-ov_{d,p}}{o^2+2ou}, -\frac{-ou\bar{d}+ov_{d,p}+\iota u}{o^2+2ou}) & \text{if } \iota < \zeta_d \end{cases}, \tag{A.70}$$

where we used Eq. A.62 to derive $q^*_{d,p,2}(v_{d,p})$ for $\iota < \zeta_d$. □

A.3.1.2 Impact of Improved Convenience under Partial Demand Learning

In this section, we use the shorthand for $q^*_{d,p}(v_{d,p})$ and $q^*_{d,p,2}(v_{d,p})$ as $q^*_{d,p}$ and $q^*_{d,p,2}$, respectively. We also explicitly write the expected plate waste, the total amount of food taken, and non-monetary cost as functions of the penalty fee unless otherwise stated.

Lemma A.10. *Under partial demand learning, Theorem 2.3 continues to hold.*

Proof: Due to partial demand learning, waste can occur in both the first and second consumption rounds.

$$\begin{aligned} \omega^*_{d,p} &= \mathbb{E}[(q^*_{d,p} - D)^+] + P(q^*_{d,p} > D)\mathbb{E}[(q^*_{d,p,2}(q^*_{d,p}) - D')^+] \\ &= \frac{q^*_{d,p}}{\bar{d}}\frac{q^*_{d,p}}{2} + \frac{\bar{d}-q^*_{d,p}}{\bar{d}}\frac{q^*_{d,p,2}}{\bar{d}-q^*_{d,p}}\frac{q^*_{d,p,2}}{2} = \frac{(q^*_{d,p})^2 + (q^*_{d,p,2})^2}{2\bar{d}}. \end{aligned} \tag{A.71}$$

By Eq. A.70, when $\iota \geq \zeta_d$, it is easy to verify that $\frac{d\omega^*_{d,p}}{d\iota} = 0$. When $\iota < \zeta_d$,

$$\begin{aligned} \omega^*_{d,p} &= \frac{\left(ou\bar{d} + \iota(o+u) - ov_{d,p}\right)^2 + \left(-ou\bar{d} + ov_{d,p} + \iota u\right)^2}{2\bar{d}\left(o^2+2ou\right)^2}, \\ \frac{d\omega^*_{d,p}}{d\iota} &= \frac{o^2\left(u\bar{d} - v_{d,p}\right) + \iota\left((o+u)^2 + u^2\right)}{o^2\bar{d}(o+2u)^2} > 0. \end{aligned} \tag{A.72}$$

where the last inequality is due to $v_{d,p} \leq u\bar{d}$ (the discussion can be found in the note of the proof of Lemma A.4). Theorem 2.3 under partial demand learning is a special case of $v_{d,p} = 0$. □Now we define the buffet operator's expected profit when diners have partial demand learning.

(1) When $\iota \geq \zeta_d$, a diner never refills, and $\pi^*_{d,p} = \pi^*_p$ given in Eq. A.51 where Δ^*_p, τ^*_p are given by Equations A.45 and A.46 respectively.

(2) When $\iota < \zeta_d$, a diner has a positive chance of refilling. By definition,

$$\begin{aligned}\tau_{d,p} &= q^*_{d,p} + \mathbb{P}(q^*_{d,p} < D) q^*_{d,p,2} = q^*_{d,p} + \frac{\bar{d} - q^*_{d,p}}{\bar{d}} q^*_{d,p,2} \\ &= \frac{o^2\left(\bar{d}(o+u)\left(\iota - 2v_{d,p}\right) + u\bar{d}^2(2o+3u) + v_{d,p}\left(\iota - v_{d,p}\right)\right) + \iota^2 u(o+u)}{o^2\bar{d}(o+2u)^2}. \end{aligned} \tag{A.73}$$

$$\begin{aligned}\Delta_{d,p} &= g_{d,p}(\varphi_{d,p}, v_{d,p}, q^*_{d,p}) - \varphi_{d,p} - \mu_{d,p}(q^*_{d,p}, q^*_{d,p,2}, v_{d,p}) \\ &= \iota + o\frac{(q^*_{d,p})^2}{2\bar{d}} + \frac{\bar{d} - q^*_{d,p}}{\bar{d}}\left(\iota + o\frac{(q^*_{d,p,2})}{\bar{d}}\frac{q^*_{d,p,2}}{\bar{d} - q^*_{d,p}} + u\frac{\bar{d} - q^*_{d,p} - q^*_{d,p,2}}{2}\frac{(\bar{d} - q^*_{d,p}) - q^*_{d,p,2}}{\bar{d} - q^*_{d,p}}\right) \\ &= \frac{o\bar{d}\left(ou\bar{d} + 4\iota o + 6\iota u\right) + 2ov^2_{d,p} - (\iota^2(o+u))}{2o\bar{d}(o+2u)}. \end{aligned} \tag{A.74}$$

Here, $\mu_{d,p}(q^*_{d,p}, q^*_{d,p,2}, v_{d,p}) = v_{d,p}\left(\mathbb{P}(D < q^*_{d,p}) + \mathbb{P}(D > q^*_{d,p})\mathbb{P}(D' < q^*_{d,p,2}(q^*_{d,p}))\right)$ is the expected penalty collected per-meal. The buffet operator's expected profit is given by

$$\pi_{d,p}(\varphi_{d,p}, v_{d,p}) \equiv \pi_{d,p}(\varphi_{d,p}, v_{d,p}, q^*_{d,p}, q^*_{d,p,2}) = \frac{1}{\bar{c}_o}(\bar{c}_o - g_{d,p}(\varphi_{d,p}, v_{d,p}, q^*_{d,p}))(\varphi_{d,p} + \mu_{d,p} - k\tau_{d,p}). \tag{A.75}$$

When $v_{d,p} = 0$, similar to the base model, $\pi^*_d = \frac{1}{4\bar{c}_o}(\bar{c}_o - \Delta^*_d - \tau^*_d)^2$ (Eq. 2.5), where $\Delta^*_d = \Delta_{d,p}|_{v_{d,p}=0}$ and $\tau^*_d = \tau_{d,p}|_{v_{d,p}=0}$. □

Lemma A.11. *Under partial demand learning, Theorem 2.7 continues to hold if and only if the unit overage cost is sufficiently low. Formally, $\exists 0 < o_d \leq u$, such that: $d\pi^*_d/d\iota < 0$ if and only if $o < o_d$.*

Proof: (1) When $\iota \geq \zeta_d$, $\pi^*_d = \pi^*$ given in Eq. 2.5, where Δ^*, τ^* are given by Equations A.20 and A.21, respectively. By Theorem 2.7, $\frac{d\pi^*_d}{d\iota} < 0$. In this case, $o_d = u$.
(2) When $\iota < \zeta_d$, $\pi^*_d = \frac{1}{4\bar{c}_o}(\bar{c}_o - \Delta^*_d - \tau^*_d)^2$ by Eq. A.73, τ^*_d increases in ι. By A.74,

$$\begin{aligned}\frac{d\Delta^*_d}{d\iota} = \frac{d\Delta_{d,p}|_{v_p=0}}{d\iota} &= \frac{o\bar{d}(2o+3u) - \iota(o+u)}{o\bar{d}(o+2u)} > \frac{o\bar{d}(2o+3u) - \iota(o+u)}{o\bar{d}(o+2u)}\Big|_{\iota=\zeta_d} \\ &= \frac{o}{\sqrt{o(o+2u)}} + 1 > 0. \end{aligned}$$

Therefore, Δ^*_d increases in ι, which implies π^*_d decreases in ι when $\iota < \zeta_d$.

(3) When $\iota = \zeta_d$, $\pi^*_{d,p}$ is not continuous. By Equations A.46 and A.73,

$$\begin{aligned}\delta\tau^*_d &= \lim_{\iota\to\zeta_d^+}\tau^*_d - \lim_{\iota\to\zeta_d^-}\tau^*_d = -\frac{\Big(ou\bar{d}+\iota(o+u)\Big)\Big(o\bar{d}\,(o^2+ou-u^2)+\iota u(o+u)\Big)}{o^2\bar{d}(o+u)(o+2u)^2}|_{\iota=\zeta_d}\\ &= \bar{d}\left(\frac{\frac{2u^2}{o+u}+\sqrt{o(o+2u)}}{o+2u}-1\right) > 0 \Leftrightarrow o < (\sqrt{2}-1)u.\end{aligned}$$

When $o = o_d < (\sqrt{2}-1)u$, $\delta\tau^*_d > 0$. Moreover, since Δ^*_d is continuous at ζ_d by Eq. A.69, $\delta\pi^*_d = \lim_{\iota\to\zeta_d^+}\pi^*_d - \lim_{\iota\to\zeta_d^-}\pi^*_d < 0$. □

A.3.1.3 Performance Assessment of Flat-rate Penalty under Partial Demand Learning

Lemma A.12. *Under partial demand learning, Theorem 2.10 continues to hold.*

Proof: Comparing Eq. A.75 with Eq. A.34, it is easy to see that the buffet operator's optimal decision on the flat-rate entrance fee and the penalty fee can still be described by Equations A.37 and A.38. To show that Theorem 2.10 continues to hold, it is sufficient to show that $\exists v_{d,p}$, such that $\Delta_{d,p} + k\tau_{d,p}$ decreases.
(1) When $\iota \geq \zeta_d$, the analysis is similar to the base model. (2) When $\iota < \zeta_d$, by Equations A.73 and A.74,

$$\frac{d\Delta_{d,p}+k\tau_{d,p}}{dv_{d,p}}|_{v_{d,p}=0} = \frac{-2k\bar{d}(o+u)+\iota k+2v_{d,p}(-k+o+2u)}{\bar{d}(o+2u)^2}|_{v_{d,p}=0} < 0.$$

□**Note:** When $k < o+2u$, the second-order condition (for a unique interior solution) is satisfied and we have $v^*_{d,p} = \frac{k\left(\iota-2\bar{d}(o+u)\right)}{2(k-o-2u)}$ by the first-order condition. By Eq. A.72,

$$\frac{d\omega^*_{d,p}(v^*_{d,p})}{dk} = \frac{d\omega^*_{d,p}}{dv^*_{d,p}}\frac{dv^*_{d,p}}{dk} = \left(-\frac{2u\bar{d}+\iota-2v_{d,p}}{\bar{d}(o+2u)^2}\right)\left(\frac{(o+2u)(2\bar{d}(o+u)-\iota)}{2(-k+o+2u)^2}\right) < 0.$$

Therefore, Lemma 2.9 continues to hold.

A.3.2 Insatiable Diners (Model Extension in Section 2.6.2)

In this extension, we use subscript 'B' to refer to the measures that relate to the "insatiable diners". Based on Section 2.6.2, an insatiable diner's minimum expected total cost of dining

in the buffet is given by

$$g_B(\varphi_B, (q_t^*)_t) = \varphi_B + J_1(\beta\bar{d}),$$
$$\text{where} \quad J_t(D) = \min\left\{uD, \min_{0<q_t\leq\bar{d}} \iota + o(q_t - D)^+ + J_{t+1}(D - q_t)\cdot \mathbf{1}_{D\geq q_t}\right\}. \tag{A.76}$$

In Eq. A.76, $t-1 \geq 0$ stands for the number of times an insatiable diner refills, and D is the amount of food an insatiable diner still needs to fulfill her hunger. q_t^* is the amount of food the insatiable diner takes in the round t that minimizes the expected underage, overage cost, and the cost-to-go $J_{t+1}(D - q_t)$. In each period t, the insatiable diner chooses the minimum between paying the underage cost by leaving directly and the lowest cost she can achieve by taking another round of food. an insatiable diner will join if $\beta\bar{c}_o > g_B(\varphi_B, (q_t^*)_t)$. The buffet operator's profit consists of the profit from regular diners and insatiable diners, i.e.,

$$\pi_B(\varphi_B) \equiv \pi_B(\varphi_B, (q_t^*)_t) = (1-m)\pi(\varphi_B) + m(\varphi_B - k\sum_{t=1}^{\tau} q_t), \tag{A.77}$$

where $\tau_B = \sum_{t=1}^{\tau} q_t$ denotes the number of plates an insatiable diner takes before she leaves the buffet. Lastly, to avoid an unrealistic case, where a regular diner joins the buffet but takes no food, we assume $\iota < u\bar{d}$ (more discussion is available in the note of the derivations for Assumption 2.1).

A.3.2.1 Equilibrium Analysis and Impact of Improved Convenience with Insatiable Diners

Lemma A.13. *An insatiable diner always takes and eats the amount of food that she demands. Formally, $\tau_B^* = \beta\bar{d}$.*

Proof: Define $\delta q \in [0, \bar{d})$. We first prove that $J_t(n\bar{d} + \delta q) \geq n\iota$ by induction with respect to $n \in \mathbb{Z}_+$ (positive integers).
(Step 1) When $n = 1$, it is easy to show that $\tau \leq 2$. Denote the amount of food an insatiable diner takes in $t = \{1, 2\}$ as q_1 and q_2. If $q_1 + q_2 \neq 0$, then the insatiable diner takes in total a non-zero amount of food and $J_t(n\bar{d} + \delta q) \geq \iota$. If $q_1 + q_2 = 0$, by the assumption of $\iota < u\bar{d}$,

$$J_t(n\bar{d} + \delta q) = J_t(\bar{d} + \delta q) = (\bar{d} + \delta q)u \geq \iota.$$

(Step 2) Suppose $J_t(n\bar{d} + \delta q) \geq n\iota$ holds for $n \in \mathbb{Z}_+$ and $n > 1$.
(Step 3) For $n + 1$, it is easy to show that $\tau \leq n + 2$. Suppose an insatiable diner takes q_1

on her first plate. If $q_1 = 0$ (the diner quits directly), by the assumption $\iota < u\bar{d}$,

$$J_t\left((n+1)\bar{d}+\delta q\right) = \left((n+1)\bar{d}+\delta q\right)u \geq (n+1)\iota.$$

If $q_1 \neq 0$, by the definition of J_t (Eq. A.76) and $q_1 \leq \bar{d}$,

$$J_t((n+1)\bar{d}+\delta q) > \iota + J_{t+1}\left((n+1)\bar{d}+\delta q - q_1\right) > \iota + J_{t+1}\left(n\bar{d}+\delta q\right) \geq (n+1)\iota.$$

Hence, we proved by induction, that we have $J_t(n\bar{d}+\delta q) \geq n\iota$ for any $n \in \mathbb{Z}_+$. Therefore, $J_1(\beta\bar{d})$ (Eq. A.76) satisfies $J_1(\beta\bar{d}) \geq \beta\iota$. Moreover, it is easy to see that, when $q_1 = q_2 = \cdots = q_\beta = \bar{d}$, $J_1(\beta\bar{d}) = \beta\iota$. Hence, $\tau = \beta$ and $\tau_B^* = \beta\bar{d}$. □

Lemma A.14. *A flat-rate buffet facing insatiable diners is profitable only if the proportion of insatiable diners is sufficiently small. Formally, $\forall\beta$, $\exists\hat{m}(\beta) > 0$, such that, if $m < \hat{m}(\beta)$, then $\pi_B^* > 0$.*

Proof: By Eq. A.77 and Lemma A.13, the first-order condition $\frac{d\pi_B(\varphi)}{d\varphi}$ gives $\varphi_B^* = \varphi^* + \frac{m}{2(1-m)}\bar{c}_o$. Hence,

$$\begin{aligned}\pi_B(\varphi_B^*) = \frac{1-m}{\bar{c}_o}\frac{1}{4}&\left(\bar{c}_o - \Delta^* - k\tau^* - \frac{m\bar{c}_o}{1-m}\right)\left(\bar{c}_o - \Delta^* - k\tau^* + \frac{m\bar{c}_o}{1-m}\right)\\ &+ \frac{m}{2}\left(\bar{c}_o - \Delta^* - k\tau^* + \frac{m\bar{c}_o}{1-m}\right) - mk\beta\bar{d},\end{aligned} \tag{A.78}$$

where Δ^* and τ^* are given by Equations A.22 and A.24 when $\iota < \zeta$, and Equations A.20 and A.21 when $\iota \geq \zeta$. For the buffet to operate: First, the proportion of regular diners cannot be negative, i.e.,

$$\bar{c}_o - \Delta^* - k\tau^* - \frac{m\bar{c}_o}{1-m} > 0 \Leftrightarrow m < m_1 = \frac{\bar{c}_o - (\Delta^* + k\tau^*)}{2\bar{c}_o - (\Delta^* + k\tau^*)} < \frac{1}{2} \wedge \bar{c}_o > \Delta^* + k\tau^*.$$

Second, there must exist some ι such that the buffet is profitable. Since the profit from insatiable diners ($m(\varphi - k\beta\bar{d})$) is not a function of ι, by the Envelope Theorem and Eq. A.77,

$$\frac{d\pi_B^*}{d\iota} = \frac{\partial\pi_B^*}{\partial\iota} = m\frac{\partial\pi^*}{\partial\iota} = m\frac{d\pi^*}{d\iota} < 0. \tag{A.79}$$

Therefore, π_B^* strictly decreases in ι and we require $\pi_B^*|_{\iota=0} > 0$, i.e.,

$$\begin{aligned}
&\pi_B(\varphi_B^*)|_{\iota=0} > 0 \Leftrightarrow 0 \leq m < m_2(\beta) \\
&= \begin{cases} \frac{\bar{d}k - 2\bar{c}_o}{\bar{d}k - 4\bar{c}_o} & \text{if } 0 < k < \frac{\bar{c}_o}{\bar{d}} \wedge \beta \leq \frac{\bar{c}_o}{\bar{d}k} \\ \frac{\bar{d}k\left((8\beta-6)\bar{c}_o + \bar{d}k\right) - 4\sqrt{2}\bar{c}_o\sqrt{(2\beta-1)\left(\beta\bar{d}k - \bar{c}_o\right)\bar{d}k}}{8\bar{d}k(2\beta-1)\bar{c}_o + \bar{d}k} & \text{if } \left(\left(\beta > \frac{\bar{c}_o}{\bar{d}k} \wedge k \leq \frac{\bar{c}_o}{\bar{d}}\right) \vee \left(\frac{\bar{c}_o}{\bar{d}} \leq k < \frac{2\bar{c}_o}{\bar{d}}\right)\right) \end{cases}.
\end{aligned}$$

Define $\hat{m}(\beta) = \min\{m_1, m_2(\beta)\}$ completes the proof. □

Lemma A.15. *When the proportion of insatiable diners is sufficiently small, Therorems 2.3 and 2.7 hold.*

Proof: By Lemma A.13, only regular diners generate plate waste. Therefore, the per-capita plate waste is given by $\omega_B^* = (1-m)\omega^*$ (Eq. A.13), and Therorem 2.3 holds. Theorem 2.7 holds due to Lemma A.14 and Eq. A.79.

□

A.3.2.2 Performance Assessment of Flat-rate Penalty and Education with Insatiable Diners

Lemma A.16. *When the proportion of insatiable diners is sufficiently small, Theorem 2.10 continues to hold.*

Proof: First, insatiable diners are never penalized since they do not generate plate waste (Lemma A.13). Second, the effect of the penalty on regular diners is the same as in Section 2.5.1. Therefore, under a flat-rate penalty, the buffet operator's expected profit is given by

$$\pi_{B,p}(\varphi_{B,p}, v_{B,p}) = (1-m)\pi_p(\varphi_{B,p}, v_{B,p}) + m(\varphi_{B,p} - k\beta\bar{d}),$$

where π_p is given by Eq. 2.8. To show that Theorem 2.10 holds, it is sufficient to show that there exists $v_{B,p}$, such that $\pi_{B,p}(\varphi_{B,p}^*, v_{B,p}) > \pi_B^*$. By the first-order condition with respect to $\varphi_{B,p}$, we obtain $\varphi_{B,p}^* = \varphi_p^* + \frac{1}{2}\frac{m}{1-m}\bar{c}_o$. By the first order condition with respect to $v_{B,p}$,

$$\begin{aligned}
&\frac{d\pi_{B,p}}{dv_{B,p}}|_{\varphi_{B,p}=\varphi_{B,p}^*, v_{B,p}=0} \\
&= -\frac{d\Delta_p^* + k\tau_p^*}{dv_{B,p}}|_{v_{B,p}=0} - \frac{m\bar{c}_o}{(1-m)(\bar{c}_o - k\tau^* - \Delta^*)}\frac{d(2\mu_p - k\tau_p^* + \Delta_p^*)}{dv_{B,p}}|_{v_{B,p}=0}. \qquad \text{(A.80)}
\end{aligned}$$

By Equations A.42 and A.47, $\frac{d\Delta_p^*+k\tau_p^*}{dv_{B,p}}|_{v_{B,p}=0} < 0$. Moreover, since $\frac{d(2\mu_p-k\tau_p^*+\Delta_p^*)}{dv_{B,p}}|_{v_{B,p}=0}$ is finite and $\frac{m}{1-m} \in (0,\infty)$, there exists $m_p > 0$, such that Eq. A.80 remains negative. □

Lemma A.17. *When the proportion of insatiable diners is sufficiently small, Theorem 2.13 continues to hold.*

Proof: By Eq. A.77, Lemma A.13, and Eq. A.78, we have

$$\pi_B^* - \pi_B^*|_{o=o_{B,e}} = \frac{(\Delta_e^* - \Delta^* - k(\tau^* - \tau_e^*))\,(2\bar{c}_o - (1-m)(\Delta^* + k(\tau^* + \tau_e^*) + \Delta_e^*))}{4\bar{c}_o},$$

where Δ^*, Δ_e^*, τ^*, and τ_e^* are defined Equations A.22, A.20, A.24, and A.21. Since $\bar{c}_o > \Delta^* + k\tau^*$ and $\bar{c}_o > \Delta_e^* + k\tau_e^*$, we have $2\bar{c}_o - (1-m)(\Delta^* + k(\tau^* + \tau_e^*) + \Delta_e^*) > 0$ and the sign of $\pi_B^* - \pi_B^*|_{o=o_{B,e}}$ fully depends on $\Delta_e^* - \Delta^* - k(\tau^* - \tau_e^*)$. By Eq. A.56, to prove Theorem 2.13 continues to hold, we only need to show that there exists some k such that $\pi_B^* - \pi_B^*|_{o=o_{B,e}} < 0$. Similar to Eq. A.80, the first-order derivative of π_B with respect to $o_{B,e}$ is given by

$$\frac{d\pi_{B,e}}{do_{B,e}}|_{\varphi_{B,e}=\varphi_{B,e}^*, o_{B,e}=o} = -\frac{d\Delta_e^* + k\tau_e^*}{do_{B,e}}|_{o_{B,e}=o} - \frac{m\bar{c}_o}{(1-m)(\bar{c}_o - k\tau^* - \Delta^*)} \frac{d(-k\tau_e^* + \Delta_e^*)}{do_{B,e}}|_{o_{B,e}=o}. \tag{A.81}$$

By Theorem 2.13, when $k > k_e(o_e)$, $\frac{d\Delta_e^*+k\tau_e^*}{do_{B,e}}|_{o_{B,e}=o} < 0$. Since $\frac{d(-k\tau_e^*+\Delta_e^*)}{do_{B,e}}|_{o_{B,e}=o}$ is finite and $\frac{m}{1-m} \in (0,\infty)$, there exists $m_e > 0$ such that Eq. A.81 remains negative. □

A.3.3 Environmentally Conscious Diners (Model Extension in Section 2.6.3)

In what follows, we use subscript G to denote measures related to environmentally conscious diners.

Lemma A.18. *When diners are environmentally conscious, i.e., $o > u$, Theorem 2.3 continues to hold. Theorem 2.7 continues to hold if and only if the unit production cost is sufficiently small ($k < k_G^\pi$).*

Proof: When diners are environmentally conscious, the equilibrium size of the first plate q_G^* has the same expression as that of regular diners (q^*). Similarly, ζ_G, τ_G^*, Δ_G^*, ω_G^*, φ_G^*, and π_G^* have the same expression as ζ, τ^*, Δ^*, ω^*, φ^*, and π^* respectively. Therefore, Theorem 2.3 continues to hold. By Equations A.17, A.22, A.24, A.20 and A.21, when $o > u$,

$$\frac{d\pi_G^*}{d\iota} < 0 \Leftrightarrow k \geq k_G^\pi = \frac{ou}{o-u} \wedge \frac{2o^2u^2\left(\vec{d}\right)}{(o+u)(k(o-u)+ou)} < \iota < \zeta. \tag{A.82}$$

Therefore, if $k < k_G^{\pi}$, Theorem 2.7 continues to hold. □

Lemma A.19. *When diners are environmentally conscious, Theorems 2.10 and 2.13 continue to hold.*

Proof: All our definitions and analyses in Section 2.5.1 of the flat-rate penalty can be applied to environmentally conscious diners. In particular, the expression of the optimal penalty fee with environmentally conscious diners is the same as the base model (Lemma A.4). Therefore, Theorem 2.10 continues to hold.

Similarly, when the buffet operator educates environmentally conscious diners, Eq. A.56 continues to hold. Therefore, to prove Theorem 2.13 continues to hold, we only need to show that there exists some k, such that $\pi_{G,e}^* - \pi^* < 0$. Mimicking the proof of Appendix Lemma A.7, we first derive the optimal post-education unit overage cost.

By the proof of Appendix Lemma A.7, o_e^* minimizes $\Delta_{G,e}^* + k\tau_{G,e}^*$. The expression for $\Delta_{G,e}^* + k\tau_{G,e}^*$ with respect to the post-education unit overage cost $o_{G,e}$ has the same expression as a similar expression in the base model where education was applied to the diners who have $o < u$ (Equations A.17, A.22, A.24, A.20, and A.21). However, since environmentally conscious diners have $o_G > u$, the discussions of the sensitivity of $\Delta_{G,e}^* + k\tau_{G,e}^*$ with respect to $o_{G,e}$ is different from that of Figure A.6. Since $o > u$, we have $\frac{1}{2}u\bar{d} < \zeta_G$. We consider the following two cases:

(1) $\iota < \zeta_G$ (Figure A.6a and A.6b). For any $o_{G,e} > o$: if $k < \frac{o}{2}$, profit decreases; if $k \geq \frac{o}{2}$, profit increases for $k < \frac{o}{2}$ and decreases otherwise.

(2) $\iota \geq \zeta_G$ (Figure A.6c). If $k < \frac{u}{2}$, profit decreases for $o_{G,e} > o$; If $\frac{u}{2} \leq k < \frac{o}{2}$, profit increases for $o < o_{G,e} < \frac{u\iota}{u\bar{d}-\iota}$ and decreases for $\frac{u\iota}{u\bar{d}-\iota} < o_{G,e}$; If $k \geq \frac{o}{2}$, profit increases for $o < o_{G,e} < 2k$ and decreases if $o_{G,e} \geq 2k$. $o_{G,e}^* = \min\{\frac{u\iota}{u\bar{d}-\iota}, 2k\}$. Figure A.6 summarizes whether optimal education increases or reduces profit for different combinations of self-service and unit production costs. From (i) and (ii), for any $o_{G,e} > o$, there exists a large enough $k > \underline{k}_{G,e}(\iota)$, such that $\pi_{G,e^*} - \pi^* > 0$. Here, $\underline{k}_{G,e}(\iota) = \frac{o}{2}$ if $\iota < \zeta_G$ and $\underline{k}_{G,e}(\iota) = \frac{u}{2}$ if $\iota \geq \zeta_G$. Figure A.7a illustrates $\underline{k}_{G,e}(\iota)$. □

Lemma A.20. *When diners are environmentally conscious, Theorem 2.14 continues to hold.*

Proof: We skip the proof of $\pi_{G,p}^* > \pi_{G,e}^*$ as it is identical to the case where diners have $o < u$. To compare $\omega_{G,e}^*$ and $\omega_{G,p}^*$, we consider two cases:

(1) When $\iota < \zeta_G$, Figure A.5 can still be used to compare $\omega_{G,e}^*$ and $\omega_{G,p}^*$. $\omega_{G,e}^*$ and $\omega_{G,p}^*$ intersect at $k = o$. Let $\bar{o}_{G,e}$ be the maximum post-education unit overage cost. If $\bar{o}_{G,e} < 2o$, $\omega_{G,e}^*$ and $\omega_{G,p}^*$ never intersect and $\omega_{G,e}^* > \omega_{G,p}^*$. If $\bar{o}_{G,e} > 2o$, $\omega_{G,e}^*$ and $\omega_{G,p}^*$ intersect at $k = o$, and $\bar{o}_{G,e} - o = k$. $\omega_{G,e}^* > \omega_{G,p}^*$ for $k < o$ and $k > \bar{o}_{G,e} - o$, and $\omega_{G,e}^* < \omega_{G,p}^*$ otherwise.

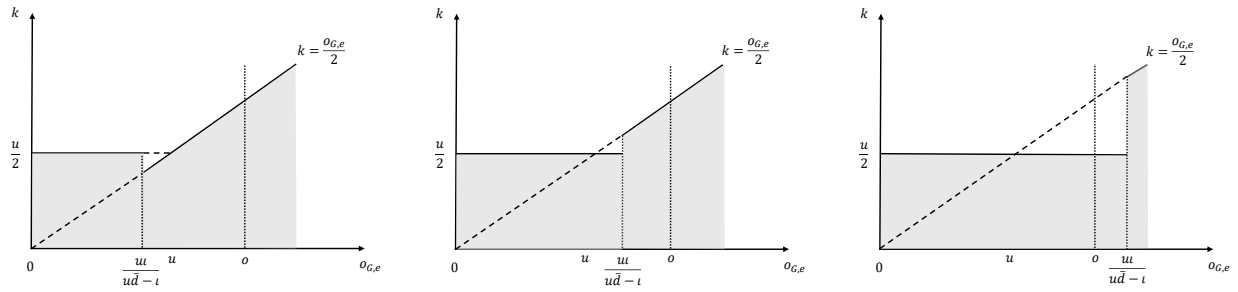

(a) $\iota < \frac{1}{2}u\bar{d}$ (b) $\frac{1}{2}u\bar{d} \leq \iota < \zeta_G$ (c) $\zeta_G \leq \iota$

Figure A.6: (When diners are environmentally conscious.) The marginal impact of consumer education on the buffet operator's expected profit for different combinations of values of this cost ($o_{G,e}$) and the values of the unit production cost (k), under small, moderate, and large self-service cost values.

Note: Profit decreases ($\frac{d\pi_{G,e}(o_{G,e})}{do_{G,e}} < 0$) in gray regions and increases ($\frac{d\pi_{G,e}(o_{G,e})}{o_{G,e}} > 0$) otherwise. The associated definitions and expressions are provided in the proof of Lemma A.19.

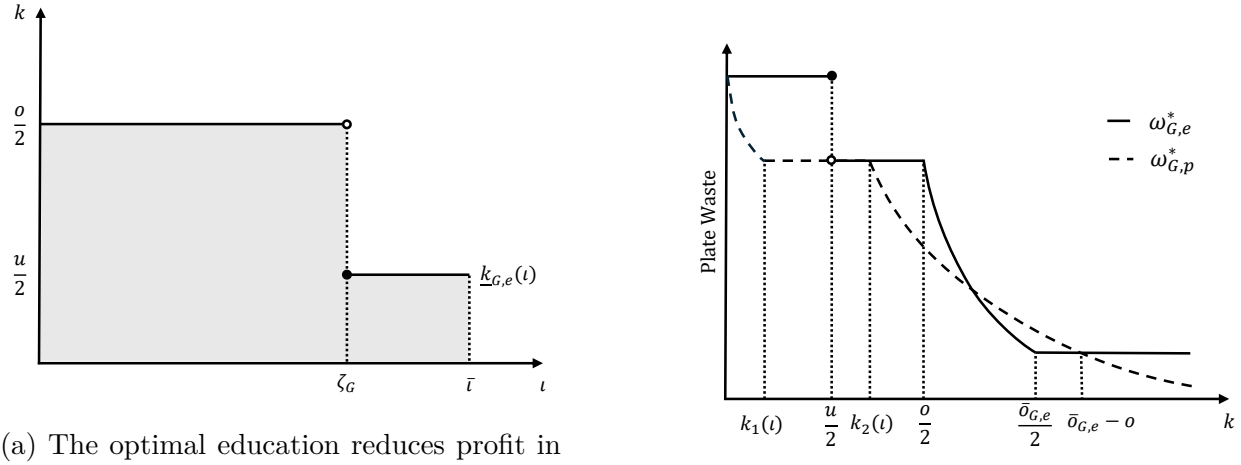

(a) The optimal education reduces profit in the gray region and increases profit in the white region, i.e., $\pi^*_{G,e} > \pi^*_G$ if and only if $k > \underline{k}_{G,e}$.

(b) Plate waste comparison as a function of the unit production cost k (drawn for $\iota > \zeta_G$, $o > k_2(\iota)$, and $\bar{o}_{G,e} > 2o$).

Figure A.7: (When diners are environmentally conscious.) Effectiveness of consumer education as an instrument and an example of waste comparison under the optimal flat-rate penalty ($\omega^*_{G,p}$) and the optimal post-education unit overage cost ($\omega^*_{G,e}$).

Note: Relevant analyses and definitions are provided in the proofs of Lemmas A.19 and A.20.

(2) When $\iota \geq \zeta_G$. We make the proof by offering two observations that are sufficient to determine the comparison. Interested readers can verify by going through all the subcases by comparing $k = \frac{u}{2}, \frac{o}{2}, k_1, k_2, o$ and $\bar{o}_{G,e}$. The observations are:

First, when $k_1(\iota) = (\frac{\iota}{\zeta_G} - 1)o < k < k_2(\iota) = \frac{(u+o)\iota - ou\bar{d}}{u\bar{d}-\iota}$, $q^*_{G,p} = \bar{d} - \frac{\iota}{u}$, and when $\frac{o}{2} < k < \frac{u}{2}$, $q^*_{G,e} = \bar{d}/(\frac{\iota u}{u\bar{d}-\iota}) = \bar{d} - \frac{\iota}{u}$. Figure A.7b provides an example of $k_1 < \frac{u}{2} < k_2 < \frac{o}{2} < o < \bar{o}_{G,e} - o$. Second, the relationship between $\omega^*_{G,e}$ and $\omega^*_{G,p}$ can only be changed at three points, i.e., $k = \frac{u}{2}$ (when $\frac{u}{2} < k_1$), $k = o$ (when $\max\{\frac{o}{2}, k_2\} < o < \frac{\bar{o}_{G,e}}{2}$) and $k = \bar{o}_{G,e} - o$ (when $\frac{\bar{o}_{G,e}}{2} > k_2$).

When $k < \frac{u}{2}, \omega^*_{G,e} > \omega^*_{G,p}$. When $k > \infty, \omega^*_{G,e} > \omega^*_{G,p}$. Hence, given that $\frac{u}{2} < o < \bar{o}_{G,e}$, $\omega^*_{G,e}$, $\omega^*_{G,p}$ cannot cross an odd number of times. We either have $\omega^*_{G,e} > \omega^*_{G,p}$ (they cross zero times) or $\omega^*_{G,e}$ and $\omega^*_{G,p}$ cross twice. The first crossing is either at $\frac{u}{2}$ or o. The second crossing is at $\bar{o}_{G,e} - o$. Lastly, $\underline{k}^{\omega}_G = \min\{\frac{u}{2}, o\}$ and $\bar{k}^{\omega}_G = \bar{o}_{G,e} - o$. □

A.3.4 Plate Size Restriction (Model Extension in Section 2.6.4)

In this extension, we consider plates to have a limited capacity for holding food. Specifically, we assume a diner using the plate can load at most M unit of food each time and diners are restricted to one plate each round of taking (or refilling) food. To be consistent with our assumption of the number of rounds of taking food in the baseline model, we assume $M > \frac{\bar{d}}{2}$ so that the maximum amount of food a diner can take in the buffet is larger than the maximum amount of food a diner would demand, i.e., $2M \geq \bar{d}$.

In section A.3.4.1, we analyze the decisions made by a diner at the buffet and incorporate the constraint in the plate size. In detail, the diner is subjected to the constraint on the maximum amount of food that can be taken in the first round, i.e., $q_M \leq M$, and she solves the following cost-minimization problem in the buffet

$$g_M(\varphi, q_M) = \varphi + \iota + \mathbb{E}_D[o \cdot (q_M - D)^+] + E_D[\min\{u \cdot (D-q)^+, \iota + u \cdot (D - M - q_M)^+\}] \quad \text{(A.83)}$$

where $(D - M - q_M)^+$ is the diner's residual hunger after refilling.

In section A.3.4.2, we examine our main results with respect to the self-service cost. We also look at how the choice of different plate sizes affects a diner's amount of food taken, the expected individual plate waste generated, and the restaurant's expected profit.

A.3.4.1 Equilibrium Analysis of Diners' Decisions under the Plate Size Restriction

Lemma A.21. *(i) With the plate size restriction, a diner's decision of whether to refill or not is described by Proposition 2.2-(ii).*

(ii) The amount of food the diner takes on her first plate and the expected total amount of

food a diner takes in the buffet weakly decreases in the plate size when the self-service cost is sufficiently high, and increases in the plate size when the self-service cost is sufficiently low. Otherwise, it is invariant to the plate size. Formally, when $\iota > \zeta$, $dq_M^*/dM \geq 0, d\tau_M^*/dM \geq 0$*; when* $\iota < o(\bar{d} - M)$, $dq_M^*/dM < 0, d\tau_M^*/dM < 0$*; otherwise,* $dq_M^*/dM = d\tau_M^*/dM = 0$.

Proof: Since $o < u$, by the assumption of $M \geq \frac{\bar{d}}{2}$, we have $\zeta = \frac{o}{o+u} u\bar{d} < u\frac{\bar{d}}{2} < uM$.
(1) When $\iota \geq uM$, since $u \cdot (D - q_M)^+ - u \cdot (D - M - q_M)^+ \leq uM < \iota$, a diner will never refill. In this case, Eq. A.83 is simplified to the expression of Eq. A.4 when $q > \bar{d} - \frac{\iota}{u}$. Therefore, $g_M(\varphi, q_M)$ is convex in q_M. By Eq. A.9 and the constraint in the plate size, when $\iota \geq uM$, $q_M^* = \min\{\frac{u}{u+o}\bar{d}, M\}$.
(2) When $\zeta \leq \iota < uM$, Eq. A.83 is again simplified to the expression of Eq. A.4 when $q_M > \bar{d} - \frac{\iota}{u}$. Therefore, $q_M^* = \min\{\frac{u}{u+o}\bar{d}, M\}$.
(3) When $\iota < \zeta$, define $d_M(q_M) = q_M + M$.
(3-i) If $d \geq d_M$, we have $\iota + u \cdot (d - M - q_M)^+ = \iota + u(d - q_M) - uM < u(d - q_M) = u(d - q_M)^+$. Therefore, a diner will refill with M unit of food.
(3-ii) If $q + \frac{\iota}{u} \leq d < d_M$, we have $\iota + u \cdot (d - M - q_M)^+ = \iota < u(d - q_M) = u(d - q_M)^+$. Therefore, a diner will refill with $d - q_M$ unit of food.
(3-iii) If $d < q_M + \frac{\iota}{u}$, $\iota + u \cdot (d - M - q_M)^+ > u(d - q_M)^+$ and the diner will not refill.
According to (3-i), (3-ii) and (3-iii), when $\iota < \zeta$, $g_M(\varphi, q_M)$ can be rewritten as,

$$\begin{aligned} g_M(\varphi, q_M) &= \varphi + \iota + o\int_0^{q_M} (q_M - d)^+ \frac{\mathrm{d}d}{\bar{d}} + u\int_{q_M}^{q_M + \frac{\iota}{u}} (d - q_M)^+ \frac{\mathrm{d}d}{\bar{d}} \\ &\quad + \int_{q_M + \frac{\iota}{u}}^{d_M(q_M)} \iota \frac{\mathrm{d}d}{\bar{d}} + \int_{d_M(q_M)}^{\bar{d}} \iota + u(d - M)\frac{\mathrm{d}d}{\bar{d}} \\ &= \varphi + \Delta(q_M) + u\frac{((\bar{d} - d_M(q_M))^+)^2}{2\bar{d}} \equiv \varphi + \Delta_M(q_M), \end{aligned} \tag{A.84}$$

where $\Delta(q_M)$ is given by Eq. A.22. By the first order condition,

$$\frac{dg_M(\varphi, q_M)}{dq_M} = \frac{d\Delta(q)}{dq} - \frac{u}{\bar{d}}(\bar{d} - d_M(q_M))^+ = 0.$$

By Eq. A.5, we have, $q_M^* = \frac{\iota}{o}$ if $\iota \geq o(\bar{d} - M)$ and $q_M^* = \frac{\iota + u(\bar{d} - M)}{u + o}$ if $\iota < o(\bar{d} - M)$. In sum, with a plate size restriction, the equilibrium amount of food a diner takes on her first plate

is given by

$$q_M^* = \begin{cases} \frac{\iota + u(\bar{d} - M)}{u+o} & \text{if } \iota < o(\bar{d} - M) \\ \frac{\iota}{o} & \text{if } o(\bar{d} - M) \leq \iota < \zeta \\ \min\{\frac{u}{u+o}\bar{d}, M\} & \text{if } \zeta \leq \iota \end{cases}. \tag{A.85}$$

Next, we compute the expected total amount of food taken by the diner. When $\iota \geq \zeta$, since a diner does not refill, $\tau_M^* = q_M^*$. When $\iota < \zeta$, a diner will refill (with $q_r = \min\{d - q_M^*, M\}$ unit of food) if and only if $d > q_M^* + \frac{\iota}{u}$. In particular, when $o(\bar{d} - M) \leq \iota < \zeta$, since $\bar{d} - q_M^* < M$, we have $q_r = d - q_M^*$ and τ_M^* has the same expression as Eq. A.24. When $\iota < o(\bar{d} - M)$, by the definition of τ_M^*, we have,

$$\begin{aligned} \tau_M^* &= \int_0^{Mq^* + \frac{\iota}{u}} q_M^* \frac{\mathrm{d}d}{\bar{d}} + \int_{q_M^* + \frac{\iota}{u}}^{\bar{d}} q_M^* + \min\{d - q_M^*, M\} \frac{\mathrm{d}d}{\bar{d}} \\ &= q_M^* + \int_{q_M^* + \frac{\iota}{u}}^{d_M(q_M^*)} (d - q_M^*) \frac{\mathrm{d}d}{\bar{d}} + \int_{d_M(q_M^*)}^{\bar{d}} M \frac{\mathrm{d}d}{\bar{d}}. \\ &= q_M^* + \frac{M - \frac{\iota}{u}}{\bar{d}} \frac{M + \frac{\iota}{u}}{2} + M \frac{\bar{d} - q_M^* - M}{\bar{d}} \\ &= \frac{2u^2\bar{d}\left(u\bar{d} + \iota + M(o-u)\right) + M^2u^2(u-o) - 2\iota M u^2 - \iota^2(o+u)}{2u^2\bar{d}(o+u)}. \\ \frac{d\tau_M^*}{dM} &= \frac{-\iota + (M - \bar{d})(u-o)}{\bar{d}(o+u)} < 0. \end{aligned} \tag{A.86}$$

Part (i) of Lemma A.21 follows from the discussion of cases (1), (2), and (3). Part (ii) of Lemma A.21 follows from equations A.85 and A.86. In particular, when ι switches from the range $[o(\bar{d} - M), \zeta)$ to $(0, o(\bar{d} - M'))$ as M decreases to M', q_M^* increases in M because $\frac{\iota + u(\bar{d} - M')}{u+o} - \frac{\iota}{o} = \frac{u(o(\bar{d} - M') - \iota)}{(u+o)o} > 0$. □

A.3.4.2 Impact of the Improved Convenience and Plate Size under the Plate Size Restriction

Lemma A.22. *When the size of the plate is restricted, Theorems 2.3 and 2.7 continue to hold.*

Proof: Theorem 2.3 under the plate size restriction continues to hold because of Eq. A.13 and part (ii) of Lemma A.21. Following a similar proof of the auxiliary Lemma A.3 and proposition 2.6, the optimal flat-rate entrance fee and the buffet operator's equilibrium profit are $\varphi^* = (\bar{c}_o - \Delta_M(q_M^*) + k\tau^*)/2$ and $\pi^* = (\bar{c}_o - \Delta_M(q_M^*) - k\tau^*)^2/(4\bar{c}_o)$ respectively.

Therefore, to verify Theorem 2.7 under the plate size restriction, it is sufficient to verify the monotonicity of $k\tau^* + \Delta_M(^*_M)$ with respect to the self-service cost. Specifically,

$$\frac{dk\tau^* + \Delta_M(q_M^*)}{d\iota} = k\frac{d\tau_M^*}{d\iota} + \frac{\partial \Delta_M(q_M)}{\partial \iota}|_{q_M=q_M^*} + \frac{\partial \Delta_M(q_M)}{\partial q_M}|_{q_M=q_M^*}\frac{dq_M^*}{d\iota}.$$

Observe that, when $q_M^* \neq M$, due to the Envelope Theorem, $\frac{\partial \Delta_M(q_M)}{\partial q_M}|_{q_M=q_M^*}$, when $q_M^* = M$, $\frac{dq_M^*}{d\iota}$. Therefore, by Eq. A.84,

$$\frac{dk\tau^* + \Delta_M(q_M^*)}{d\iota} = k\frac{d\tau_M^*}{d\iota} + \frac{\partial \Delta_M(q_M)}{\partial \iota}|_{q_M=q_M^*} = k\frac{d\tau_M^*}{d\iota} + \frac{\partial \Delta(q_M)}{\partial \iota}|_{q_M=q_M^*}.$$

By the proof of Theorem 2.7, $\frac{\partial \Delta(q_M)}{\partial \iota} < 0$. Moreover, by Eq. A.86 and the proof of part (ii) of Lemma A.21, $\frac{d\tau_M^*}{d\iota} \leq 0$. In sum, $\frac{dk\tau^* + \Delta_M(q_M^*)}{d\iota} < 0$ and Theorem 2.7 continues to hold. □

Lemma A.23. *When the size of the plate is restricted:*
(i) If the self-service cost is not prohibitively high ($\iota < \zeta$), there exists a threshold in the plate size such that the per-diner expected plate waste monotonically decreases and the buffet operator's profit monotonically increases in the plate size if the plate size is below the threshold; otherwise, the per-diner expected plate waste and the buffet operator's profit are invariant to the plate size. Formally, if $\iota < \zeta$, when $M < \bar{d} - \frac{\iota}{o}$, $\frac{d\omega_M^*}{dM} < 0$ and $\frac{d\pi^*}{dM} > 0$; otherwise, $\frac{d\omega_M^*}{dM} = \frac{d\pi^*}{dM} = 0$.
(ii) If the self-service cost is prohibitively high ($\iota \geq \zeta$), there exists a threshold in the plate size such that the per-diner expected plate waste monotonically increases and the buffet operator's profit first increases then decreases in the plate size if the plate size is below the threshold; otherwise, the per-diner expected plate waste and the buffet operator's profit are invariant to the plate size. Formally, if $\iota \geq \zeta$, when $M < \frac{u\bar{d}}{u+o}$, $\frac{d\omega_M^*}{dM} > 0$, $\frac{d\pi^*}{dM} > 0$ if $M < \frac{u-k}{u+o}\bar{d}$, and $\frac{d\pi^*}{dM} < 0$ if $M \geq \frac{u-k}{u+o}\bar{d}$; otherwise, when $M \geq \frac{u}{u+o}\bar{d}$, $\frac{d\omega_M^*}{dM} = \frac{d\pi^*}{dM} = 0$.

Proof: The impact of the plate size on the expected individual waste follows from Eq. A.13 and part (ii) of Lemma A.21. To study the impact of the plate size on the buffet operator's expected profit, similar to the proof of Lemma A.22, consider

$$\frac{dk\tau^* + \Delta_M(q_M^*)}{dM} = k\frac{d\tau_M^*}{dM} + \frac{\partial \Delta_M(q_M)}{\partial M}|_{q_M=q_M^*} + \frac{\partial \Delta_M(q_M)}{\partial q_M}|_{q_M=q_M^*}\frac{dq_M^*}{dM}.$$

(1) When $\iota < o(\bar{d} - M) \Leftrightarrow M < \bar{d} - \frac{\iota}{o}$, by the Envelope Theorem and Eq. A.84,

$$\frac{dk\tau^* + \Delta_M(q_M^*)}{dM} = k\frac{d\tau_M^*}{dM} + \frac{\partial \frac{u(\bar{d} - d_M(q_M)^2)}{2}}{\partial M}|_{q_M=q_M^*} < 0,$$

where the last inequality is because of Eq. A.86 and $d_M(q_M) = q_M + M$.
(2) When $o(\bar{d} - M) \leq \iota < \zeta$, by Eq. A.85 and part (ii) of Lemma A.21, $\frac{dq_M^*}{dM} = \frac{d\tau_M^*}{dM} = 0$. Moreover, since $\bar{d} - d_M(q_M^*) \leq 0$, $\frac{\partial \Delta_M(q_M)}{\partial M} = 0$. Therefore, $\frac{dk\tau^* + \Delta_M(q_M^*)}{dM}$=0.
(3) When $\iota \geq \zeta$.
(3-i) If $M \geq \frac{u}{u+o}\bar{d}$, similar to the case (2), $\tau_M^* = q_M^* = \frac{u}{u+o}\bar{d}$, and $\bar{d} - d_M(q_M^*) \leq 0$. Therefore, $\frac{dk\tau^* + \Delta_M(q_M^*)}{dM}$=0.
(3-ii) If $M < \frac{u}{u+o}\bar{d}$, $\tau_M^* = q_M^* = M$, $\bar{d} - d_M(q_M^*) \leq 0$ and $\Delta_M(q_M) = \Delta(q_M)$, therefore,

$$\frac{dk\tau^* + \Delta_M(q_M^*)}{dM} = k + \frac{d\Delta(q_M)}{dq_M}|_{q_M=M} = k - u + \frac{o+u}{\bar{d}}M,$$

where the last equality is due to Eq. A.7. $\frac{dk\tau^* + \Delta_M(q_M^*)}{dM} > 0$ (i.e.,$d\pi^*/dM < 0$) if and only if $M > \frac{u-k}{u+o}\bar{d}$. □

A.4 Auxiliary Results and Discussions

Here we provide explanations, proofs, and facts for the additional results that we mentioned in Sections 2.2-2.5.

A.4.1 Endogenous Refill Quantity

When a diner takes q_1 on her first plate and the realized demand is d, a diner's expected total cost of refilling q_2 amount of food on her plate is given by

$$g_r(q_1, d, q_2) = \iota + o(q_2 + q_1 - d)^+ + u(d - q_1 - q_2)^+, \tag{A.87}$$

where the above consists of the self-service cost and the underage and overage costs.

Lemma A.24. *If $q = q_1 < \bar{d}$ and $D = d > q_1$, define $q_2^* = \arg\min_{0 \leq q_2 \leq \bar{d} - q_1} g_r(q_1, d, q_2)$, then $q_2^* = d - q_1$.*

Proof: $g_r(q_1, d, q_2) \geq \iota$ and the equality is uniquely achieved at $q_2 = d - q_1$, therefore, $q_2^* = d - q_1$. □

A.4.2 Minimum Consumption Needs and the Impact of Demand Uncertainty

In the base model, we assumed that the lowest possible realization of the diner's demand is zero. As an extension, we assume that all diners eat at least $\underline{d}$ in the buffet (subscript

'a'). Formally, we consider the new demand distribution $D_a \sim U[\underline{d}, \bar{d}]$ and refer to $\underline{d}$ as the adjusted minimum consumption need. All our main results (Theorem 2.3 to 2.14) continue to hold. The derivation of equilibrium results and analyses are a mere replication of those provided in Sections A.2.1 and A.2.2.

The mean and standard deviation of the demand D_a can be expressed by the minimum consumption need and maximum consumption level, i.e., $\mu_{D_a} = \frac{\underline{d}+\bar{d}}{2}$ and $\sigma_{D_a} = \frac{\bar{d}-\underline{d}}{2\sqrt{3}}$.

Lemma A.25. *$\exists \iota_a > 0$ such that $\frac{d\omega_a^*}{d\sigma_{D_a}} < 0$ if and only if $\iota < \iota_a$.*

Proof: Replicating the proofs of Appendix Lemma A.1 and Proposition 2.2, we have

$$q_a^* = \begin{cases} \frac{\iota}{o} + \underline{d} = \frac{\iota}{o} + \mu_{D_a} - \sqrt{3}\sigma_{D_a} & \text{if } \iota < \zeta_a \\ \frac{u}{u+o}\bar{d} + \frac{o}{u+o}\underline{d} = \frac{u-o}{u+o}(\sqrt{3}\sigma_{D_a}) + \mu_{D_a} & \text{if } \iota \geq \zeta_a \end{cases},$$

$$\omega_a^* = \frac{(q_a^*)^2}{2(\bar{d}-\underline{d})} = \begin{cases} \frac{\iota^2}{2o^2(2\sqrt{3}\sigma_{D_a})} & \text{if } \iota < \zeta_a \\ \frac{u^2(2\sqrt{3}\sigma_{D_a})}{2(u+o)^2} & \text{if } \iota \geq \zeta_a \end{cases},$$

where $\zeta_a = \frac{ou}{o+u}(\bar{d}-\underline{d})$. Define $\iota_a = \zeta_a$. When food accessibility is not too low and a diner has a positive chance of refilling, greater uncertainty in her demand encourages her to take less on the first plate and opt for a refill instead. When food accessibility is low and a diner never refills, the higher cost of falling short of demand compared to generating waste ($u > o$) leads to increased waste as demand uncertainty grows—the more uncertain the demand, the more food the diner takes, and the more waste is generated. □

A.4.3 Endogenous Self-service Cost via Congestion Feedback

In this extension, we endogenize the self-service cost ι by allowing it to depend both on the buffet's physical design and the level of crowding. While ι was treated as exogenous in the main model, in practice, it is shaped by two components: by operator choices on food access, and congestion from the number of diners present. We model this by letting ι be the product of:

- ι_c: an operator-controlled inconvenience level representing the physical inconvenience embedded in the buffet's layout (e.g., number and placement of food stations, walking distance). Lower values of ι_c correspond to more convenient designs. And
- $\phi(\iota_c)$: the equilibrium mass of diners who choose to enter the buffet given the operator-controlled inconvenience level.

Formally, we define:

$$\iota = \iota_c \cdot \phi(\iota_c). \tag{A.88}$$

This formulation captures a key tradeoff: reducing ι_c (e.g., by improving layout or adding food stations) makes the buffet more convenient by design, but may attract more diners and increase congestion, which offsets some of the gains. The diner's expected total cost remains governed by Eq. 2.1, with ι now endogenously defined by (A.88). We restrict attention to pure-strategy Nash equilibria.

In Section A.4.3.1, we derive the equilibrium refill decisions under this formulation. In Section A.4.3.2, we prove that Theorem 2.7 continues to hold, and we provide an upper bound on the unit production cost that is sufficient for Theorem 2.3 to continue to hold.

A.4.3.1 Equilibirum Analysis

Lemma A.26. *For any operator-controlled inconvenience level ι_c, there exists a threshold in the flat-rate entrance fee such that all diners who join the buffet have the same probability of making a refill if the flat-rate entrance fee is above the threshold; and never refilling if the flat-rate entrance fee is below the threshold. Formally, there exists $\bar{\varphi}_c(\iota_c) \equiv \bar{\varphi}_c$ such that $\mathbb{P}(\textit{refill}) > 0$ if and only if $\varphi_c > \bar{\varphi}_c$. Moreover, $\mathbb{P}(\textit{refill})$ is the same for all diners who join the buffet.*

Proof: Lemma A.26 is proved in three steps. First, we show that equilibrium strategies are symmetric across diners. Second, we derive the two possible symmetric strategies diners may follow. Third, we characterize which of these equilibria the buffet admits under different flat-rate entrance fees.

Step 1. Since diners minimize their expected total cost of dining at the buffet, and the outside option cost is sunk once a diner joins, and given that diners' demands are independently drawn from the same distribution D, all diners adopt symmetric strategies in equilibrium. That is, every diner takes the same amount of food on the first plate and makes the same refill decision for any realized demand.

Step 2. It is straightforward to show that if a diner with outside option cost c_o chooses to join the buffet in equilibrium, then any diner with a higher outside option cost also joins. Let c_o^e denote the outside option cost of the marginal diner who is indifferent between joining and not joining.

(1) We first consider the symmetric equilibrium in which no diner refills. In this case, the following conditions must hold:

$$c_o^c = \varphi_c + \Delta(q_c^*)|_{\iota=\iota_c\phi}. \tag{A.89}$$

$$\phi = 1 - \frac{c_o^c}{\bar{c}_o}. \tag{A.90}$$

$$q_c^* = \arg\min_q \Delta(q_c), \tag{A.91}$$

where $\Delta(q_c)$ is defined in Eq. A.20. The solution to equations A.89 to A.91 is

$$c_o^c = \frac{\bar{c}_o}{\bar{c_o} + \iota_c}\left(\iota_c + \varphi_c + \frac{ou}{2(o+u)}\bar{d}\right). \tag{A.92}$$

$$q_c^* = \frac{u}{u+o}\bar{d}. \tag{A.93}$$

(2) We next consider the symmetric equilibrium in which all diners who join the buffet have a positive probability of refilling. As in the baseline model, it is straightforward to show that all diners follow a threshold-based refill policy: each diner refills if and only if $d \geq d_c$, where d_c is the threshold demand level satisfying $\iota = u(d_c - q_c)$. The solution to equations A.89 to A.91 with $\Delta(q_c)$ defined by Eq. A.22 is

$$c_o^c = \bar{c}_o - \left(\bar{c}_o + 2\iota_c - \sqrt{(\bar{c}_o + 2\iota_c)^{\,2} - 2\iota_c^2\frac{o+u}{ou\bar{d}}\,(\bar{c}_o - \varphi_c)}\right)\frac{\bar{c}_o ou\bar{d}}{\iota_c^2(o+u)}. \tag{A.94}$$

$$q_c^* = \frac{\iota_c}{o}\left(1 - \frac{c_o^c}{\bar{c}_o}\right), \tag{A.95}$$

where we discarded the negative root of c_o^c.

Step 3. We now consider the case in which diners are indifferent between refilling and not refilling—that is, between the two symmetric strategies described in Step 2. In this case, equilibrium requires that both strategies yield the same expected in-buffet cost. By equations A.20 and A.22, the following two conditions must hold in equilibrium:

$$c_o^c = \iota + \varphi_c + \frac{ou\bar{d}}{2(o+u)} = \varphi_c + 2\iota - \frac{(o+u)\iota^2}{2ou\bar{d}}.$$

$$\iota = \iota_c(1 - \frac{c_o^c}{\bar{c}_o}).$$

By Eq. A.92, the solution is

$$\iota_c = \frac{\bar{c}_o\xi}{(\bar{c}_o - \varphi_c) - \frac{3}{2}\xi} \Leftrightarrow \bar{\varphi}_c = \bar{c}_o - \frac{3\iota_c + 2\bar{c}_o}{2\iota_c}\xi, \tag{A.96}$$

where $\xi = \frac{ou\bar{d}}{o+u}$. Denote c_o^c given by Eq. A.92 and Eq. A.94 by $c_o^{c,n}$ and $c_o^{c,r}$ respectively. Since $\frac{\iota}{u} + \frac{\iota}{o}$ monotonically decreases in φ_c, and $\frac{\iota}{u} + \frac{\iota}{o} = \bar{d}$ when $\varphi_c = \bar{\varphi}_c$, no diner refills when $\varphi_c \leq \bar{\varphi}_c$. In contrast, when $\varphi_c > \bar{\varphi}_c$, the marginal outside option cost satisfies $c_o^{c,r} < c_o^{c,n}$. Therefore, all diners who join the buffet have a positive probability of refilling when $\varphi_c > \bar{\varphi}_c$. □

A.4.3.2 Impact of Improved Convenience under Endogenous Self-service Cost

Lemma A.27. *Under endogenous self-service cost, Theorem 2.7 continues to hold.*

Proof: We first prove that there exists a unique optimal flat-rate entrance fee that maximizes the buffet operator's expected profit.
(1) We first consider the operator's profit function when $\varphi_c \leq \bar{\varphi}_c$, by definition and Eq. A.92,

$$\begin{aligned}\pi_c(\varphi_c) &= \pi(\varphi_c, q_c^*)|_{\iota=\iota_c\phi} = \phi(\varphi_c - kq^*) = \left(1 - \frac{c_o^c}{\bar{c}_o}\right)\left(\varphi_c - k\frac{u\bar{d}}{u+o}\right) \\ &= \frac{\bar{c}_o - \varphi_c - \frac{ou}{2(o+u)}\bar{d}}{c_o + \iota_c}\left(\varphi_c - k\frac{u\bar{d}}{u+o}\right). \end{aligned} \tag{A.97}$$

By the first-order condition, φ_c^* has the same expression as Eq. A.16.
(2) When $\varphi_c > \bar{\varphi}_c$, the expression of the expected total amount of food taken by a diner τ_c has the same expression as Eq. A.24. Therefore,

$$\pi_c(\varphi_c) = \left(1 - \frac{c_o^c}{\bar{c}_o}\right)\left(\varphi_c - k\left(\frac{\bar{d}}{2} + \frac{\iota^2}{2\bar{d}}\left(\frac{1}{o^2} - \frac{1}{u^2}\right)\right)\right). \tag{A.98}$$

By solving the first-order condition (and applying Eq. A.94),

$$\frac{d\pi_c}{d\varphi_c} = -\frac{1}{\bar{c}_o}\frac{\partial c_o^c}{\partial \varphi_c}(\varphi_c - k\tau) + (1 - \frac{c_o^c}{\bar{c}_o})\left(1 - k\frac{\iota}{\bar{d}}(\frac{1}{o^2} - \frac{1}{u^2})\frac{\partial \iota}{\partial \varphi_c}\right) = 0,$$

we have,

$$\begin{aligned}\varphi_c^* = \frac{1}{18\iota_c^4(o+u)^2(k(o-u)+ou)^2}\Big(&6\iota_c^4(o+u)^2\bar{c}_o(k(o-u)+ou)(3k(o-u)+2ou) \\ &+ \iota_c^2 ou\bar{d}(o+u)(\iota_c^2\left(3k^2(o-u)(o+u) + 3kou(17u-15o) - 16o^2u^2\right) \\ &- 4ou\bar{c}_o(3k(o-u)+ou)\left(\bar{c}_o + 4\iota_c\right) + 2ou\bar{d}(3k(o-u)+ou)^2\left(\bar{c}_o + 2\iota_c\right)\iota_c^2(o+u) \\ &\times\sqrt{\frac{f_{\varphi_c}}{\bar{d}(3k(o-u)+ou)^2}}\Big), \end{aligned} \tag{A.99}$$

where we discarded the negative root and

$$f_{\varphi_c} = \Bigg(\bar{d} \left(4o^2u^2\bar{c}_o \left(\bar{c}_o + 4\iota_c \right) + \iota_c^2 \left(3k^2(o-u)(o+u) + 3kou(o+u) + 16o^2u^2 \right) \right) - 6\iota_c^2(o+u)\bar{c}_o(k(o-u)+ou) \Bigg).$$

Moreover, π_c at φ_c^* is a maximal since $\pi_c|_{\varphi_c \to \bar{c}_o} < 0$.
Since π_c has at most one maximum in each of $\varphi_c \leq \bar{\varphi}_c$ and $\varphi_c > \bar{\varphi}_c$, and π_c is continuously differentiable, π_c has a unique maximal in its domain.
Next, we prove that $\pi_c(\varphi_c^*) \equiv \pi_c^*$ monotonically decreases in ι_c. When diners never refills, by Eq. A.97 and the Envelope Theorem, $\frac{d\pi_c^*}{d\iota_c} = \frac{\partial \pi_c^*}{\partial \iota_c} < 0$. When diners have a positive probability of refilling. By Eq. A.88, A.90 and A.94,

$$\frac{\partial \iota}{\partial \iota_c} = \frac{oud\bar{c}_o \left(-\sqrt{\frac{2\iota_c^2(o+u)(\varphi_c - \bar{c}_o)}{oud} + \left(\bar{c}_o + 2\iota_c\right)^2} + \bar{c}_o + 2\iota_c \right)}{\iota_c^2(o+u)\sqrt{\frac{2\iota_c^2(o+u)(\varphi_c - \bar{c}_o)}{oud} + \left(\bar{c}_o + 2\iota_c\right)^2}} > 0. \tag{A.100}$$

Moreover, by Eq. A.88,

$$\frac{\partial c_o^c}{\partial \iota_c} = -\frac{\frac{d(\varphi_c + \Delta(q_c) - c_o^c)}{d\iota_c}}{\frac{\partial(\varphi_c + \Delta(q_c) - c_o^c)}{\partial c_o^c}} = -\frac{\frac{\partial(\varphi_c + \Delta(q_c) - c_o^c)}{\partial \iota}\frac{\partial \iota}{\partial \iota_c}}{\frac{\partial(\varphi_c + \Delta(q_c) - c_o^c)}{\partial \iota}\frac{\partial \iota}{\partial c_o^c} - 1} > 0.$$

In the last inequality, $\frac{\partial(\varphi_c + \Delta(q_c) - c_o^c)}{\partial \iota} > 0$ is because of Theorem 2.7, and $\frac{\partial \iota}{\partial c_o^c} < 0$ is because of Eq. A.88 and A.90.
Therefore, by the Envelope Theorem and Eq. A.98,

$$\frac{d\pi_c^*}{d\iota_c} = \frac{\partial \pi_c^*}{\partial \iota_c} = -\frac{1}{\bar{c}_o}(\varphi_c - k\tau_c)\frac{\partial c_o^c}{\partial \iota_c} - k\left(1 - \frac{c_o}{\bar{c}_o}\right)\frac{d\tau_c}{d\iota}\frac{\partial \iota}{\partial \iota_c} < 0.$$

Finally, since π_c^* is continuous, π_c^* monotonically decreases in ι_c. □

Lemma A.28. *When the self-service cost is endogenized, there exists a threshold for the operator-controlled inconvenience level, such that diners who join the buffet have a positive probability of refilling if and only if the operator-controlled inconvenience level is lower than the threshold hold. Formally, $\exists \hat{\iota}_c$ such that $\mathbb{P}(refill) > 0$ if and only if $\iota_c < \hat{\iota}_c$.*

Proof: Consider the case where the flat-rate entrance fee set by the buffet operator leads diners to be indifferent between refilling and not refilling. Under such a flat-rate entrance

fee, by Eq. A.96 and the expression for φ^* in Eq. A.16,

$$\hat{\iota}_c = \frac{\bar{c}_o\xi}{(\bar{c}_o-\varphi_c)-\frac{3}{2}\xi}|_{\varphi^*} = \frac{\bar{c}_o\xi}{\left(\bar{c}_o-\frac{\bar{c}_o-\frac{1}{2}\xi-\frac{k}{o}\xi}{2}\right)-\frac{3}{2}\xi} = \frac{2\bar{c}_o\xi}{\bar{c}_o-\left(\frac{5}{2}-\frac{k}{o}\right)\xi}.$$

We prove the result by contradiction. When $\iota_c < \hat{\iota}_c$, suppose diners never refill in equilibrium, we must have $\varphi_c^*(\iota_c) < \bar{\varphi}_c(\iota_c)$ according to Lemma A.26. However, according to Eq. A.96, $\bar{\varphi}_c(\iota_c)$ increases in ι_c. Therefore, have $\bar{\varphi}_c(\iota_c) < \bar{\varphi}_c(\hat{\iota}_c) = \varphi^* = \varphi_c^*(\iota_c)$. Contradicted. Diners have a positive probability of refilling when $\iota_c < \hat{\iota}_e$. A similar argument shows that diners never refill when $\iota_c \geq \hat{\iota}_c$. □

Lemma A.29. *If the unit production cost is sufficiently small, under endogenous self-service cost, Theorem 2.3 continues to hold.*

Proof: Eq. A.13 continuous to hold. By Lemma A.27, Lemma A.28 and Eq. A.13, when $\iota_c \geq \hat{\iota}_c$, $q_c^* = \frac{u}{o+u}\bar{d}$. Therefore, $\frac{d\omega^*}{d\iota_c} = 0$ when $\iota \geq \iota_c$.
When $\iota < \hat{\iota}_c$,

$$\frac{d\omega^*}{d\iota_c} = \frac{d\omega}{dq^*}\frac{dq^*}{d\iota_c} = \frac{d\omega}{dq^*}\frac{dq^*}{d\iota_c} = \frac{d\omega}{dq^*}\frac{dq^*}{d\iota}\left(\frac{\partial\iota}{\partial\iota_c}+\frac{\partial\iota}{\partial\varphi_c^*}\frac{d\varphi_c^*}{d\iota_c}\right) \tag{A.101}$$

By inequality A.100, $\frac{\partial\iota}{\partial\iota_c} > 0$. By Eq. A.99,

$$\frac{d\varphi_c^*}{d\iota_c} = -\frac{(\bar{c}_o-\varphi_c^*)-(\varphi_c^*-k\mathbf{E}[D])}{ou-3k(u-o)}\frac{ou\bar{c}_o}{\iota}\sqrt{\frac{\bar{d}(3k(o-u)+ou)^2}{A_e}},$$

where

$$A_e = \Bigg(\bar{d}\left(4o^2u^2\bar{c}_o\left(\bar{c}_o+4\iota_c\right)+\iota_c^2\left(3k^2(o-u)(o+u)+3kou(o+u)+16o^2u^2\right)\right) - 6\iota_c^2(o+u)\bar{c}_o(k(o-u)+ou)\Bigg).$$

When $\iota_c = 0$, $\iota = 0$. In this case, the equilibrium under the endogenous self-service cost is the same as that of the baseline model. As a result, $\varphi_c^* = \frac{1}{2}(\bar{c}_o + k\tau - \Delta(q_c^*))$ and $(\bar{c}_o-\varphi_c^*)-(\varphi_c^*-k\mathbf{E}[D]) > 0$. If $ou-3k(u-o) > 0 \Leftrightarrow k < \frac{ou}{3(u-o)}$, $\frac{d\varphi_c^*}{d\iota_c} < 0$ and φ_c^* monotonically decreases in ι_c (which guarantees that $(\bar{c}_o-\varphi_c^*)-(\varphi_c^*-k\mathbf{E}[D])$ remains positive).
Finally, by Eq. A.88, A.90, and A.94, we have $\frac{\partial\iota}{\partial\varphi_c} < 0$. Therefore, $\frac{d\omega^*}{d\iota_c} > 0$ (Eq. A.101) when $\iota < \hat{\iota}_c$ and Theorem 2.3 holds. □

A.4.4 Equivalence Between Monetary Instruments

In addition to the flat-rate penalty discussed in Section 2.5.1, we consider a per-weight penalty, where diners are penalized based on the amount of their leftovers. We also consider a reward system, where diners receive a flat-rate reward if they do not create any waste. A diner's expected total cost of dining at the buffet, as well as the buffet operator's expected profit under a per-weight penalty, are given by:

$$\begin{aligned} g_w(\varphi_w, v_w, q_w) &= \varphi_w + \mu_w(v_w) + \iota + \mathbb{E}_D[o \cdot (q_w - D)^+] + \mathbb{E}_D[\min\{\iota, u \cdot (D - q_w)^+\}] \\ &= \varphi_w + \mu_w(v_w) + \Delta(q_w), \\ \pi_w(\varphi_w, v_w) &= \frac{1}{\bar{c}_o}(\bar{c}_o - g_w(\varphi_w, v_w, q_w^*(\varphi_w, v_w))(\varphi_w + \mu_w(v_w) - k\tau(q_w^*(\varphi_w, v_w))), \end{aligned}$$

where v_w is the per-weight penalty fee and $\mu_w = v_w \cdot \omega_w^*$ is the expected total per-weight penalty collected. For the reward, let v_r be the flat-rate reward and $\mu_r = v_r \cdot \mathbb{P}(\bar{d} \geq q_r)$ be the expected reward paid to a diner:

$$\begin{aligned} g_r(\varphi_r, v_r, q_r) &= \varphi_r - \mu_r(v_r) + \Delta(q_r), \\ \pi_r(\varphi_r, v_r) &= \frac{1}{\bar{c}_o}(\bar{c}_o - g_r(\varphi_r, v_r, q_r^*(\varphi_r, v_r))(\varphi_r - \mu_r(v_r) - k\tau(q_r^*(\varphi_r, v_r))). \end{aligned} \tag{A.102}$$

All terms with subscripts 'w' and 'r' are defined similarly to terms with subscripts 'p' in Section 2.5.1.

Lemma A.30. *$\forall \iota, k$, $q_{d,p}^* = q_w^* = q_r^*$, $\pi_p^* = \pi_w^* = \pi_r^*$, and $\omega_p^* = \omega_w^* = \omega_r^*$.*

Proof: We prove the result in two steps. First, we show that similar to the flat-rate penalty (Eq. A.51), the equilibrium profit under the per-weight penalty and a reward fully depend on the expected non-monetary cost and per-diner expected production cost.
Repeating the proof of Lemma A.4, by the first-order condition, we have

$$\begin{aligned} \varphi_w^* &= \frac{1}{2}(\bar{c}_o - \Delta_w(v_w^*) + k\tau_w(v_w^*)) - \mu_w(v_w^*) \Rightarrow \pi_w^* = \frac{(\bar{c}_o - \Delta_w(v_w^*) - k\tau_w(v_w^*))^2}{4\bar{c}_o}, \\ \varphi_r^* &= \frac{1}{2}(\bar{c}_o - \Delta_r(v_r^*) - k\tau_r(v_r^*)) + \mu_r(v_r^*) \Rightarrow \pi_r^* = \frac{(\bar{c}_o - \Delta_r(v_r^*) - k\tau_r(v_r^*))^2}{4\bar{c}_o}. \end{aligned}$$

Second, we show that, for any given flat-rate penalty v_p, there exists a per-weight penalty v_w and a flat-rate reward v_r, such that $q_p(v_p) = q_w(v_w) = q_r(v_r)$. And the reverse is also true. This result implies that $\Delta_p^* = \Delta_w^* = \Delta_r^*$ and $\tau_p^* = \tau_w^* = \tau_r^*$. It is easy to check that if v_p satisfies the first and second order condition, so do v_w and v_r. In other words, $\pi_p^* = \pi_w^* = \pi_r^*$ and $\omega_p^* = \omega_w^* = \omega_r^*$.

(i) Let $o_w = o + v_w$. g_w has the same expression as g (Eq. 2.2) by replacing o with o_w. Therefore, $q_w^* = \frac{\min\{\iota, \zeta_w\}}{o_w}$, where $\zeta_w = \frac{o_w u \bar{d}}{o_w + u}$. When $q_p^* = \frac{\iota - v_p}{o}$, let $v_w = \frac{v_p}{\iota - v_p} o$, then $q_w^* = q_p^*$. When $q_p^* = \frac{u\bar{d} - v_p}{u + o}$, let $v_w = \frac{v_p}{u\bar{d} - v_p}(u + o)$, then $q_w^* = q_p^*$. Last we need to check that if $q_p^* = \frac{\iota - v_p}{o}$, $\iota < \zeta_w$ (so that $q_w^* = \frac{\iota}{o_w}$), which is true as

$$q_p^* = \frac{\iota - v_p}{o} \Leftrightarrow v_p > \tilde{v}_p = o\bar{d}(\frac{\iota}{\zeta} - 1) \Leftrightarrow \iota - \frac{o\iota}{o + v_w} > o\bar{d}(\frac{\iota}{\zeta} - 1) \Leftrightarrow \iota < \frac{u\bar{d}(o + v_w)}{o + u + v_w} = \zeta_w.$$

(ii) Consider $v_r = v_p$, we have $\mu_r = v_r \mathbb{P}(\bar{d} \geq q_r) = v_r - v_r \mathbb{P}(\bar{d} < q_r) = v_p - v_p \mathbb{P}(\bar{d} < q_r)$. By Equations 2.7 and A.102, we have

$$\frac{dg_r(q_r)}{dq_r} = \frac{d(g_p(q_r) - v_r)}{dq_r} = \frac{dg_p(q_r)}{dq_r}.$$

Therefore, $q_r^*(v_r) = q_p^*(v_p)|_{v_p = v_r}$. □

APPENDIX B

Proofs in Chapter 3

For ease of exposition, we denote the Signal-Noise Ratio by $\sigma := \frac{\xi}{\epsilon}$.

Proof of Assumption 3.1

Proof. Proof. Since the seller experiences stockout and overstock in High seasons ($\xi_t = \xi$) and Low seasons ($\xi_t = -\xi$), by condition 3.8,

$$\begin{aligned}
&0 < \Pr(d_t > \hat{q}_S \mid \xi_t = -\xi) < \Pr(d_t > \hat{q}_S \mid \xi_t = \xi) < 1 \\
\implies &0 < \Pr(d_t > \hat{q}_S \mid \xi_t = -\xi) < 1 \wedge 0 < \Pr(d_t > \hat{q}_S \mid \xi_t = \xi) < 1 \\
\implies &\hat{q}_S \in (\mu - \xi - \epsilon, \mu - \xi + \epsilon) \wedge \hat{q}_S \in (\mu + \xi - \epsilon, \mu + \xi + \epsilon) \\
\implies &\mu + \xi - \epsilon < \mu - \xi + \epsilon \implies \xi < \epsilon.
\end{aligned}$$

This further simplifies the probability density function of $\xi + \epsilon$ to

$$f_{\xi+\epsilon}(x) = \Pr(\xi_t = \xi) f_\epsilon(x - \xi) + \Pr(\xi_t = -\xi) f_\epsilon(x + \xi) = \frac{1}{4\epsilon}$$

when $d_t < \mu + \xi - \epsilon$ or when $d_t > \mu - \xi + \epsilon$; and to

$$f_{\xi+\epsilon}(x) = \Pr(\xi_t = \xi) f_\xi(x - \xi) + \Pr(\xi_t = -\xi) f_\xi(x + \xi) = \frac{1}{4\epsilon} + \frac{1}{4\epsilon} = \frac{1}{2\epsilon}$$

when $d_t \in (\mu + \xi - \epsilon, \mu - \xi + \epsilon)$. The cumulative distribution function $F_{\xi+\epsilon}$ is therefore given as,

$$F_{\xi+\epsilon}(x) = \begin{cases} \frac{x-(-\xi-\epsilon)}{4\epsilon} & \text{if } \; x \le \xi - \epsilon; \\ \frac{2\xi}{4\epsilon} + \frac{x-(\xi-\epsilon)}{2\epsilon} & \text{if } \; \xi - \epsilon < x \le -\xi + \epsilon; \\ \frac{2\xi}{4\epsilon} + \frac{(-\xi+\epsilon)-(\xi-\epsilon)}{2\epsilon} + \frac{x-(-\xi+\epsilon)}{4\epsilon} & \text{if } \; -\xi + \epsilon < x. \end{cases} \tag{B.1}$$

By condition 3.8 and the definition of $\hat{q}_S$,

$$\begin{aligned}
\hat{q}_S \in (\mu + \xi - \epsilon, \mu - \xi + \epsilon) &\implies F^{-1}_{\xi+\epsilon}(r_S) \in (\xi - \epsilon, -\xi + \epsilon) \\
&\implies r_S \in (F_{\xi+\epsilon}(\xi - \epsilon), F_{\xi+\epsilon}(-\xi + \epsilon)) \\
&\implies r_S \in \left(\frac{\xi}{2\epsilon}, 1 - \frac{\xi}{2\epsilon}\right) \\
&\implies \frac{\xi}{\epsilon} < 2\min(r_S, 1 - r_S). \qquad \text{(B.2)}
\end{aligned}$$

By condition 3.9, and the definition of $q_P(\xi_t)$ and $\hat{q}_S$,

$$\begin{aligned}
\hat{q}_S \geq \min q_P(\xi_t) &\implies \mu + F^{-1}_{\xi+\epsilon}(r_S) \geq \mu - \xi + F^{-1}_{\epsilon}(r_P) \\
&\implies \mu - \epsilon + 2\epsilon r_S \geq \mu - \xi - \epsilon + 2\epsilon r_P \\
&\implies \frac{\xi}{\epsilon} \geq 2(r_P - r_S), \qquad \text{(B.3)}
\end{aligned}$$

where we substituted $F^{-1}_{\xi+\epsilon}(r_S)$ with $-\epsilon + 2\epsilon r_S$ according to equations B.1 and B.2, and substituted $F^{-1}_{\epsilon}(r_P)$ with $-\epsilon + 2\epsilon r_P$ as ϵ is uniformly distributed. Inequalities B.2 and B.3 completes the proof. □

Proof of Proposition 3.2

Proof. Proof. The proof proceeds in three parts:

Part 1: We first show that

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \alpha_t = \lim_{t\to\infty} \mathbb{E}[\alpha_t] = \rho(\eta) = \Pr[\Delta_t \Pi^S > 0].$$

Since ξ_t and ϵ_t are i.i.d., d_t is i.i.d.. Moreover, since $\bar{q}^R_t(\eta)$ depends only on ξ_t, the events $\{\Delta_t \Pi^S > 0\}$ are i.i.d. Define $p = \Pr[\Delta_t \Pi^S > 0] = \rho(\eta)$. As $\{d_t : \Delta_t \Pi^S = 0\}$ has measure zero and d_t is continuously distributed, $\Pr[\Delta_t \Pi^S = 0] = 0$.

As $\Pr[\Delta_t \Pi^S = 0] = 0$, the trust update rule simplifies to:

$$\alpha_{t+1} = \begin{cases} \delta + (1 - \delta)\alpha_t & \text{with probability } p, \\ (1 - \delta)\alpha_t & \text{with probability } 1 - p. \end{cases}$$

Taking the conditional expectation yields:

$$\mathbb{E}[\alpha_{t+1} \mid \alpha_t] = p[\delta + (1 - \delta)\alpha_t] + (1 - p)(1 - \delta)\alpha_t = p\delta + (1 - \delta)\alpha_t.$$

Taking expectations on both sides, we obtain a linear recurrence:

$$\mathbb{E}[\alpha_{t+1}] = p\delta + (1-\delta)\mathbb{E}[\alpha_t].$$

Assuming convergence to a limit $\bar{\alpha}$, and setting $\mathbb{E}[\alpha_{t+1}] = \mathbb{E}[\alpha_t] = \bar{\alpha}$, we find:

$$\bar{\alpha} = p\delta + (1-\delta)\bar{\alpha} \implies \delta\bar{\alpha} = p\delta \implies \bar{\alpha} = p = \rho(\eta).$$

Since $\alpha_t \in [0,1]$, and the update process defines an irreducible and aperiodic Markov chain (due to $\delta > 0$), the process has a unique stationary distribution. Thus, $\mathbb{E}[\alpha_t] \to \rho(\eta)$ as $t \to \infty$. By ergodicity, the time average also converges:

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \alpha_t = \rho(\eta) = \Pr[\Delta_t \Pi^S > 0].$$

Part 2: We next derive the expression for $\rho(\eta)$. First, simplify $\bar{q}_t^R(\eta)$ as follows:

$$\bar{q}_t^R(\eta) = \mu + \xi_t + \eta(2\epsilon r_P - \epsilon) + (1-\eta)(2\epsilon r_S - \epsilon) = \mu + \xi_t + k(\eta),$$

where

$$k(\eta) = 2\epsilon[\eta r_P + (1-\eta) r_S] - \epsilon. \tag{B.4}$$

Let $q^R = \bar{q}_t^R(\eta)$ and $q^S = \hat{q}_S$. Then, the profit difference is:

$$\Delta_t \Pi^S = (1-\gamma)p[\min(q^R, d_t) - \min(q^S, d_t)] - (1-\beta)c(q^R - q^S).$$

We analyze $\Delta_t \Pi^S$ via case analysis:

1. If $q^R > q^S$:
   (a) $d_t \le q^S$: $\Delta_t \Pi^S = -(1-\beta)c(q^R - q^S) < 0$,
   (b) $q^S < d_t \le q^R$: $\Delta_t \Pi^S = (1-\gamma)p(d_t - q^S) - (1-\beta)c(q^R - q^S)$,
   (c) $d_t > q^R$: $\Delta_t \Pi^S = ((1-\gamma)p - (1-\beta)c)(q^R - q^S) = (1-\gamma)pr_S(q^R - q^S) > 0$.

   The threshold where $\Delta_t \Pi^S = 0$ in the middle region is:

$$d_+ = q^S + (1 - r_S)(q^R - q^S).$$

2. If $q^R < q^S$, then $\Delta_t \Pi^S > 0$ if $d_t < d_- = q^R + (1 - r_S)(q^S - q^R)$.

3. If $q^R = q^S$, then $\alpha_{t+1} > \alpha_t$ if $d_t > d_+ = q^R = q_S = d_-$.

To compute $\rho(\eta)$, we take the expectation over ξ_t:

$$\rho(\eta) = \mathbb{E}_{\xi_t}[\Pr[\Delta_t \Pi^S > 0 \mid \xi_t]].$$

1. For $\xi_t = \xi$: Let $q^R = q^R_+ = \mu + \xi + k(\eta)$, and $q^S = \mu + F^{-1}_{\xi+\epsilon}(r_S)$, assuming $q^R_+ > q^S$. Then, the threshold is $d_+ = q^S + (1 - r_S)(q^R_+ - q^S)$, and:

$$\Pr[d_t > d_+ \mid \xi_t = \xi] = 1 - F_\epsilon(d_+ - \mu - \xi) = 1 - \frac{d_+ - \mu - \xi + \epsilon}{2\epsilon}.$$

2. For $\xi_t = -\xi$: Let $q^R_- = \mu - \xi + k(\eta)$, assuming $q^R_- < q^S$. Then, the threshold is $d_- = q^R_- + (1 - r_S)(q^S - q^R_-)$, and:

$$\Pr[d_t < d_- \mid \xi_t = -\xi] = F_\epsilon(d_- - \mu + \xi) = \frac{d_- - \mu + \xi + \epsilon}{2\epsilon}.$$

Combining the two cases, we obtain:

$$\rho(\eta) = \frac{1}{2}\left[1 - \frac{d_+ - \mu - \xi + \epsilon}{2\epsilon}\right] + \frac{1}{2} \cdot \frac{d_- - \mu + \xi + \epsilon}{2\epsilon}.$$

Substituting the expressions and simplifying yields:

$$\rho(\eta) = \frac{1}{2} + \frac{\xi}{4\epsilon} - \frac{(1 - 2r_S)(r_P - r_S)}{2}\eta, \tag{B.5}$$

which is consistent with the parameter constraints (e.g., $r_P > r_S$) and Assumption 3.1.

Part 3: We now derive the steady-state variance of the trust level, $\lim_{t\to\infty} \mathrm{Var}(\alpha_t)$. The trust update rule can be rewritten as:

$$\alpha_{t+1} = (1 - \delta)\alpha_t + \delta Z_t,$$

where Z_t is an indicator variable such that $Z_t = 1$ if $\Delta_t \Pi^S > 0$ (with probability $p = \rho(\eta)$) and $Z_t = 0$ otherwise (with probability $1 - p$). The Z_t are i.i.d. across time, and since $0 < \delta < 1$, the process is a stable linear recursion.

We unroll the recursion to obtain an explicit expression for α_t:

$$\begin{aligned}
\alpha_1 &= (1-\delta)\alpha_0 + \delta Z_0, \\
\alpha_2 &= (1-\delta)\alpha_1 + \delta Z_1 = (1-\delta)^2\alpha_0 + \delta(1-\delta)Z_0 + \delta Z_1, \\
\alpha_3 &= (1-\delta)\alpha_2 + \delta Z_2 = (1-\delta)^3\alpha_0 + \delta(1-\delta)^2 Z_0 + \delta(1-\delta)Z_1 + \delta Z_2, \\
&\vdots \\
\alpha_t &= (1-\delta)^t\alpha_0 + \delta\sum_{k=0}^{t-1}(1-\delta)^{t-1-k}Z_k.
\end{aligned}$$

As $t \to \infty$, the deterministic term $(1-\delta)^t\alpha_0 \to 0$ due to $|1-\delta| < 1$, and α_t converges in distribution to the infinite weighted sum:

$$\alpha_\infty = \delta\sum_{k=0}^{\infty}(1-\delta)^k Z_k,$$

where $Z_k \sim \text{Bernoulli}(p)$ are i.i.d. random variables. To compute the variance of α_∞, we use the independence of the Z_ks:

$$\text{Var}(\alpha_\infty) = \text{Var}\left(\delta\sum_{k=0}^{\infty}(1-\delta)^k Z_k\right) = \delta^2\sum_{k=0}^{\infty}(1-\delta)^{2k}\text{Var}(Z_k).$$

Since $Z_k \sim \text{Bernoulli}(p)$, we have $\text{Var}(Z_k) = p(1-p)$, which is constant for all k. Therefore,

$$\text{Var}(\alpha_\infty) = \delta^2 p(1-p)\sum_{k=0}^{\infty}(1-\delta)^{2k} = \delta^2 p(1-p)\cdot\frac{1}{1-(1-\delta)^2} = \frac{\delta^2\rho(\eta)(1-\rho(\eta))}{2-\delta},$$

where we used $p = \rho(\eta)$. This completes the proof. □

Proof of Proposition 3.3

Proof. Proof. Notice that

$$\begin{aligned}
\bar{\Pi}^i(\eta) &:= \lim_{t\to\infty}\mathbb{E}\left[\Pi^i(q_t, d_t)\right] \\
&= \lim_{t\to\infty}\mathbb{E}\left[\alpha_t\Pi^i(\bar{q}^R(\eta), d_t) + (1-\alpha_t)\Pi^i(\hat{q}^S, d_t)\right] \\
&= \rho(\eta)\mathbb{E}\left[\Pi^i(\bar{q}^R(\eta), d_t)\right] + (1-\rho(\eta))\mathbb{E}\left[\Pi^i(\hat{q}^S, d_t)\right] \\
&= \underbrace{\mathbb{E}\left[\Pi^i(\hat{q}^S, d_t)\right]}_{\text{profit without recommendation}} + \underbrace{\rho(\eta)}_{\substack{\text{probability of} \\ \text{adopting recommendation}}} \times \underbrace{\mathbb{E}\left[\Pi^i(\bar{q}^R(\eta), d_t) - \Pi^i(\hat{q}^S, d_t)\right]}_{\substack{\text{expected performance} \\ \text{of recommendation}}}.
\end{aligned} \tag{B.6}$$

We first compute the expected overstock level when a seller follows the platform's recommended order quantity ($q_t = q_t^R$) and when she uses the preceived optimal order quantity ($q_t = \hat{q}_S$), i.e.,

$$\begin{aligned}\mathbb{E}[(\bar{q}_t^R - d_t)^+] &= \mathbb{E}[((\mu + \xi_t + k(\eta)) - (\mu + \xi_t + \epsilon_t))^+] \qquad \text{(B.7)}\\ &= \mathbb{E}[(k(\eta) - \epsilon_t)^+] = \int_{-\epsilon}^{\epsilon} (k(\eta) - \epsilon_t)^+ f_\epsilon(\epsilon_t) d\epsilon_t \\ &= \epsilon(\eta r_P + (1-\eta) r_S))^2; \qquad \text{(B.8)}\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}[(\hat{q}^S - d_t)^+] &= \mathbb{E}[\Big((\mu + F^{-1}_{\xi+\epsilon}(r_S)) - (\mu + \xi_t + \epsilon_t)\Big)^+] \\ &= \int_{-\xi-\epsilon}^{\xi+\epsilon} (-\epsilon + 2\epsilon r_S - s_t)^+ f_{\xi+\epsilon}(s_t) ds_t \qquad \text{(B.9)}\\ &= \int_{-\xi-\epsilon}^{\xi-\epsilon} \frac{-\epsilon + 2\epsilon r_S - s_t}{4\epsilon} ds_t + \int_{\xi-\epsilon}^{-\epsilon+2\epsilon r_S} \frac{-\epsilon + 2\epsilon r_S - s_t}{2\epsilon} ds_t \\ &= \frac{2(4\epsilon^2 r_S^2 + \xi^2)}{8\epsilon} = \epsilon\left(r_S^2 + \left(\frac{\xi}{2\epsilon}\right)^2\right). \qquad \text{(B.10)}\end{aligned}$$

In the above derivation, the density function $f_{\xi+\epsilon}$ is defined in the proof of Assumption 3.1.

Without recommendation, the seller always order at the percieved optimal level. Therefore, sending recommendation is rational to the platform if and only if

$$\bar{\Pi}^P(\eta) \geq \mathbb{E}[\Pi^P(\hat{q}^S, d_t)] \iff \mathbb{E}[\Pi^P(\bar{q}^R(\eta), d_t)] \geq \mathbb{E}[\Pi^P(\hat{q}^S, d_t)], \qquad \text{(B.11)}$$

where the equivalence relationship is due to equation B.6. Rewritting the platform's expected profit, we derive,

$$\begin{aligned}\mathbb{E}[\Pi^P(q_t, d_t)] &= \mathbb{E}[\gamma p \min(q_t, d_t) - \beta c q_t] \\ &= \mathbb{E}[(\gamma p - \beta c) q_t - \gamma p (q_t - d_t)^+] \\ &= (\gamma p - \beta c)\mathbb{E}[q_t] - \gamma p \mathbb{E}[(q_t - d_t)^+] \\ &= \gamma p \left(r_P \mathbb{E}[q_t] - \mathbb{E}[(q_t - d_t)^+]\right). \qquad \text{(B.12)}\end{aligned}$$

Using equations B.8, B.10, B.12, and the fact that

$$\begin{aligned}\mathbb{E}[\bar{q}_t^R] - \mathbb{E}[\hat{q}^S] &= \mathbb{E}[\mu + \xi_t - \epsilon + k(\eta)] - \mathbb{E}[\mu + \hat{\sigma}_S] \\ &= (\mu + k(\eta)) - (\mu + \hat{\sigma}_S) \\ &= 2\epsilon\eta(r_P - r_S) > 0,\end{aligned}$$

the platform's individual rationality constraint (inequality B.11) is simplified to

$$\begin{aligned}&\mathbb{E}[\Pi^P(\bar{q}^R(\eta), d_t)] \geq \mathbb{E}[\Pi^P(\hat{q}^S, d_t)] \\ \implies &\gamma p \left(r_P \mathbb{E}[\bar{q}^R(\eta)] - \mathbb{E}[(\bar{q}^R(\eta) - d_t)^+]\right) \geq \gamma p \left(r_P \mathbb{E}[\hat{q}^S] - \mathbb{E}[(\hat{q}^S - d_t)^+]\right) \\ \implies &r_P \geq \frac{\mathbb{E}[(\bar{q}^R(\eta) - d_t)^+] - \mathbb{E}[(\hat{q}^S - d_t)^+]}{2\epsilon\eta(r_P - r_S)} = \frac{\epsilon\left(\eta(r_P - r_S) + r_S\right)^2 - \epsilon\left(r_S^2 + \left(\frac{\xi}{2\epsilon}\right)^2\right)}{2\epsilon\eta(r_P - r_S)}.\end{aligned}$$

In a similar fashion, the seller's individual rationality constraint is given by,

$$r_S \geq \frac{\epsilon\left(\eta(r_P - r_S) + r_S\right)^2 - \epsilon\left(r_S^2 + \left(\frac{\xi}{2\epsilon}\right)^2\right)}{2\epsilon\eta(r_P - r_S)}. \tag{B.13}$$

Since $r_P \geq r_S$, the recommendation of the order quality is rational to both the platform and the seller if and only if inequality B.13 holds. Simplifying inequality B.13 gives

$$\eta \leq \frac{\xi}{2\epsilon(r_P - r_S)}. \tag{B.14}$$

Under assumption 3.1 and inequality 3.10, $\frac{\xi}{2\epsilon(r_P - r_S)} \geq 1$. Therefore, for all $\eta \in [0, 1]$, inequality B.14 always hold. □

Proof of Proposition 3.4

Proof. Proof. We divide the proof of proposition 3.4 into two parts. In the first part, we derive the optimal SoC given by equation 3.18. In the second part, we prove the sensitivity of $\tilde{\eta}$ with respect to (w.r.t.) ϵ, ξ, p and c.

Part 1: To begin with, we first derive the expanded expressions of the platform's expected profit when the seller adopts the recommendation and when she uses her perceived optimal order quality. By equations B.8, B.10, and B.12,

$$\mathbb{E}\left[\Pi^P(\bar{q}^R(\eta), d_t)\right] = \gamma p \left(r_P(\mu + k(\eta)) - \epsilon(\eta r_P + (1 - \eta) r_S)^2\right); \tag{B.15}$$

$$\mathbb{E}\left[\Pi^P(\hat{q}^S, d_t)\right] = \gamma p \left(r_P(\mu - \epsilon + 2\epsilon r_S) - \epsilon\left(r_S^2 + \left(\frac{\xi}{2\epsilon}\right)^2\right)\right). \tag{B.16}$$

Inserting equations B.4, B.5, B.15 and B.16 to equation B.6, we obtain $\bar{\Pi}^P(\eta)$

$$\begin{aligned}\bar{\Pi}^P(\eta) =& -\frac{1}{2}\gamma\eta^3 p\epsilon\left(2r_S-1\right)\left(r_P-r_S\right)^3 \\ &+\frac{1}{4}\gamma\eta^2 p\left(r_P-r_S\right)^2\left(-\xi+4\epsilon\left(2r_S-1\right)\left(r_P-r_S\right)-2\epsilon\right) \\ &+\frac{\gamma\eta p\left(r_P-r_S\right)\left(-\xi^2+4\epsilon r_P(\xi+2\epsilon)+2r_S\left(\xi^2-4\epsilon^2-2\xi\epsilon\right)\right)}{8\epsilon} \\ &+\frac{\gamma\xi^2 p(\xi-2\epsilon)}{16\epsilon^2}+\gamma p r_P\left(\mu+2\epsilon r_S-\epsilon\right)-\gamma p\epsilon r_S^2,\end{aligned} \tag{B.17}$$

and its first order derivative,

$$\begin{aligned}\frac{d\bar{\Pi}^P(\eta)}{d\eta} =& -\frac{3}{2}\gamma\eta^2 p\epsilon\left(2r_S-1\right)\left(r_P-r_S\right)^3 \\ &+\frac{1}{2}\gamma\eta p\left(r_P-r_S\right)^2\left(-\xi+4\epsilon\left(2r_S-1\right)\left(r_P-r_S\right)-2\epsilon\right) \\ &+\frac{\gamma p\left(r_P-r_S\right)\left(4\epsilon(\xi+2\epsilon)\left(r_P-r_S\right)+\xi^2\left(2r_S-1\right)\right)}{8\epsilon}.\end{aligned} \tag{B.18}$$

Moreover, we observe that,

$$\frac{d\bar{\Pi}^P(\eta)}{d\eta}\mid_{\eta=1} = \frac{\gamma p\left(2r_S-1\right)\left(r_P-r_S\right)\left(\xi^2+4\epsilon^2\left(r_P-r_S\right)^2\right)}{8\epsilon}; \tag{B.19}$$

$$\frac{d\bar{\Pi}^P(\eta)}{d\eta}\mid_{\eta=0} = \frac{\gamma p\left(r_P-r_S\right)\left(4\epsilon(\xi+2\epsilon)\left(r_P-r_S\right)+\xi^2\left(2r_S-1\right)\right)}{8\epsilon}. \tag{B.20}$$

1. When $r_S > \frac{1}{2}$, by equation B.18, $\frac{d\bar{\Pi}^P(\eta)}{d\eta}$ is concave. Furthermore, by equation B.19 and B.20, $\frac{d\bar{\Pi}^P(\eta)}{d\eta}\mid_{\eta=1} > 0$ and $\frac{d\bar{\Pi}^P(\eta)}{d\eta}\mid_{\eta=0} > 0$. Hence, $\frac{d\bar{\Pi}^P(\eta)}{d\eta} > 0$ for all $\eta \in [0,1]$, $\bar{\Pi}^P$ monotonically increases in η, and $\eta_P^* = 1$.

2. When $r_S = \frac{1}{2}$, the first order derivative is reduced to

$$\frac{d\bar{\Pi}^P(\eta)}{d\eta} = \frac{1}{8}\gamma(1-\eta)p\left(1-2r_P\right)^2(\xi+2\epsilon) > 0.$$

   Therefore, $\bar{\Pi}^P$ monotonically increases in η, and $\eta_P^* = 1$.

3. When $r_S < \frac{1}{2}$, $\frac{d\bar{\Pi}^P(\eta)}{d\eta}$ is convex and $\frac{d\bar{\Pi}^P(\eta)}{d\eta}\mid_{\eta=1} < 0$. Consider,

$$\frac{d\bar{\Pi}^P(\eta)}{d\eta}\mid_{\eta=0} < 0 \iff r_P - r_S < \frac{\xi^2(1-2r_S)}{4\epsilon(\xi+2\epsilon)} = \frac{\left(\frac{\xi}{2\epsilon}\right)^2}{1+\frac{\xi}{2\epsilon}}\left(\frac{1}{2}-r_S\right). \tag{B.21}$$

(a) If condition B.21 holds, $\frac{d\bar{\Pi}^P(\eta)}{d\eta} < 0$ for all $\eta \in [0,1]$, $\bar{\Pi}^P$ monotonically decreases in η, and $\eta_P^* = 0$.

(b) If condition B.21 does not hold, $\frac{d\bar{\Pi}^P(\eta)}{d\eta} \mid_{\eta=0} > 0$. There exists a unique $\tilde{\eta} \in (0,1)$ such that, $\frac{d\bar{\Pi}^P(\eta)}{d\eta} > 0$ for $\eta \in [0,\tilde{\eta})$ and $\frac{d\bar{\Pi}^P(\eta)}{d\eta} < 0$ for $\eta \in (\tilde{\eta},1]$. Therefore, $\bar{\Pi}^P(\eta)$ is maximized at $\tilde{\eta}$ and $\eta_P^* = \tilde{\eta}$. By the first order condition (FOC), i.e., $\frac{d\bar{\Pi}^P(\eta)}{d\eta} = 0$, we obtain the closed-form solution of $\tilde{\eta}$,

$$\tilde{\eta} = -\frac{-2\sqrt{A_1} - 4\epsilon\left(2r_S - 1\right)\left(r_P - r_S\right) + \xi + 2\epsilon}{2\epsilon\left(6r_S - 3\right)\left(r_P - r_S\right)}, \tag{B.22}$$

where $A_1 = 4\epsilon^2 r_P^2 \left(1 - 2r_S\right)^2 - \epsilon r_P \left(2r_S - 1\right)\left(-\xi + 8\epsilon r_S \left(2r_S - 1\right) - 2\epsilon\right) + \xi^2$.

Part 2: In part 2 of the proof, we assume $r_S < \frac{1}{2}$ and condition B.21 does not hold, so that $\eta_P^* = \tilde{\eta}$.

2.1 Sensitivity of η_P^* w.r.t. ξ. Because η_P^* satisfies the FOC, i.e., $\frac{d\bar{\Pi}^P(\eta)}{d\eta} \mid_{\eta=\eta_P^*} = 0$, by the Implicit Function Theorem,

$$\frac{d\eta_P^*}{d\xi} = -\frac{\left(\frac{\partial}{\partial\xi}\frac{d\bar{\Pi}^P(\eta)}{d\eta}\right)\mid_{\eta=\eta_P^*}}{\left(\frac{\partial}{\partial\eta}\frac{d\bar{\Pi}^P(\eta)}{d\eta}\right)\mid_{\eta=\eta_P^*}} = \frac{\gamma p\left(r_P - r_S\right)\left(\xi(1-2r_S) - 2(1-\eta_P^*)\epsilon(r_P - r_S)\right)}{4\epsilon\left(\frac{\partial}{\partial\eta}\frac{d\bar{\Pi}^P(\eta)}{d\eta}\right)\mid_{\eta=\eta_P^*}}. \tag{B.23}$$

We show that

$$\xi(1-2r_S) - 2(1-\eta_P^*)\epsilon(r_P - r_S) > 0 \iff \eta_P^* > 1 - \frac{\xi(1-2r_S)}{2\epsilon(r_P - r_S)}. \tag{B.24}$$

Consider,

$$\frac{d\bar{\Pi}^P(\eta)}{d\eta} \mid_{\eta=1-\frac{\xi(1-2r_S)}{2\epsilon(r_P-r_S)}} = -\frac{\gamma p\left(1-2r_S\right)\left(r_P - r_S\right)}{2\epsilon} A_2, \tag{B.25}$$

where $A_2 = \left(-r_S\left(2\epsilon r_P(\xi+\epsilon) + \xi(\epsilon - 3\xi)\right) + \epsilon r_P\left(\xi + \epsilon r_P\right) + r_S^2(\epsilon - \xi)(3\xi + \epsilon) - \xi(\xi + \epsilon)\right)$.

Observe that, because $r_S < \frac{1}{2}$ and $r_P > r_S$,

$$\frac{dA_2}{dr_P} = \epsilon\left(\xi + 2\epsilon r_P - 2r_S(\xi + \epsilon)\right) > \epsilon\left(\xi + 2\epsilon r_P - \xi - 2r_S\epsilon\right) > 0.$$

Therefore, A_2 monotonically increases in r_P. By assumption 3.1, $r_P < r_S + \frac{\xi}{2\epsilon}$, we have

$$\begin{aligned} A_2 &\leq A_2|_{r_P = r_S + \frac{\xi}{2\epsilon}} \\ &= -\frac{1}{4}\xi\left(\xi - 4\xi r_S\left(2 - 3r_S\right) + 4\epsilon\right) \\ &< -\frac{1}{4}\xi\left(\xi - 8\xi r_S + 4\epsilon\right) \\ &< -\frac{1}{4}\xi\left(\xi - 4\xi + 4\epsilon\right) < 0, \end{aligned}$$

where the last three inequalities are due to $r_S < \frac{1}{2}$ and $\xi < \epsilon$.

Because $A_2 < 0$, we have $\frac{d\bar{\Pi}^P(\eta)}{d\eta} \mid_{\eta = 1 - \frac{\xi(1-2r_S)}{2\epsilon(r_P - r_S)}} > 0$. By the discussion in part 1-3-(b), $\frac{d\bar{\Pi}^P(\eta)}{d\eta} > 0$ iff $\eta < \eta_P^*$. Therefore, condition B.24 holds and the numerator of $\frac{d\eta_P^*}{d\xi}$ (equation B.23) is positive. Lastly, because $\bar{\Pi}^P$ is concave in η, the denominator of $\frac{d\eta_P^*}{d\xi}$ is negative. In sum, $\frac{d\eta_P^*}{d\xi} < 0$.

2.2 Sensitivity of η_P^* w.r.t. ϵ. Replacing $\frac{\xi}{\epsilon}$ with σ, $\tilde{\eta}$ can be re-written as

$$\tilde{\eta}(\sigma) = -\frac{-2\sqrt{A_1(\sigma)} - 4\left(2r_S - 1\right)\left(r_P - r_S\right) + \sigma + 2}{2\left(6r_S - 3\right)\left(r_P - r_S\right)},$$

where $A_1 = 4r_P^2\left(1 - 2r_S\right)^2 - r_P\left(2r_S - 1\right)\left(-\sigma + 8r_S\left(2r_S - 1\right) - 2\right) + \sigma^2$. By the chain rule, there is

$$\frac{d\eta_P^*}{d\xi} = \frac{d\tilde{\eta}(\sigma)}{d\sigma}\frac{d\sigma}{d\xi} \quad \text{and} \quad \frac{d\eta_P^*}{d\epsilon} = \frac{d\tilde{\eta}(\sigma)}{d\sigma}\frac{d\sigma}{d\epsilon}.$$

Because $\frac{d\eta_P^*}{d\xi} < 0$ and the fact that $\frac{d\sigma}{d\xi} > 0$, $\frac{d\tilde{\eta}(\sigma)}{d\sigma} < 0$. Lastly, because $\frac{d\sigma}{d\epsilon} < 0$, $\frac{d\eta_P^*}{d\epsilon} > 0$.

2.3 Sensitivity of η_P^* w.r.t. p and c. Defining k, λ_P and λ_S and rewriting r_P and r_S as

$$k = \frac{p}{c},\ \lambda_P = \frac{\beta}{\gamma},\ \lambda_S = \frac{1-\beta}{1-\gamma};\ r_P = 1 - \lambda_P k,\ r_S = 1 - \lambda_S k. \tag{B.26}$$

We have

$$r_S = 1 - \lambda_S k < \frac{1}{2} \iff k > \frac{1}{2\lambda_S}. \tag{B.27}$$

Following a similar approach with part 2-2.1, we will prove that $\frac{d\eta_P^*}{dk} > 0$. We first re-write the first-order derivative by using equation B.26 and substituting r_P and r_S, which gives

$$\frac{d\bar{\Pi}^P}{d\eta} = -\frac{\gamma k p\left(\lambda_P - \lambda_S\right)}{8\epsilon} A_3(k), \tag{B.28}$$

where

$$\begin{aligned} A_3(k) = {} & 8\eta(3\eta-4)k^3\epsilon^2\lambda_S\left(\lambda_P-\lambda_S\right)^2 - 4\eta(3\eta-4)k^2\epsilon^2\left(\lambda_P-\lambda_S\right)^2 \\ & + \Big(4(\eta-1)\epsilon(\xi+2\epsilon)\left(\lambda_P-\lambda_S\right) - 2\xi^2\lambda_S\Big)k + \xi^2. \end{aligned} \tag{B.29}$$

In what follows, we show that $\frac{dA_3(k)}{dk}\,|_{\eta=\eta_P^*} < 0$ for all $k > \frac{1}{2\lambda_S}$. Consider,

$$\begin{aligned} \frac{dA_3(k)}{dk} = {} & 24\eta(3\eta-4)k^2\epsilon^2\lambda_S\left(\lambda_P-\lambda_S\right)^2 - 8\eta(3\eta-4)k\epsilon^2\left(\lambda_P-\lambda_S\right)^2 \\ & + 4(\eta-1)\epsilon(\xi+2\epsilon)\left(\lambda_P-\lambda_S\right) - 2\xi^2\lambda_S. \end{aligned} \tag{B.30}$$

When $\eta = \eta_P^*(=\tilde{\eta})$, by the first order condition and equation B.28 and B.29, $A_3(k)|_{\eta=\eta_P^*} = 0$ is equivalent to

$$\begin{aligned} & 4(\eta_P^*-1)\epsilon(\xi+2\epsilon)\left(\lambda_P-\lambda_S\right) - 2\xi^2\lambda_S \\ = & -\frac{1}{k}\left(8\eta_P^*(3\eta_P^*-4)k^3\epsilon^2\lambda_S\left(\lambda_P-\lambda_S\right)^2 - 4\eta_P^*(3\eta_P^*-4)k^2\epsilon^2\left(\lambda_P-\lambda_S\right)^2 + \xi^2\right). \end{aligned} \tag{B.31}$$

At $\eta = \eta_P^*$, applying equation B.31 to equation B.30 simplifies the first-order derivative to

$$\frac{dA_3(k)}{dk}|_{\eta=\eta_P^*} = 4\eta(3\eta-4)k\epsilon^2\left(4k\lambda_S-1\right)\left(\lambda_P-\lambda_S\right)^2 - \frac{\xi^2\left(2k\lambda_S+1\right)}{k},$$

which indicates that $\frac{dA_3(k)}{dk}$ is concave and symmetric w.r.t.

$$k = -\frac{8\eta(3\eta-4)\epsilon^2\left(\lambda_P-\lambda_S\right)^2}{2\times\left(24\eta(3\eta-4)\epsilon^2\lambda_S\left(\lambda_P-\lambda_S\right)^2\right)} = \frac{1}{6\lambda_S}.$$

Because $k > \frac{1}{2\lambda_S}$ (inequality B.27), $\frac{dA_3(k)}{dk}$ monotonically decreases in k. Therefore,

$$\frac{dA_3(k)}{dk}|_{\eta=\eta_P^*} < \frac{dA_3(k)}{dk}|_{\eta=\eta_P^*,k=\frac{1}{2\lambda_s}} = \frac{2\eta(3\eta-4)\epsilon^2\left(\lambda_P-\lambda_S\right)^2}{\lambda_S} - 4\xi^2\lambda_S < 0. \tag{B.32}$$

Because $A_3(k)|_{\eta=\eta_P^*} = 0$, by the Implicit Function Theorem, and with a slight abuse of notation

$$\frac{d\eta_P^*}{dk} = -\frac{\frac{\partial A_3(\eta,k)}{\partial k}|_{\eta=\eta_P^*}}{\frac{\partial A_3(\eta,k)}{\partial \eta}|_{\eta=\eta_P^*}} > 0. \tag{B.33}$$

The above inequality is due to (1) the numerator is negative (inequality B.32) and (2)

the denominator is positive because of the concavity of $\frac{d\bar{\Pi}^P}{d\eta}$ and equation B.28.

Finally, because η_P^* (equation B.22) is not an explicit function of p and c,

$$\frac{d\eta_P^*}{dp} = \frac{d\eta_P^*}{dk}\frac{dk}{dp} > 0 \quad \text{and} \quad \frac{d\eta_P^*}{dc} = \frac{d\eta_P^*}{dk}\frac{dk}{dc} < 0,$$

which follow directly from the definition of k. □

Proof of Proposition 3.5

Proof. Proof. Let $\bar{\Pi}^{P^*} = \bar{\Pi}^P(\eta_P^*)$, we prove $\frac{d\bar{\Pi}^{P^*}}{d\xi} < 0$ and $\frac{d\bar{\Pi}^{P^*}}{d\epsilon} < 0$ for $\eta_P^* \in (0,1)$, $\eta_P^* = 0$ and $\eta_P^* = 1$ in part 1, 2 and 3 respectively.

Part 1. When $\eta_P^* \in (0,1)$, we have $r_S < \frac{1}{2}$ and condition B.21 does not hold throughout part 1.

By the Envelope Theorem,

$$\frac{d\bar{\Pi}^{P^*}}{d\xi} = \frac{\partial\bar{\Pi}^{P^*}}{\partial\xi} = \frac{\gamma p\left(\xi(3\xi-4\epsilon) - 4\eta\epsilon\left(r_P - r_S\right)\left(\xi(1-2r_S) + (\eta-2)\epsilon(r_P - r_S)\right)\right)}{16\epsilon^2} \equiv \frac{\gamma p}{16\epsilon^2}A_4.$$

We will prove $A_4 < 0$. First, observe that A_4 is convex in r_P. Therefore, it suffices to prove that $A_4|_{\min\{r_P\}} < 0$ and $A_4|_{\max\{r_P\}} < 0$. By assumption 3.1, there is $r_P \in (r_S, r_S + \frac{\xi}{2\epsilon})$. Further,

$$\begin{aligned} A_4|_{\min\{r_P\}} &= A_4|_{r_P = r_S} = \xi(3\xi - 4\epsilon) < 0; \quad \text{and} \\ A_4|_{\max\{r_P\}} &= A_4|_{r_P = r_S + \frac{\xi}{2\epsilon}} \\ &= \xi\left(-\left(\eta^2 - 3\right)\xi + 4\eta\xi r_S - 4\epsilon\right) \\ &= \xi\left(-\xi(\eta - 2r_S)^2 + (3 + 4r_S^2)\xi - 4\epsilon\right) < 0. \end{aligned} \tag{B.34}$$

The inequality in the first line is due to $\epsilon > \xi$; and the inequality in the last line is due to $r_S < \frac{1}{2}$ and $\epsilon > \xi$.

Similarly, by the Envelope Theorem,

$$\frac{d\bar{\Pi}^{P^*}}{d\epsilon} = \frac{\partial\bar{\Pi}^{P^*}}{\partial\epsilon} = \frac{p\gamma}{8}A_5,$$

where

$$\begin{aligned} A_5 = \frac{1}{\epsilon^3}\Bigg(& -4(\eta-2)\eta^2\epsilon^3 r_P^3\left(2r_S - 1\right) + 4(\eta-2)\eta\epsilon^3 r_P^2\left(3\eta r_S\left(2r_S - 1\right) - 1\right) \\ & + \epsilon r_P\left(\eta\xi^2 + 2r_S\left(-\eta\xi^2 - 6(\eta-2)\eta^2\epsilon^2 r_S\left(2r_S - 1\right) + 4((\eta-2)\eta + 2)\epsilon^2\right) - 8\epsilon^2\right) \\ & + \epsilon r_S\left(2r_S\left(\eta\xi^2 + 2(\eta-2)\eta^2\epsilon^2 r_S\left(2r_S - 1\right) - 2((\eta-2)\eta + 2)\epsilon^2\right) - \eta\xi^2\right) + \xi^2(\epsilon - \xi)\Bigg). \end{aligned}$$

Substitute ξ with $\sigma\epsilon$ and $r_P - r_S$ with t, and with a slight abuse of notation, we obtain

$$A_5 = -\frac{1}{2}(2-\eta)\eta^2 t^3 (1-2r_S) + \frac{1}{2}(2-\eta)\eta t^2 - \frac{1}{8}t\left(2-\eta\sigma^2\right)(1-2r_S) - (1-r_S)\, r_S + \frac{1}{8}(1-\sigma)\sigma^2. \tag{B.35}$$

Because $t = r_P - r_S \leq \frac{\xi}{2\epsilon} = \frac{\sigma}{2}$, $r_S < \frac{1}{2}$, $\sigma = \frac{\xi}{\epsilon} < 2r_S$ and $\eta < 1$,

$$\begin{aligned}
A_5 &< \frac{1}{2}(2-\eta)\eta t^2 - (1-r_S)\, r_S + \frac{1}{8}(1-\sigma)\sigma^2 \\
&< \frac{1}{2}(2-\eta)\eta \frac{\sigma^2}{4} - (1-r_S)\, r_S + \frac{1}{8}(1-\sigma)\sigma^2 \\
&< \frac{1}{2}\frac{\sigma^2}{4} - (1-r_S)\, r_S + \frac{1}{8}(1-\sigma)\sigma^2 \\
&< \frac{1}{2}\frac{\sigma^2}{4} - \left(1-\frac{\sigma}{2}\right)\frac{\sigma}{2} + \frac{1}{8}(1-\sigma)\sigma^2 \\
&= -\frac{1}{8}(\sigma-2)^2\sigma < 0,
\end{aligned}$$

where the third line is due to $\min(2-\eta)\eta = 1$, and the fourth line is due to $\max(1-r_S)r_S = \left(1-\frac{\sigma}{2}\right)\frac{\sigma}{2}$.

Part 2. When $\eta_P^* = 0$, by Proposition 3.4, $r_S < \frac{1}{2}$. Therefore, $\frac{d\bar{\Pi}^{P^*}}{d\xi} < 0$ and $\frac{d\bar{\Pi}^{P^*}}{d\epsilon} < 0$ continue to hold from the proof of part 1 by treating $\eta = 0$ as a special case.

Part 3. When $\eta_P^* = 1$, by Proposition 3.4, $r_S > \frac{1}{2}$. Because $\eta_P^* = 1$, we again have $\frac{d\bar{\Pi}^{P^*}}{d\xi} = \frac{\partial\bar{\Pi}^{P^*}}{\partial\xi}$ and $\frac{d\bar{\Pi}^{P^*}}{d\epsilon} = \frac{\partial\bar{\Pi}^{P^*}}{\partial\epsilon}$. Therefore, similar to the proof of part 1, to show $\frac{d\bar{\Pi}^{P^*}}{d\xi} < 0$ and $\frac{d\bar{\Pi}^{P^*}}{d\epsilon} < 0$, it suffices to show that $A_4 < 0$ and $A_5 < 0$.

Inserting $\eta_P^* = 1$ into equation B.34, we have

$$A_4|_{\max r_P} = 2\xi\left(\xi + 2\xi r_S - 2\epsilon\right) < 2\xi\left(\xi + 2\xi(1-\frac{\epsilon}{2\xi}) - 2\epsilon\right) = 6\xi(\xi-\epsilon) < 0,$$

because $r_S < 1 - \frac{\epsilon}{2\xi}$ and $\xi < \epsilon$. Inserting $\eta_P^* = 1$ into equation B.35, we again have

$$\begin{aligned}
A_5 &< \frac{1}{2}\frac{\sigma^2}{4} - (1-r_S)\, r_S + \frac{1}{8}(1-\sigma)\sigma^2 < \frac{1}{2}\frac{\sigma^2}{4} - \frac{\sigma}{2}\left(1-\frac{\sigma}{2}\right) + \frac{1}{8}(1-\sigma)\sigma^2 \\
&= -\frac{1}{8}(\sigma-2)^2\sigma < 0,
\end{aligned}$$

because $r_S < 1 - \frac{\xi}{2\epsilon} = 1 - \frac{\sigma}{2}$. □

Proof of Lemma 3.7

Proof. Proof. With censored demand ($q_t \leq d_t$), the performance metric in (3.21) equals

$$\Delta_t \hat{\Pi}^S(q_t, q_t^R, \hat{q}_S) = \int_{q_t-\mu}^{\infty} \Big[\Pi^S(q_t^R, \mu + z) - \Pi^S(\hat{q}_S, \mu + z)\Big] \frac{f_{\xi+\epsilon}(z)}{1 - F_{\xi+\epsilon}(q_t - \mu)} \, \mathrm{d}z. \tag{B.36}$$

By definition 3.6, the demand is unambiguous if $q_t = \max\{q_t^R, \hat{q}^S\}$. Because $q_t = \max\{q_t^R, \hat{q}^S\}$, there is $\min(q_t^R, d) = q_t^R$ and $\min(\hat{q}_S, d) = \hat{q}_S$ for all $d > q_t$. Therefore, by definition (equation 3.4),

$$\Pi^S(q_t^R, \mu + z) - \Pi^S(\hat{q}_S, \mu + z) = \Big[(1-\gamma)p - (1-\beta)c\Big](q_t^R - \hat{q}_S).$$

Because $(1-\gamma)p > (1-\beta)$ (otherwise, the seller is always unprofitable), the integrand of equation B.36 is strictly positive over the full domain of integration if $q_t^R > \hat{q}^S$ and strictly negative otherwise. This observation completes the proof of Lemma 3.7.

□

Proof of Lemma 3.8

Proof. Proof. We prove this result in three steps. In the first step, we derive the posterior distributions of $\xi + \epsilon$ for $q_t < \mu - \epsilon + \xi$ and $q_t > \mu - \epsilon + \xi$ respectively. In the second step, we prove that the seller always decreases her trust in a low season if the censored demand is ambiguous. In the third step, we derive the conditions for the seller to increase her trust in the platform in a high season if the censored demand is ambiguous.

Step 1. When $q_t \in (\mu + \xi - \epsilon, \mu - \xi - \epsilon)$, by equation B.1, there is

$$F_{\xi+\epsilon}(q_t - \mu) = \frac{2\xi}{4\epsilon} + \frac{q_t - \mu - (\xi - \epsilon)}{2\epsilon}.$$

In this case, when $d_t > q_t$, equation 3.20 becomes,

$$\hat{f}_{\xi+\epsilon}(x) = \frac{f_{\epsilon+\xi}(x)}{1 - \frac{q_t - \mu + \epsilon}{2\epsilon}} = \begin{cases} \frac{1}{\epsilon + \mu - q_t} & x < \epsilon - \xi \\ \frac{1}{2}\frac{1}{\epsilon + \mu - q_t} & x > \epsilon - \xi. \end{cases} \tag{B.37}$$

When $q_t < \mu - \epsilon + \xi$, by equation B.1,

$$F_{\xi+\epsilon}(q_t - \mu) = \frac{q_t - \mu + \xi + \epsilon}{4\epsilon}.$$

In this case, when $d_t > q_t$, equation 3.20 becomes,

$$\hat{f}_{\xi+\epsilon}(x) = \frac{f_{\epsilon+\xi}(x)}{1 - \frac{q_t - \mu + \epsilon + \xi}{4\epsilon}} = \begin{cases} \frac{2}{4\epsilon - (q_t + \epsilon + \xi - \mu)} & -\epsilon + \xi < x < \epsilon - \xi \\ \frac{1}{4\epsilon - (q_t + \epsilon + \xi - \mu)} & x > \epsilon - \xi \vee x < -\epsilon + \xi \end{cases}. \tag{B.38}$$

Step 2. When the censored demand is ambiguous, we have $q_t = \min\{q_t^R, \hat{q}^S\}$. In a low season, we have $\hat{q}^S > q_t^R$. Therefore, we assume $q_t = q_t^R < d_t$ holds throughout step 2.

- When $\bar{q}_t^R(\eta)|_{\xi_t=-\xi} > \mu - \epsilon + \xi \iff \eta > \frac{\frac{\xi}{\epsilon} - r_S}{r_P - r_S}$. The posterior distribution of $\xi + \epsilon$ follows equation B.37. Let $q^R = \bar{q}_t^R(\eta)$. By definition, $\hat{\Pi}^S(q^R) = (1-\gamma)pr_S q^R$ because $\hat{d}_t \geq q^R$. Because $\hat{q}^S \in (\mu - \epsilon + \xi, \mu + \epsilon - \xi)$,

$$\begin{aligned} \hat{\Pi}^S(\hat{q}^S) &= \mathbb{E}[\Pi^S(\hat{q}^S, \hat{d}_t)] \\ &= \mathbb{E}[((1-\gamma)p - (1-\beta)c)\,\hat{q}^S - (1-\gamma)p(\hat{q}^S - \hat{d}_t)^+] \\ &= (1-\gamma)pr_S\hat{q}^S - (1-\gamma)p\frac{1}{\epsilon + \mu - q_t}\frac{(q^R - \hat{q}_S)^2}{2} \end{aligned}$$

and

$$\Delta\hat{\Pi}^S = \hat{\Pi}^S(q^R) - \hat{\Pi}^S(\hat{q}^S) = -(1-\gamma)p(\hat{q}_S - q^R)\left(r_S - \frac{\hat{q}_S - q^R}{2(\epsilon + \mu - q_t)}\right).$$

Inserting the expression of $\hat{q}^S$ and $q_t^R|_{\xi_t=-\xi}$, to show $\Delta\hat{\Pi}^S < 0$, it suffices to show that

$$r_S - \frac{\hat{q}^S - \hat{q}^R}{2(\epsilon + \mu - \hat{q}^R)} = r_S - \frac{1}{2}\frac{(\mu - \epsilon + 2\epsilon r_S) - (\mu - \xi - \epsilon + 2\epsilon(\eta r_P + (1-\eta)r_S))}{\epsilon + \mu - (\mu - \xi - \epsilon + 2\epsilon(\eta r_P + (1-\eta)r_S))} \equiv A_6 > 0.$$

Observe that A_6 monotonically increases in η because

$$\frac{dA_6}{d\eta} = \frac{2\epsilon^2(1 - r_S)(r_P - r_S)}{(\xi - 2\eta\epsilon r_P + 2(\eta - 1)\epsilon r_S + 2\epsilon)^2} > 0.$$

Therefore, we have

$$\begin{aligned}
A_6 &> A_6|_{\eta=\min\{0,\frac{\frac{\xi}{\epsilon}-r_S}{r_P-r_S}\}} \\
&> A_6|_{\eta=0} \\
&= \frac{\sigma}{r_S-2(\sigma+2)}+r_S \\
&> \frac{\sigma}{r_S-2(\sigma+2)}+r_S|_{\sigma=\min\{1,2-2r_S,2r_S\}} \\
&= \frac{\sigma}{r_S-2(\sigma+2)}+r_S|_{\sigma=\min\{2-2r_S,2r_S\}} \\
&= \begin{cases} \frac{r_S}{2} & \text{if } r_S<\frac{1}{2} \\ r_S-\frac{1}{4} & \text{if } r_S>\frac{1}{2} \end{cases},
\end{aligned}$$

which is always positive. Here, the second inequality is because $A_6|_{\eta=0}$ monotonically decreases in σ ($\frac{d}{d\sigma}\left(A_6|_{\eta=0}\right)=\frac{r_s-1}{(-2r_s+\sigma+2)^2}<0$) and because $\frac{\sigma}{2}<r_S<1-\frac{\sigma}{2} \implies \sigma<\min\{2r_S,2-2r_S\}$.

- When $\bar{q}_t^R(\eta)|_{\xi_t=-\xi}<\mu-\epsilon+\xi \iff \eta<\frac{\frac{\xi}{\epsilon}-r_S}{r_P-r_S}$. The posterior distribution of $\xi+\epsilon$ follows equation B.38. By definition, $\hat{\Pi}^S(q^R)=(1-\gamma)pr_Sq^R$ because $\hat{d}_t\geq q^R$. Because $\hat{q}^S\in(\mu-\epsilon+\xi,\mu+\epsilon-\xi)$,

$$\begin{aligned}
&\hat{\Pi}^S(\hat{q}^S)=\mathbb{E}[\Pi^S(\hat{q}^S,\hat{d}_t)] \\
&=\mathbb{E}[((1-\gamma)p-(1-\beta)c)\,\hat{q}^S-(1-\gamma)p(\hat{q}^S-\hat{d}_t)^+] \\
&=(1-\gamma)pr_S\hat{q}^S \\
&-(1-\gamma)p\left(\int_{q^R}^{\mu-\epsilon+\xi}\frac{(\hat{q}^S-\hat{d}_t)\mathrm{d}\hat{d}_t}{4\epsilon-(q_t+\epsilon+\xi-\mu)}+\int_{\mu-\epsilon+\xi}^{\hat{q}^S}\frac{2(\hat{q}^S-\hat{d}_t)\mathrm{d}\hat{d}_t}{4\epsilon-(q_t+\epsilon+\xi-\mu)}\right) \\
&=(1-\gamma)pr_S\hat{q}^S-(1-\gamma)p\frac{(\hat{q}_S-q^R)^2}{2(4\epsilon-(q_t+\epsilon+\xi-\mu))}-(1-\gamma)p\frac{(\hat{q}^S-\mu+\epsilon-\xi)^2}{2(4\epsilon-(q_t+\epsilon+\xi-\mu))},
\end{aligned}$$

and

$$\Delta\hat{\Pi}^S=\hat{\Pi}^S(q^R)-\hat{\Pi}^S(\hat{q}^S)=-(1-\gamma)p\left[r_S(\hat{q}^S-q^R)-\frac{(\hat{q}^S-\mu+\epsilon-\xi)^2+(q^R-\hat{q}^S)^2}{2(4\epsilon-(q_t+\epsilon+\xi-\mu))}\right].$$

To show $\Delta\hat{\Pi}^S < 0$, it suffices to show that

$$\begin{aligned}
&r_S(\hat{q}^S - q^R) - \frac{(\hat{q}^S - \mu + \epsilon - \xi)^2 + (q^R - \hat{q}^S)^2}{2(4\epsilon - (q_t + \epsilon + \xi - \mu))} \\
&= r_S\left(\xi + 2\eta\epsilon\left(r_S - r_P\right)\right) + \frac{\left(\xi + 2\eta\epsilon\left(r_S - r_P\right)\right)^2 + \left(\xi - 2\epsilon r_S\right)^2}{4\epsilon\left(\eta r_P - (\eta - 1)r_S - 2\right)} = A_7 > 0.
\end{aligned}$$

We first show that A_7 monotonically decreases in η. Consider,

$$\begin{aligned}
\frac{dA_7}{d\eta} &= \frac{4\left(\eta\left(r_P - r_S\right) + r_S - 2\right)\left(2\eta\left(r_P - r_S\right) - \sigma\right) - \left(\sigma - 2\eta\left(r_P - r_S\right)\right)^2}{4\left(\eta\left(r_P - r_S\right) + r_S - 2\right)^2} \\
&\quad - \frac{8r_S\left(\eta\left(r_P - r_S\right) + r_S - 2\right)^2 + \left(\sigma - 2r_S\right)^2}{4\left(\eta\left(r_P - r_S\right) + r_S - 2\right)^2} \\
&= \frac{A_8}{4\left(\eta\left(r_P - r_S\right) + r_S - 2\right)^2}.
\end{aligned}$$

To show A_7 decreases in η is equivalent to show A_8 is always negative. Observe that

$$\begin{aligned}
\frac{dA_8}{d\eta} &= 8(r_P - r_S)\left(2r_S - 1\right)\left(2 - r_S - \eta(r_P - r_S)\right) \\
\frac{dA_8}{dr_P} &= 8\eta\left(2r_S - 1\right)\left(2 - r_S - \eta(r_P - r_S)\right) \\
\frac{dA_8}{d\sigma} &= 8 - 4\sigma
\end{aligned}$$

If $r_S > \frac{1}{2}$, A_8 monotonically increases in η, r_P and σ, and $\sigma < \min\{2r_S, 2 - 2r_S\} = 2 - 2r_S$. Therefore,

$$A_8 < A_8|_{\eta=1, r_P = r_S + \frac{\sigma}{2}, \sigma = 2 - 2r_S} = -8\left(2r_S - 1\right)r_S - 4 < 0;$$

If $r_S < \frac{1}{2}$, A_8 monotonically decreases in η, r_P and monotonically increases in σ, and $\sigma < \min\{2r_S, 2 - 2r_S\} = 2r_S$. Therefore,

$$A_8 < A_8|_{\eta=0, r_P = r_S, \sigma = 2r_S} = -4r_S\left(r_S\left(2r_S - 5\right) + 4\right) < 0,$$

where the last inequality is because $r_S\left(2r_S - 5\right) + 4 > r_S\left(2r_S - 5\right)|_{r_S = 5/4} + 4 > 0$.

Because A_8 is always negative, A_7 monotonically decreases in η.

Second, we show that $A_7|_{\eta = \max\{1, \frac{\frac{\xi}{\epsilon} - r_S}{r_P - r_S}\}}$ is positive. Specifically, we consider the following two cases.

1. $r_P < \sigma$. In this case, $\max\{1, \frac{\frac{\xi}{\epsilon} - r_S}{r_P - r_S}\} = 1$, and

$$\begin{aligned} A_7|_{\eta=1} &= \frac{(-2r_P + 2r_S + \sigma)^2 + (\sigma - 2r_S)^2 + 4(r_P - 2) r_S (-2r_P + 2r_S + \sigma)}{4(r_P - 2)} \\ &= \frac{A_9}{4(r_P - 2)}. \end{aligned}$$

Observe that A_9 monotonically increases in r_P, specifically, consider

$$\frac{dA_9}{dr_P} = 2\left(r_P(2 - 4r_S) + r_S(2r_S + +2) - \sigma(1 - r_S)\right).$$

If $r_S > \frac{1}{2}$, $\frac{dA_9}{dr_P}$ monotonically decreases in r_P and $\sigma < 2 - 2r_S$. Therefore,

$$\frac{dA_9}{dr_P} > \frac{dA_9}{dr_P}|_{r_P=\sigma, \sigma=2-2r_S} = 4r_S(4r_S - 3) + 4 > 0;$$

If $r_S < \frac{1}{2}$, $\frac{dA_9}{dr_P}$ monotonically increases in r_P and $\sigma < 2r_S$. Therefore,

$$\frac{dA_9}{dr_P} > \frac{dA_9}{dr_P}|_{r_P=r_S, \sigma=2r_S} = 4r_S > 0.$$

Because A_9 monotonically increases in r_P, we have

$$\begin{aligned} A_9 &< A_9|_{\max r_P} < A_9|_{r_P=\sigma} = -2\sigma^2 r_S + 4(\sigma - 1)r_S^2 + \sigma^2 \\ &< -2\sigma^2 r_S + 4(\sigma - 1)r_S^2 + \sigma^2|_{r_S=1-\frac{\sigma}{2}} = -2(1 - \sigma)((\sigma - 1)^2 + 1) < 0. \end{aligned}$$

The second line is because $-2\sigma^2 r_S + 4(\sigma - 1)r_S^2 + \sigma^2$ monotonically increases in r_S. Finally, because $A_9 < 0$ (and $r_P - 2 < 0$), $A_7 > A_7|_{\eta=1} > 0$.

2. $r_P > \sigma$. In this case, $\max\{1, \frac{\frac{\xi}{\epsilon} - r_S}{r_P - r_S}\} = \frac{\frac{\xi}{\epsilon} - r_S}{r_P - r_S}$, and

$$A_7|_{\eta=\frac{\frac{\xi}{\epsilon} - r_S}{r_P - r_S}} = \frac{\epsilon(-2\sigma^2 r_S + 4(\sigma - 1)r_S^2 + \sigma^2)}{2(\sigma - 2)} > 0.$$

The inequality is because of $(-2\sigma^2 r_S + 4(\sigma - 1)r_S^2 + \sigma^2)$ monotonically decreasing in r_S and $(-2\sigma^2 r_S + 4(\sigma - 1)r_S^2 + \sigma^2)|_{r_S=\frac{\sigma}{2}} = 0$.

Because A_7 is decreasing in η and $A_7|_{\max \eta} > 0$, we proved that $\Delta\hat{\Pi} < 0$ if $\bar{q}_t^R(\eta)|_{\xi_t=-\xi} < \mu + \xi - \epsilon$.

In sum, by proving $A_6 > 0$ and $A_7 > 0$, we have successfully shown that $\Delta\hat{\Pi}^S < 0$ in a low season for all $q_t^R|_{\xi_t=-\xi} < \hat{q}^S$.

Step 3. When the censored demand is ambiguous, we have $q_t = \max\{q_t^R, \hat{q}^S\}$. In a high season, we have $\hat{q}^S < q_t^R$. Therefore, $q_t = \hat{q}^S < d_t$ holds throughout step 3. The proof is further divided into three steps. We first show that $\Delta\hat{\Pi}$ (in a high season with censored and ambiguous demand) is concave in η. Second, we show that $\Delta\hat{\Pi} > 0$ at $\eta = 0$. Third, because of the first two steps, we derive the condition of $\Delta\hat{\Pi} < 0$ by studying the sign of $\Delta\hat{\Pi}$ at $\eta = 1$.

(Step 3.1) Because $\hat{d}_t \geq \hat{q}^S$, $\hat{\Pi}^S(\hat{q}^S) = (1-\gamma)pr_S\hat{q}^S$. Because $\hat{q}^S < \mu + \epsilon - \xi$, the posterior distribution of $\xi + \epsilon$ follows equation B.37.

- When $q^R < \mu + \epsilon - \xi$,

$$\begin{aligned}\hat{\Pi}^S(q^R) &= \mathbb{E}[\Pi^S(q^R, \hat{d}_t)] = \mathbb{E}[((1-\gamma)p - (1-\beta)c)\, q^R - (1-\gamma)p(q^R - \hat{d}_t)^+] \\ &= (1-\gamma)pr_S q^R - (1-\gamma)p\frac{1}{\epsilon+\mu-q_t}\frac{(q^R-\hat{q}_S)^2}{2}. \\ \Delta\hat{\Pi}^S &= \hat{\Pi}^S(q^R) - \hat{\Pi}^S(\hat{q}^S) = (1-\gamma)p(q^R-\hat{q}_S)\left(r_S - \frac{q^R-\hat{q}_S}{2(\epsilon+\mu-q_t)}\right). \end{aligned} \tag{B.39}$$

  which is concave in q^R. Since $\bar{q}^R|_{\xi_t=\xi}(\eta) = \mu + \xi - \epsilon + 2\epsilon(\eta r_P + (1-\eta)r_S)$ is linearly increasing in η, $\Delta\hat{\Pi}$ is concave in η.

- When $q^R > \mu + \epsilon - \xi$, we have

$$\begin{aligned}&\hat{\Pi}^S(q^R) = \mathbb{E}[\Pi^S(q^R, \hat{d}_t)] \\ &= \mathbb{E}[((1-\gamma)p - (1-\beta)c)\, q^R - (1-\gamma)p(q^R - \hat{d}_t)^+] \\ &= (1-\gamma)pr_S q^R - (1-\gamma)p\left(\int_{\hat{q}_S}^{\mu+\epsilon-\xi}(q^R-\hat{d}_t)\frac{\mathrm{d}\hat{d}_t}{\epsilon+\mu-q_t} + \int_{\mu+\epsilon-\xi}^{q^R}(q^R-\hat{d}_t)\frac{\mathrm{d}\hat{d}_t}{2(\epsilon+\mu-q_t)}\right) \\ &= (1-\gamma)pr_S q^R - (1-\gamma)p\frac{1}{\epsilon+\mu-q_t}\frac{(q^R-\hat{q}_S)^2}{2} + (1-\gamma)p\frac{1}{2(\epsilon+\mu-q_t)}\frac{(q^R-\mu-\epsilon+\xi)^2}{2},\end{aligned}$$

  and

$$\begin{aligned}\Delta\hat{\Pi}^S &= \hat{\Pi}^S(q^R) - \hat{\Pi}^S(\hat{q}^S) \\ &= \frac{(1-\gamma)p}{4(\epsilon+\mu-q_t)}\left[(q^R-\hat{q}_S)\left(4(\epsilon+\mu-q_t)r_S - 2(q^R-\hat{q}_S)\right) + (q^R-\mu-\epsilon+\xi)^2\right],\end{aligned} \tag{B.40}$$

$$\frac{d^2\Delta\hat{\Pi}}{d(q^R)^2} = -\frac{(1-\gamma)p}{2(\mu-\hat{q}^S+\epsilon)} = -\frac{(1-\gamma)p}{4\epsilon\,(1-r_S)} < 0.$$

  Because $\bar{q}^R(\eta)|_{\xi_t=\xi}$ is linearly increasing in η, $\Delta\hat{\Pi}$ is concave in η.

(Step 3.2)

- If $q^R = \mu + \xi - \epsilon + 2\epsilon r_P < \mu - \xi + \epsilon \iff r_P < 1 - \sigma$, inserting the expression of $\hat{q}^S$ and $q_t^R|_{\xi_t=\xi}$ to equation B.39, we have

$$\Delta\hat{\Pi} = \frac{(\xi + 2\eta\epsilon\,(r_P - r_S))\,(\xi + 2\eta\epsilon r_P - 2(\eta+2)\epsilon r_S + 4\epsilon r_S^2)}{4\epsilon\,(r_S - 1)}. \tag{B.41}$$

  Letting $\eta = 0$, we have

$$\Delta\hat{\Pi}|_{\eta=0} = \frac{1}{4}\sigma\epsilon\left(\frac{\sigma}{r_S - 1} + 4r_S\right).$$

  Observe that $4\,(r_S - 1)\,r_S + \sigma < 4\,(r_S - 1)\,r_S + \sigma|_{r_S=1-\frac{\sigma}{2}} = 4\,(r_S - 1)\,r_S + \sigma|_{r_S=\frac{\sigma}{2}} = (\sigma - 1)\sigma < 0$. Therefore, $\Delta\hat{\Pi}|_{\eta=0} > 0$.

- If $q^R > \mu + \epsilon - \xi \iff r_P > 1 - \sigma$, inserting the expression of $\hat{q}^S$ and $q_t^R|_{\xi_t=\xi}$ to equation B.40, we have

$$\begin{aligned}\Delta\hat{\Pi} = \frac{(\gamma-1)p}{4\epsilon\,(r_S - 1)}\Bigg(&(-2\xi - 2\eta\epsilon r_P + 2(\eta-1)\epsilon r_S + 3\epsilon)^2 \\ &- 2\,(\xi + 2\eta\epsilon\,(r_P - r_S))\left(\xi + 2\eta\epsilon r_P - 2(\eta+3)\epsilon r_S + 4\epsilon r_S^2\right)\Bigg)\end{aligned} \tag{B.42}$$

  Letting $\eta = 0$, we have

$$\Delta\hat{\Pi}|_{\eta=0} = \frac{(\gamma-1)p\epsilon\,(-2(2\sigma-1)\,(r_S-2)\,r_S + (\sigma-4)\sigma + 2)}{4\,(r_S - 1)} = \frac{(\gamma-1)p\epsilon}{4\,(r_S-1)}A_{10}.$$

  Consider,

$$\frac{dA_{10}}{dr_S} = -4(2\sigma - 1)\,(r_S - 1).$$

  When $\sigma > \frac{1}{2}$, A_{10} monotonically increases in r_S and

$$A_{10} > A_{10}|_{r_S=1-\frac{\sigma}{2}} = \frac{1}{2}(3 - 2\sigma)\sigma^2 > 0.$$

  When $\sigma < \frac{1}{2}$, A_{10} monotonically decreases in r_S and

$$A_{10} > A_{10}|_{r_S=\frac{\sigma}{2}} = 2 - \frac{1}{2}(\sigma-4)\sigma(2\sigma-3) > 2 - \frac{1}{2}(\sigma-4)\sigma(2\sigma-3)|_{\sigma=\frac{1}{2}} = \frac{1}{4} > 0,$$

where the second inequality is because

$$\frac{d}{d\sigma}\left(2-\frac{1}{2}(\sigma-4)\sigma(2\sigma-3)\right)=-(\sigma-3)(3\sigma-2)<0.$$

In sum, $A_{10}>0$ and $\Delta\hat{\Pi}>0$.

Step 3.3 Because $\Delta\hat{\Pi}$ is concave in η and positive at $\eta=0$ in a high season, there exists $\bar{\eta}^c$ such that $\Delta\hat{\Pi}<0$ for $\eta>\bar{\eta}^c$ if and only if $\Delta\hat{\Pi}|_{\eta=1}<0$.

When $r_P<1-\sigma$, by equation B.41, there is

$$\Delta\hat{\Pi}|_{\eta=1}<0 \iff 2r_P+4r_S^2-6r_S+\sigma>0, \text{ and} \tag{B.43}$$

$$\Delta\hat{\Pi}=0 \implies \bar{\eta}_1^c=\frac{-\xi+4\epsilon(1-r_S)r_S}{2\epsilon(r_P-r_S)}. \tag{B.44}$$

When $r_P>1-\sigma$ by equation B.42,

$$\begin{aligned}&\Delta\hat{\Pi}|_{\eta=1}<0\\ &\iff -4r_S^2\left(2r_P+\sigma+3\right)+16r_Pr_S-2r_P\left(r_P+2\right)+8\sigma r_S+8r_S^3+(\sigma-4)\sigma+2<0,\end{aligned} \tag{B.45}$$

and

$$\begin{aligned}&\Delta\hat{\Pi}=0\\ &\implies \bar{\eta}_2^c=\frac{\sqrt{2}\sqrt{4r_S\left(r_S\left(2\left(r_S-3\right)r_S-\sigma+7\right)+2(\sigma-2)\right)+(\sigma-2)^2}-4r_S^2+6r_S-2}{2\left(r_P-r_S\right)}.\end{aligned} \tag{B.46}$$

We skip the details of the derivation, using assumption 3.1, conditions B.43 and B.45 can be combined and simplified to

- When $\sigma\leq\frac{4}{9}$,

$$r_S<\frac{1}{2}-\frac{1}{2}\sqrt{1-2\sigma}\wedge\frac{1}{2}\left(-4r_S^2+6r_S-\sigma\right)<r_P;$$

- When $\frac{4}{9} < \sigma \leq 2 - \sqrt{2}$,

$$\begin{cases} r_S \leq \frac{3}{4} - \frac{1}{4}\sqrt{4\sigma+1} \wedge \frac{1}{2}\left(-4r_S^2 + 6r_S - \sigma\right) < r_P \\ \left(\frac{3}{4} - \frac{1}{4}\sqrt{4\sigma+1} < r_S < \frac{7\sigma-2}{2(4\sigma-1)} - \sqrt{\frac{\sigma^3}{(4\sigma-1)^2}} \right. \\ \left. \wedge \frac{\sqrt{-4\sigma r_S^2 + 8\sigma r_S + 8r_S^4 - 24r_S^3 + 28r_S^2 - 16r_S + \sigma^2 - 4\sigma + 4}}{\sqrt{2}} - 2r_S^2 + 4r_S - 1 < r_P \right) \end{cases}$$

- $2 - \sqrt{2} \leq \sigma < 1$

$$r_S < \frac{7\sigma-2}{2(4\sigma-1)} - \sqrt{\frac{\sigma^3}{(4\sigma-1)^2}}$$
$$\wedge \frac{\sqrt{-4\sigma r_S^2 + 8\sigma r_S + 8r_S^4 - 24r_S^3 + 28r_S^2 - 16r_S + \sigma^2 - 4\sigma + 4}}{\sqrt{2}} - 2r_S^2 + 4r_S - 1 < r_P.$$

Define

$$\underline{r_S} = \min\left\{\frac{1}{2} - \frac{1}{2}\sqrt{1-2\sigma}, \frac{3}{4} - \frac{1}{4}\sqrt{4\sigma+1}, \frac{7\sigma-2}{2(4\sigma-1)} - \sqrt{\frac{\sigma^3}{(4\sigma-1)^2}}\right\}, \text{ and}$$

$$\bar{r}_P = \max\left\{\frac{1}{2}\left(-4r_S^2 + 6r_S - \sigma\right), \right.$$
$$\left. \frac{\sqrt{-4\sigma r_S^2 + 8\sigma r_S + 8r_S^4 - 24r_S^3 + 28r_S^2 - 16r_S + \sigma^2 - 4\sigma + 4}}{\sqrt{2}} - 2r_S^2 + 4r_S - 1\right\}.$$

When $r_S < \underline{r_S}$ and $r_P > \bar{r}_P$, define

$$\bar{\eta}^c = \begin{cases} \bar{\eta}_1^c & \text{if } \Delta\hat{\Pi}|_{\eta=\min\{\frac{1-\sigma-r_S}{r_P-r_S},1\}} < 0; \\ \bar{\eta}_2^c & \text{if } \Delta\hat{\Pi}|_{\eta=\min\{\frac{1-\sigma-r_S}{r_P-r_S},1\}} > 0. \end{cases} \tag{B.47}$$

where $\eta = \frac{1-\sigma-r_S}{r_P-r_S}$ is the solution of $\bar{q}^R(\eta)|_{\xi_t=\xi} = \mu + \epsilon - \xi$ and $\Delta\hat{\Pi}$ is given by equation B.41. □

Proof of Lemma 3.9

Proof. Proof. We first summarize the seller's decision of trust updates under different order quantities and demand realizations. Write $q^R = \bar{q}_t^R(\eta)$.

1. High season ($\xi_t = \xi \implies q^R > \hat{q}^S$).

   (a) $q_t = q^R = \max\{\hat{q}^S, q^R\}$.

i. When $d_t \geq q_t = q^R > \hat{q}_S$, the demand is censored. By definition 3.6, the demand is unambiguous. By Lemma 3.7, $q^R > \hat{q}_S \implies \Delta\hat{\Pi} > 0$ and the seller increases her trust.

ii. When $d_t < \hat{q}_S < q^R = q_t$, the demand is uncensored. We have $\bar{\Pi}^S(q^R) = (1-\gamma)pd_t - (1-\beta)cq^R < (1-\gamma)pd_t - (1-\beta)c\hat{q}^S = \bar{\Pi}^S(\hat{q}^S)$. Therefore, the seller decreases her trust.

iii. When $\hat{q}_S < d_t \leq q^R = q_t$, the demand is uncensored. Similar to the discussion in part 2-1-(c) of the proof of Proposition 3.2, the seller increases her trust iff $d_t > d_+ = \hat{q}^S + (1-r_S)(q^R - \hat{q}^S)$.

(b) $q_t = \hat{q}^S = \min\{\hat{q}^S, q^R\}$.

i. When $d_t < \hat{q}^S < q^R$, similar to case 1-(a)-ii, the seller decreases her trust.

ii. When $d_t \geq \hat{q}^S$, the demand is censored. By definition 3.6, the demand is ambiguous. By Lemma 3.8, the seller decreases her trust in the platform if and only if $r_S < \underline{r}_S$, $r_P > \bar{r}_P$, and $\eta > \bar{\eta}^c$. Otherwise, she decreases her trust.

2. Low season ($\xi_t = -\xi \implies q^R < \hat{q}^S$).

(a) $q_t = \hat{q}^S = \max\{\hat{q}^S, q^R\}$. This case is similar to case 1-(a). When $d_t > q_t = \hat{q}^S > q^R$, the seller increases her trust in the platform; when $d_t < q^R < \hat{q}^S = d_t$, the seller decreases her trust in the platform; when $q^R < d_t \leq \hat{q}^S$, similar to the discussion in part 2-2 of the proof of Proposition 3.2, the seller increases her trust in the platform iff $d_t < d_- = q^R + (1-r_S)(\hat{q}^S - q^R)$.

(b) $q_t = q^R = \min\{\hat{q}^S, q^R\}$. This case is similar to case 1-(b). When $d_t < q_t = q^R < \hat{q}^S$, the seller increases her trust in the platform. When $d_t > q_t$, the demand is censored. By definition 3.6, the demand is unambiguous. By Lemma 3.8, the seller always decreases her trust in the platform because of the low season.

By extracting the cases where the seller increases her trust in the platform from the above discussion, when the seller's trust in the platform is α_t and when the DoS is η, the seller's unconditional probability of increasing her trust in the platform before the demand

realization can be expressed as,

$$
\begin{aligned}
\rho^c(\alpha_t,\eta) &= \mathbb{E}[\rho^c(\alpha_t,\eta)|\xi_t] \\
&= \frac{1}{2}\mathbb{E}[[\rho^c(\alpha_t,\eta)|\xi_t=\xi]|q_t] + \frac{1}{2}\mathbb{E}[[\rho^c(\alpha_t,\eta)|\xi_t=-\xi]|q_t] \\
&= \frac{1}{2}\Big(\alpha_t\mathbb{E}[[\rho^c(\alpha_t,\eta)|\xi_t=\xi,q_t=q^R]|d_t] + (1-\alpha_t)\mathbb{E}[[\rho^c(\alpha_t,\eta)|\xi_t=\xi,q_t=\hat{q}^S]|d_t]\Big) \\
&+ \frac{1}{2}\Big(\alpha_t\mathbb{E}[[\rho^c(\alpha_t,\eta)|\xi_t=-\xi,q_t=q^R]|d_t] + (1-\alpha_t)\mathbb{E}[[\rho^c(\alpha_t,\eta)|\xi_t=-\xi,q_t=\hat{q}^S]|d_t]\Big) \\
&= \frac{1}{2}\Big(\alpha_t\Pr(d_t>d_+|\xi_t=\xi) + (1-\alpha_t)\Pr(d_t\geq\hat{q}^S|\xi_t=\xi)\mathbf{1}_{\eta<\bar{\eta}^c\vee r_S>\underline{r}_S\vee r_P<\bar{r}_P}\Big) \\
&+ \frac{1}{2}\Big(\alpha_t\Pr(d_t<\hat{q}^R|\xi_t=-\xi) + (1-\alpha_t)\Pr(d_t<d_-|\xi_t=-\xi)\Big)
\end{aligned}
$$

Inserting the expressions of $q^R, \bar{q}^S$, d_+ and d_-, we derive

$$
\begin{aligned}
\rho^c(\alpha_t,\eta) = \frac{1}{2}\Bigg(&(\alpha_t\Pr(\epsilon_t>\hat{q}^S+(1-r_S)(q^R-\hat{q}^S)-\mu-\xi_t|\xi_t=\xi) \\
&+(1-\alpha_t)\Pr(\epsilon_t>\hat{q}^S-\mu-\xi_t|\xi_t=\xi)\mathbf{1}_{\eta<\bar{\eta}^c\vee r_S>\underline{r}_S\vee r_P<\bar{r}_P}\Bigg) \\
&+\frac{1}{2}\Bigg(\alpha_t\Pr(\epsilon_t<q^R-\mu-\xi_t|\xi_t=-\xi) \qquad \text{(B.48)} \\
&+(1-\alpha_t)\Pr(\epsilon_t<q^R+(1-r_S)(\hat{q}^S-q^R)-\mu-\xi_t|\xi_t=-\xi)\Bigg) \\
=\frac{1}{2}\Bigg(&\alpha_t\frac{\epsilon-(-\xi-(r_S-1)(\xi+2\eta\epsilon(r_P-r_S))+2\epsilon r_S-\epsilon)}{2\epsilon} \\
&+(1-\alpha_t)\frac{\epsilon-(-\xi-\epsilon+2\epsilon r_S)}{2\epsilon}\mathbf{1}_{\eta<\bar{\eta}^c\vee r_S>\underline{r}_S\vee r_P<\bar{r}_P}\Bigg) \\
&+\frac{1}{2}\Bigg(\alpha_t\frac{-\epsilon+2\epsilon(\eta r_P+(1-\eta)r_S)-(-\epsilon)}{2\epsilon} \\
&+(1-\alpha_t)\frac{(\xi-r_S(\xi+2\eta\epsilon(r_S-r_P)-2\epsilon)-\epsilon)-(-\epsilon)}{2\epsilon}\Bigg). \qquad \text{(B.49)}
\end{aligned}
$$

- When $\eta>\bar{\eta}^c$, $r_s<\underline{r}_S$, and $r_P>\bar{r}_P$, equation B.49 is simplified to

$$
\rho^c(\alpha_t,\eta) = \frac{\xi-r_S(\xi+2\eta\epsilon(r_S-r_P)-2\epsilon)-\alpha_t(\xi+2r_S(\epsilon-\xi)-2\epsilon)}{4\epsilon}. \qquad \text{(B.50)}
$$

  Because $r_P>r_S$, $\rho^c(\alpha_t,\eta)$ monotonically increases in η. Because $\epsilon>\xi$ and by assump-

tion 3.1,

$$\xi + 2r_S(\epsilon - \xi) - 2\epsilon < \xi + 2r_S(\epsilon - \xi) - 2\epsilon\big|_{r_S = 1 - \frac{\xi}{2\epsilon}} = \frac{\xi(\xi - 2\epsilon)}{\epsilon} < 0,$$

and $\rho^c(\alpha_t, \eta)$ monotonically increases in α_t.

- When $\eta < \bar{\eta}^c$ or $r_s < \underline{r}_S$ or and $r_P > \bar{r}_P$, equation B.49 is simplified to

$$\rho^c(\alpha_t, \eta) = \frac{r_S\left(-\xi + 2\eta\epsilon\left(r_P - r_S\right)\right) + 2\xi\left(r_S - 1\right)\alpha_t + 2(\xi + \epsilon)}{4\epsilon}. \tag{B.51}$$

  Because $r_P > r_S$, $\rho^c(\alpha_t, \eta)$ monotonically increases in η. Because $r_S < 1$, $\rho^c(\alpha_t, \eta)$ monotonically decreases in α_t.

□

Proof of Proposition 3.10 By the proof of Proposition 3.2, in the steady state, $\alpha_t = \bar{\alpha}^c = \rho^c(\eta, \alpha_t)$.

When $\eta > \bar{\eta}^c$, $r_s < \underline{r}_S$, and $r_P > \bar{r}_P$, by equation B.51,

$$\bar{\alpha}^c = \frac{\xi - r_S\left(\xi + 2\eta\epsilon\left(r_S - r_P\right) - 2\epsilon\right)}{\xi + 2r_S(\epsilon - \xi) + 2\epsilon}. \tag{B.52}$$

Because $r_P > r_S$, $\bar{\alpha}^c$ monotonically increases in η. Consider

$$\frac{d}{d\sigma}\bar{\alpha}^c = \frac{d}{d\sigma}\left(-\frac{r_S\left(-2\eta r_P + 2\eta r_S + \sigma - 2\right) - \sigma}{-2(\sigma - 1)r_S + \sigma + 2}\right) = \frac{A_{11}}{\left(-2(\sigma - 1)r_S + \sigma + 2\right)^2}.$$

where $A_{11} = 2r_S\left(\eta\left(2r_S - 1\right)\left(r_P - r_S\right) + r_S - 1\right) + 2$. We show that $\bar{\alpha}^c$ increasing in $\sigma = \frac{\xi}{\epsilon}$ by showing $A_{11} > 0$. Consider

$$\frac{dA_{11}}{dr_P} = 2\eta r_S\left(2r_S - 1\right).$$

- If $r_S > \frac{1}{2}$, A_{11} monotonically increases in r_P. Therefore,

$$A_{11} > A_{11}|_{r_P = r_S} = 2\left(r_S - 1\right)r_S + 2 > 2\left(r_S - 1\right)r_S + 2\big|_{r_S = \frac{1}{2}} = \frac{3}{2} > 0.$$

- If $r_S < \frac{1}{2}$, A_{11} monotonically decreases in r_P. Therefore,

$$\begin{aligned} A_{11} &> A_{11}|_{r_P = r_S + \frac{\sigma}{2}} \\ &= r_S\left(-\eta\sigma + 2(\eta\sigma+1)r_S - 2\right) + 2 \\ &\geq \min_{r_S}\left[r_S\left(-\eta\sigma + 2(\eta\sigma+1)r_S - 2\right) + 2\right] \\ &= \frac{\eta\sigma(12-\eta\sigma)+12}{8\eta\sigma+8} > 0. \end{aligned}$$

  The final equality follows from the fact that $r_S\left(-\eta\sigma + 2(\eta\sigma+1)r_S - 2\right) + 2$ is minimized at $r_S = \frac{\eta\sigma+2}{4(\eta\sigma+1)} < \frac{1}{2}$.

When $\eta < \bar{\eta}^c$ or $r_s < \underline{r}_S$, or $r_P > \bar{r}_P$, by equation B.50,

$$\bar{\alpha}^c = \frac{2(\xi+\epsilon) + r_S\left(-\xi + 2\eta\epsilon\left(r_P - r_S\right)\right)}{2\left(\xi - \xi r_S + 2\epsilon\right)}. \tag{B.53}$$

Because $r_P > r_S$, $\bar{\alpha}^c$ monotonically increases in η. $\bar{\alpha}^c$ increases in σ because

$$\frac{d}{d\sigma}\bar{\alpha}^c = \frac{d}{d\sigma}\frac{r_S\left(-2\eta r_P + 2\eta r_S + \sigma\right) - 2(\sigma+1)}{2\sigma r_S - 2(\sigma+2)} = \frac{\eta\left(r_S-1\right)r_S\left(r_P - r_S\right)+1}{\left(\sigma\left(-r_S\right)+\sigma+2\right)^2} > 0,$$

where the inequality is due to $r_S < 1, r_P - r_S < 1, \eta < 1$ and $1 - r_S < 1$. □

Proof of Proposition 3.11

Proof. Proof. The proof is completed in two steps. In the first step, we derive the optimal DoS under a censored demand (η_c^*). In the second step, we analyze the monotonicity of η_c^*.

Step 1. By the Newsvendor setup, regardless of the season, the platform's expected profit is concave in the order quantity with a maximum achieved at q_t^P defined in equation 3.5. In other words, $\mathbb{E}[\Pi^P(q_t, d_t)|\xi_t]$ monotonically increases in q_t for $q_t \leq q_t^P$. Because

$$q_t^S \leq \bar{q}_t^R(\eta) = \eta q_t^P + (1-\eta)q_t^S \leq q_t^P$$

and $q_t^R(\eta)$ monotonically increases in η, $\mathbb{E}[\Pi^P(\bar{q}_t^R(\eta), d_t)|\xi_t]$, as well as

$$\mathbb{E}[\Pi^P(\bar{q}^R(\eta), d_t)] = \mathbb{E}[\mathbb{E}[\Pi^P(q_t^R(\eta), d_t)|\xi_t]]$$

monotonically increases in η. By equation B.6, the platform's expected steady-state profit when the demand is censored can be expressed as

$$\bar{\Pi}^P(\eta) = \mathbb{E}[\Pi^P(\hat{q}_S, d_t)] + \bar{\alpha}(\eta)\left(\mathbb{E}[\Pi^P(\bar{q}^R(\eta), d_t)] - \mathbb{E}[\Pi^P(\hat{q}_S, d_t)]\right). \tag{B.54}$$

Because both $\mathbb{E}[\Pi^P(q^R(\eta), d_t)]$ and $\bar{\alpha}(\eta)$ are monotonically increases in η by the above discussion and Proposition 3.10, and because $\mathbb{E}[\Pi^P(\hat{q}_S, d_t)]$ is a constant, $\bar{\Pi}^P(\eta)$ piecewisely increases in η. Hence, $\eta_c^* \in \{\bar{\eta}^c, 1\}$.

Step 2. We look at the nontrivial case where $\eta_c^* = \bar{\eta}^c$. We further consider the following two cases based on the expression of $\bar{\eta}^c$ (equation B.47).

- $\eta_c^* = \bar{\eta}^c = \bar{\eta}_1^c$. In this case, $\bar{\eta}_c$ is monotonically decreasing in r_P by its expression (equation B.44). Moreover, $\bar{\eta}^c = \bar{\eta}_1^c$ implies

$$\bar{q}_t^R|_{\xi_t=\xi,\eta=\bar{\eta}_c} < \mu + \epsilon - \xi \implies \sigma < 2 - 2r_S\,(3 - 2r_S)\,.$$

  We aim to show

$$\frac{d\bar{\eta}_1^c}{dr_S} = \frac{4\,(r_P - r_S)^{\,2} - 4\,(r_P - 1)\,r_P + \sigma}{2\,(r_P - r_S)^{\,2}} = \frac{4r_S^2 + 4r_P(1 - 2r_S) - \sigma}{2\,(r_P - r_S)^{\,2}} = \frac{A_{12}}{2\,(r_P - r_S)^{\,2}} > 0.$$

  - If $r_S > \frac{1}{2}$, A_{12} monotonically decreases in r_P and

$$\begin{aligned} A_{12} &> A_{12}|_{r_P = r_S + \frac{\sigma}{2}} = (1 - 4r_S)\sigma - 4r_S\,(r_S - 1) \\ &> (1 - 4r_S)\sigma - 4r_S\,(r_S - 1)\,|_{\sigma = 2 - 2r_S(3 - 2r_S)} = 2\,(1 - r_S)\left(4r_S^2 + (2r_S - 1)^2\right) > 0. \end{aligned}$$

  - If $r_S < \frac{1}{2}$, A_{12} monotonically increases in r_P and

$$\begin{aligned} A_{12} &> A_{12}|_{r_P = r_S} = 4\,(1 - r_S)\,r_S - \sigma \\ &> 4\,(1 - r_S)\,r_S - \sigma|_{r_S = \frac{\sigma}{2}} = 4\,(1 - r_S)\,r_S - \sigma|_{r_S = 1 - \frac{\sigma}{2}} = \sigma(1 - \sigma) > 0. \end{aligned}$$

  - If $r_S = \frac{1}{2}$, $A_{12} = 1 - \sigma > 0$.

- $\eta_c^* = \bar{\eta}^c = \bar{\eta}_2^c$. By the expression of $\bar{\eta}_2^c$ (equation B.46), $\bar{\eta}_2^c$ monotonically decreases in r_P. It remains to show that $\bar{\eta}_2^c$ is increasing in r_S. Because the denominator of $\bar{\eta}_2^c$ monotonically decreases in r_S, it suffices to show that the numerator of $\bar{\eta}_2^c$ monotonically

increases in r_S. In other words, we show that

$$\begin{aligned}
&\frac{d}{dr_S}\left(\sqrt{2}\sqrt{\left(\sigma-2\left(r_S-1\right)^{2}\right)^{2}+4\left(r_S-1\right)^{2}r_S^2}-2\left(r_S-1\right)\left(2r_S-1\right)\right)\\
&=-\frac{4\sqrt{2}\left(r_S-1\right)\left(r_S\left(5-4r_S\right)+\sigma-2\right)}{\sqrt{\left(\sigma-2\left(r_S-1\right)^{2}\right)^{2}+4\left(r_S-1\right)^{2}r_S^2}}-8r_S+6>0\\
&\iff \left(6-8r_S\right)^{2}\left(\left(\sigma-2\left(r_S-1\right)^{2}\right)^{2}+4\left(r_S-1\right)^{2}r_S^2\right)\\
&-\left(4\sqrt{2}\left(r_S-1\right)\left(r_S\left(5-4r_S\right)+\sigma-2\right)\right)^{2}>0\\
&\iff 4\left(2r_S+\sigma-2\right)\left(2(4\sigma-2)r_S^2+(6-8\sigma)r_S+\sigma-2\right)>0.
\end{aligned}$$

The last inequality holds because $2r_S+\sigma-2=(\sigma-1)+(2r_S-1)<0$ as $r_S<1-\frac{\sigma}{2}$, and because

$$\begin{aligned}
&2(4\sigma-2)r_S^2+(6-8\sigma)r_S+\sigma-2\\
&<\max\{\sigma-2,2(4\sigma-2)r_S^2+(6-8\sigma)r_S+\sigma-2|_{r_S=\frac{1}{2}}\}\\
&=\max\{\sigma-2,-\sigma\}<0,
\end{aligned}$$

where the second line is due to $2(4\sigma-2)r_S^2+(6-8\sigma)r_S+\sigma-2$ is maximized at its endpoints.

□

Proof of Proposition 3.12

Proof. Proof. We prove this result in two steps. In the first step, we derive the expression of $\bar{\sigma}$. In the second step, we analyze the sensitivity of $\bar{\sigma}$.

Step 1. By Proposition 3.4 and Proposition 3.11, $r_S>\underline{r}_S$ or $r_P<\bar{r}_P$ implies $\eta_c^*=1\geq\eta_P^*$ and

$$\mathbb{E}[\Pi^P(\bar{q}_t^R(\eta_c^*),d_t)]=\mathbb{E}[\Pi^P(q_t^P,d_t)]\geq\mathbb{E}[\Pi^P(\bar{q}_t^R(\eta_P^*),d_t)].$$

By equation B.6 and B.54, it suffices to show that for any r_S and r_P, there exists some $\frac{\xi}{\epsilon}$ such that,

$$\bar{\alpha}^c(\eta_c^*)>\bar{\alpha}(\eta_P^*)=\rho(\eta_P^*).$$

Here, $\bar{\alpha}^c(\eta)$ is given by equation B.53 and $\bar{\alpha}(\eta_P^*)=\rho(\eta_P^*)$ is given by equation 3.15.

- $r_S<\frac{1}{2}$. By equation 3.15, $\rho(\eta_P^*)<\frac{1}{2}+\frac{\xi}{4\epsilon}$. Therefore, $\bar{\alpha}^c(\eta_c^*)>\frac{1}{2}+\frac{\xi}{4\epsilon}\implies\bar{\alpha}^c(\eta_c^*)>$

$\bar{\alpha}(\eta_P^*)$. Specifically, replacing $\frac{\xi}{\epsilon}$ with σ,

$$\bar{\alpha}^c(\eta_c^*) > \frac{1}{2}+\frac{\xi}{4\epsilon} \iff \frac{-r_S\left(4r_P+\sigma^2\right)+4r_S^2+\sigma^2}{4\sigma r_S-4(\sigma+2)} > 0 \iff \sigma < \bar{\sigma}_1 = 2\sqrt{\frac{r_S\left(r_P-r_S\right)}{1-r_S}}. \tag{B.55}$$

- $r_S \geq \frac{1}{2}$. By proposition 3.4, $\eta_P^* = 1$. Therefore,

$$\begin{aligned}
&\bar{\alpha}^c(\eta_c^*) - \bar{\alpha}^c(\eta_P^*) = \bar{\alpha}^c(1) - \bar{\alpha}^c(1) \\
&= \frac{-r_S\left(4r_P+\sigma^2\right)+4r_S^2+\sigma^2}{4\sigma r_S-4(\sigma+2)} + \frac{1}{2}\left(1-2r_S\right)\left(r_P-r_S\right) > 0 \\
&\implies \sigma < \bar{\sigma}_2 \\
&= \sqrt{\left(r_P-r_S\right)\left(r_P\left(1-2r_S\right)^2 - r_S\left(1-2r_S\right)^2+4\right)} + r_P\left(1-2r_S\right) + r_S\left(2r_S-1\right).
\end{aligned} \tag{B.56}$$

Define

$$\bar{\sigma} = \min\{\bar{\sigma}_1, \bar{\sigma}_2\} = \begin{cases} \bar{\sigma}_1 & r_S < \frac{1}{2} \\ \bar{\sigma}_2 & r_S \geq \frac{1}{2} \end{cases}. \tag{B.57}$$

$\bar{\sigma}$ is continuous because $\bar{\sigma}_1 = \bar{\sigma}_2$ at $r_S = \frac{1}{2}$.

Step 2.

- When $r_S < \frac{1}{2}$, $\bar{\sigma} = \bar{\sigma}_1$ monotonically increases in r_P by its expression. When $r_S \geq \frac{1}{2}$,

$$\begin{aligned}
\frac{d\bar{\sigma}_2}{dr_P} &= \frac{r_P\left(1-2r_S\right)^2 - r_S\left(1-2r_S\right)^2+2}{\sqrt{\left(r_P-r_S\right)\left(r_P\left(1-2r_S\right)^2 - r_S\left(1-2r_S\right)^2+4\right)}} - 2r_S + 1 \\
&= \frac{(1-2r_S)\bar{\sigma}_2+2}{\sqrt{\left(r_P-r_S\right)\left(r_P\left(1-2r_S\right)^2 - r_S\left(1-2r_S\right)^2+4\right)}} > 0.
\end{aligned}$$

  The last inequality is because of $\bar{\sigma}_2 < 1$ and $1-2r_S > 1-2 = -1$. In sum, $\bar{\sigma}$ increases in r_P.

- First, observe that $\bar{\sigma}_1$ is concave in r_S. In detail, consider

$$\frac{d^2\bar{\sigma}_1}{dr_S^2} = \frac{-6r_Pr_S^2 + r_P^2\left(4r_S-1\right) - \left(\left(r_S-4\right)r_S^3\right)}{4\left(r_S-1\right)^4\left(\frac{r_S(r_S-r_P)}{r_S-1}\right)^{3/2}}.$$

  Let $A_{13}(r_P) = -6r_Pr_S^2 + r_P^2\left(4r_S-1\right) - \left(\left(r_S-4\right)r_S^3\right)$. Observe that $A_{13}(r_S) = -\left(r_S-1\right)^2r_S^2 < 0, A_{13}(1) = -\left(r_S-1\right)^4 < 1$. Furthermore, $\frac{dA_{13}(r_P)}{dr_P}|_{r_P=r_S} = -2(1-$

$r_S)r_S < 0$. Because A_{13} is quadratic in r_P, we must have $A_{13} < 0$ for all $r_P \in (r_S, 1)$, which further implies that $\bar{\sigma}_1$ is concave in r_S.

Second, observe that $\bar{\sigma}_1|_{r_S=0} = 0$ and $\bar{\sigma}_1|_{r_S=r_P} = 0$. Therefore, $\bar{\sigma}_1$ first increase then decrease in r_S in $(0, r_P)$.

Third, we show that $\bar{\sigma}_2$ monotonically decreases in r_S. To see this, consider

$$\begin{aligned}\frac{d\bar{\sigma}_2}{dr_S} &= \frac{-6\left(2r_P+1\right)r_S^2 + 4r_P\left(r_P+2\right)r_S - r_P\left(2r_P+1\right) + 8r_S^3 + r_S - 2}{\sqrt{\left(r_P - r_S\right)\left(r_P\left(1-2r_S\right)^2 - r_S\left(1-2r_S\right)^2 + 4\right)}} - 2r_P + 4r_S - 1 \\ &= -t + \frac{\left(2r_S-1\right)\left(\sqrt{t\left(t\left(1-2r_S\right)^2+4\right)} - t\left(2r_S-1\right) + 2t^2\right) - 2}{\sqrt{t\left(r_P\left(1-2r_S\right)^2 - r_S\left(1-2r_S\right)^2 + 4\right)}},\end{aligned}$$

where $t = r_P - r_S \in (0, 1 - r_S)$. To show $\frac{d\bar{\sigma}_2}{dr_S} < 0$, it suffices to show that

$$A_{14} = \left(2r_S-1\right)\left(\sqrt{t\left(t\left(1-2r_S\right)^2+4\right)} - t\left(2r_S-1\right) + 2t^2\right) - 2$$

is always negative. We show this in two steps.

Step 1. $\sqrt{t\left(t\left(1-2r_S\right)^2+4\right)} < 2 - t(2r_S - 1)$. To see this,

$$\begin{aligned}&\sqrt{t\left(t\left(1-2r_S\right)^2+4\right)} < 2 - t(2r_S-1) \\ &\iff t\left(t\left(1-2r_S\right)^2+4\right) < \left(2-t(2r_S-1)\right)^2 \iff 2tr_S < 1.\end{aligned}$$

Because $t < 1 - r_S$, $2tr_S < 2(1-r_S)r_S < \frac{1}{2} < 1$ and the above inequality holds.

Step 2. By step 1,

$$\begin{aligned}A_{14} &< (2r_S-1)(2 - 2t(2r_S-1) + 2t^2) - 2 \\ &< \max\{2\left((2r_S-1)(1-(1-r_S)(2r_S-1) + (1-r_S)^2) - 1\right), 2(2r_S-1) - 2\} \\ &= \max\{2\left((2r_S-1)(r_S^2 - r_S + 1) - 1\right), 2(r_S-1)\} < 0,\end{aligned}$$

where the second inequality is because $2 - 2t(2r_S-1) + 2t^2$ is maximized at its endpoints, and the last inequality is because $2r_S - 1 < 1$ and $r_S^2 - r_S + 1 < 1$.

Finally, because $\bar{\sigma}_1$ is concave in r_S and $\bar{\sigma}_2$ monotonically decreases in r_S, and by the definition of $\bar{\sigma}$ (equation B.57), $\bar{\sigma}$ first increases in r_S and then decreases.

□

Bibliography

[1] Vibhanshu Abhishek, Kinshuk Jerath, and Z John Zhang. Agency selling or reselling? channel structures in electronic retailing. Management Science, 62(8):2259–2280, 2016.

[2] Arzum Akkaş and Vishal Gaur. Om forum—reducing food waste: An operations management research agenda. Manufacturing & Service Operations Management, 24(3):1261–1275, 2022.

[3] Arzum Akkaş and Dorothee Honhon. Shipment policies for products with fixed shelf lives: Impact on profits and waste. Manufacturing & Service Operations Management, 24(3):1611–1629, 2022.

[4] Arzum Akkaş and Nachiketa Sahoo. Reducing product expiration by aligning sales-force incentives: A data-driven approach. Production and Operations Management, 29(8):1992–2009, 2020.

[5] Amazon. Help for Amazon Sellers, 2024. Accessed: 2025-10-01.

[6] Aramark. Aramark Study: Trayless Dining a Win/Win, Jul 2008. Accessed: 2025-10-01.

[7] Kenneth J Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. Econometrica: Journal of the Econometric Society, pages 250–272, 1951.

[8] Ekaterina Astashkina, Elena Belavina, and Simone Marinesi. The environmental impact of the advent of online grocery retailing. Available at SSRN 3358664, 2019.

[9] Zumbul Atan, Dorothee Honhon, and Xiajun Amy Pan. Displaying and discounting perishables: Impact on retail profits and waste. Available at SSRN 4369956, 2023.

[10] Tal Avinadav, Tatyana Chernonog, and Noam Shamir. Information provision from a platform to competing sellers: The role of strategic ambiguity. Manufacturing & Service Operations Management, 27(1):269–286, 2025.

[11] Elena Belavina. Grocery store density and food waste. Manufacturing & Service Operations Management, 23(1):1–18, 2021.

[12] Elena Belavina, Karan Girotra, and Ashish Kabra. Online grocery retail: Revenue models and environmental impact. Management Science, 63(6):1781–1799, 2017.

[13] Saif Benjaafar, Guangwen Kong, Xiang Li, and Costas Courcoubetis. Peer-to-peer product sharing: Implications for ownership, usage, and social welfare in the sharing economy. Management Science, 65(2):477–493, 2019.

[14] Dimitris Bertsimas and Romy Shioda. Restaurant revenue management. Operations Research, 51(3):472–486, 2003.

[15] Yiwen Bian, Hang Wang, Shuai Yan, and Xiaohua Han. Platform channel strategy of selling format: Effect of asymmetric service information. Transportation Research Part E: Logistics and Transportation Review, 185:103508, 2024.

[16] Kostas Bimpikis and Giacomo Mantegazza. Strategic release of information in platforms: Entry, competition, and welfare. Operations Research, 71(5):1619–1635, 2023.

[17] Kostas Bimpikis, Yiangos Papanastasiou, and Wenchang Zhang. Information provision in two-sided platforms: Optimizing for supply. Management Science, 70(7):4533–4547, 2024.

[18] John Blundell, Cees De Graaf, Toine Hulshof, Susan Jebb, Barbara Livingstone, Anne Lluch, David Mela, Samir Salah, Ewoud Schuring, Henk Van Der Knaap, et al. Appetite control: methodological aspects of the evaluation of foods. Obesity reviews, 11(3):251–270, 2010.

[19] Richard Boso, Mahyar Eftekhar, Xinming Liu, and Yanchong Zheng. Reducing food waste in production: A field experiment in ghana. Available at SSRN 5063255, 2024.

[20] Gerard P Cachon, Tolga Dizdarer, and Gerry Tsoukalas. Pricing control and regulation on online service platforms. Management Science, 2025.

[21] Gérard P Cachon and Fuqiang Zhang. Procuring fast delivery: Sole sourcing with information asymmetry. Management Science, 52(6):881–896, 2006.

[22] John T Cacioppo and Richard E Petty. Effects of message repetition and position on cognitive response, recall, and persuasion. Journal of Personality and Social Psychology, 37(1):97, 1979.

[23] Colin Camerer and Teck Hua Ho. Experience-weighted attraction learning in normal form games. Econometrica, 67(4):827–874, 1999.

[24] Gabriel Carroll. Robustness and linear contracts. American Economic Review, 105(2):536–563, 2015.

[25] Juan Camilo Castillo, Dan Knoepfle, and E Glen Weyl. Matching and pricing in ride hailing: Wild goose chases and how to solve them. Management Science, 71(5):4377–4395, 2025.

[26] Amazon Seller Central. Auto-Replenishment, 2025. Accessed: 2025-09-12.

[27] Yevvon Yi-Chi Chang. All you can eat or all you can waste? effects of alternate serving styles and inducements on food waste in buffet restaurants. Current Issues in Tourism, 25(5):727–744, 2022.

[28] Yevvon Yi-Chi Chang, Jo-Hui Lin, and Chiu-Han Hsiao. Examining effective means to reduce food waste behaviour in buffet restaurants. International Journal of Gastronomy and Food Science, 29:100554, 2022.

[29] Manlu Chen, Ming Hu, and Jianfu Wang. Food delivery service and restaurant: Friend or foe? Management Science, 68(9):6539–6551, 2022.

[30] Nan Chen and Hsin-Tien Tsai. Price competition under information (dis) advantage. Available at SSRN 4420175, 2023.

[31] Austin Clowes, Peter Mitchell, and Graig Hanson. The Business Case for Reducing Food Loss and Waste: Hotels, 2018. Accessed: 2024-09-12.

[32] Lauren Collines. The Hottest Restaurant in France Is an All-You-Can-Eat Buffet, 2024. Accessed: Sep 12, 2024.

[33] Yann Cornil. Mind over stomach: A review of the cognitive drivers of food satiation. Journal of the Association for Consumer Research, 2(4):419–429, 2017.

[34] Ayşen Coşkun and Raife Meltem Yetkin Özbük. What influences consumer food waste behavior in restaurants? an application of the extended theory of planned behavior. Waste Management, 117:170–178, 2020.

[35] Claudia Cozzio, Oksana Tokarchuk, and Oswin Maurer. Minimising plate waste at hotel breakfast buffets: an experimental approach through persuasive messages. British Food Journal, 123(9):3208–3227, 2021.

[36] Pascale Crama, Yangfang Helen Zhou, and Manman Wang. Impact of cosmetic standards on food loss. Available at SSRN 4452800, 2023.

[37] Zachary Crockett. The economics of all-you-can-eat buffets, 2020. Accessed: 2024-09-12.

[38] Zoë Cullen and Chiara Farronato. Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms. Management Science, 67(7):3985–4003, 2021.

[39] Helen Davidson. China to bring in law against food waste with fines for promoting overeating. The Guardian, 2020. Accessed: 2024-09-12.

[40] Yiting Deng, Christopher S Tang, Wei Wang, and Onesun Steve Yoo. Can third-party sellers benefit from a platform's entry to the market? Service Science, 15(4):233–249, 2023.

[41] Sara Dolnicar, Emil Juvan, and Bettina Grün. Reducing the plate waste of families at hotel buffets–a quasi-experimental field study. Tourism Management, 80:104103, 2020.

[42] Neda Ebrahim-Khanjari, Wallace Hopp, and Seyed MR Iravani. Trust and information sharing in supply chains. Production and Operations Management, 21(3):444–464, 2012.

[43] Brenna Ellison, Erica Whitney Nehrling, Cassandra J Nikolaus, and Brittany RL Duff. Evaluation of a food waste reduction campaign in a univ. dining hall. Journal of Nutrition Education and Behavior, 49(7):S9–S10, 2017.

[44] Emplicit. Amazon vs Walmart: Marketplace Fee Comparison, 2025. Accessed: 2025-09-12.

[45] Marjorie R Freedman and Carolina Brochado. Reducing portion size reduces food intake and plate waste. Obesity, 18(9):1864–1866, 2010.

[46] Red Stag Fulfillment. How Many Amazon Marketplace Sellers Use Fulfillment by Amazon (FBA)? 2025 Statistics, 2025. Accessed: 2025-09-12.

[47] Fei Gao and Xuanming Su. Omnichannel service operations with online and offline self-order technologies. Management Science, 64(8):3595–3608, 2018.

[48] Mattia Garutti, Marianna Sirico, Claudia Noto, Lorenzo Foffano, Mark Hopkins, and Fabio Puglisi. Hallmarks of appetite: A comprehensive review of hunger, appetite, satiation, and satiety. Current Obesity Reports, 14(1):1–20, 2025.

[49] Canran Gong, Joshua Ignatius, Huaming Song, Junwu Chai, and Steven James Day. The impact of platform's information sharing on manufacturer encroachment and selling format decision. European Journal of Operational Research, 317(1):141–155, 2024.

[50] Jenny Gustavsson, Christel Cederberg, Ulf Sonesson, Robert Van Otterdijk, and Alexandre Meybeck. Global food losses & food waste, 2011. Accessed: 2025-10-01.

[51] Albert Y Ha, Huajiang Luo, and Weixin Shang. Supplier encroachment, information sharing, and channel structure in online retail platforms. Production and Operations Management, 31(3):1235–1251, 2022.

[52] Albert Y Ha and Shilu Tong. Contracting and information sharing under supply chain competition. Management science, 54(4):701–715, 2008.

[53] Albert Y Ha, Shilu Tong, and Yunjie Wang. Channel structures of online retail platforms. Manufacturing & Service Operations Management, 24(3):1547–1561, 2022.

[54] Hanna Hałaburda and Yaron Yehezkel. Platform competition under asymmetric information. American Economic Journal: Microeconomics, 5(3):22–68, 2013.

[55] Pelle Guldborg Hansen, Andreas Maaløe Jespersen, and Laurits Rohden Skov. Size matter! a choice architectural field experiment in reducing food waste. Menu. The Journal of Food and Hospitality Research, 4:11–15, 2015.

[56] Linda Himelstein. Hotel buffets, a culprit of food waste, get downsized, 2017. Accessed: 2024-09-12.

[57] Chris Hock. Mastering 5 Critical Elements in Buffet Restaurant Design and Construction, 2023. Accessed: 2024-09-12.

[58] Stephen S Holden, Natalina Zlatevska, and Chris Dubelaar. Whether smaller plates reduce consumption depends on who's serving and who's looking: a meta-analysis. Journal of the Association for Consumer Research, 1(1):134–146, 2016.

[59] Dorothee Honhon, Iva Rashkova, and P Shahsavand. It's still good! expiration dates & food waste, 2025. Working paper, UTD.

[60] M. Hopkins, K. Beaulieu, C. Gibbons, J. Stubbs, and G. Finlayson. The Control of Food Intake in Humans. In K. R. Feingold, B. Anawalt, A. Boyce, and et al., editors, Endotext. MDText.com, Inc., South Dartmouth (MA), 2022.

[61] Shu Hu and Ke Fu. Equilibrium analysis of the seller's fulfillment channels and sales channels. European Journal of Operational Research, 320(3):544–558, 2025.

[62] Shu Hu, Mike Mingcheng Wei, and Shiliang Cui. The role of product and market information in an online marketplace. Production and Operations Management, 32(10):3100–3118, 2023.

[63] Yihong Hu, Guo Li, Mengqi Liu, and Shengnan Qu. Information sharing and personalized pricing in online platforms. Production and Operations Management, page 10591478231225178, 2024.

[64] Hu Huang, Tian Li, and Hongtao Zhang. Information sharing in the online marketplace when sellers incur nonlinear costs. Available at SSRN 3883770, 2021.

[65] IBISWorld. Buffet Restaurants Industry in the US - Market Research Report, 2022. Accessed: 2024-09-12.

[66] Rafay Ishfaq and Uzma Raja. Evaluation of order fulfillment options in retail supply chains. Decision Sciences, 49(3):487–521, 2018.

[67] Garud Iyengar, Yuanzhe Ma, Thomas J Rivera, Fahad Saleh, and Jay Sethuraman. The distributional effects of'fulfilled by amazon'(fba). Available at SSRN 4365855, 2023.

[68] Baojun Jiang, Kinshuk Jerath, and Kannan Srinivasan. Firm strategies in the "mid tail" of platform-based retailing. Marketing Science, 30(5):757–775, 2011.

[69] Gapper John. The strange new economy of all-you-can-eat buffets, 2024. Accessed: 2024-09-12.

[70] Emil Juvan, Bettina Grün, Petra Zabukovec Baruca, and Sara Dolnicar. Drivers of plate waste at buffets: A comprehensive conceptual model based on observational data and staff insights. Annals of Tourism Research Empirical Insights, 2(1):100010, 2021.

[71] Emil Juvan, Bettina Grün, and Sara Dolnicar. Biting off more than they can chew: Food waste at hotel breakfast buffets. Journal of Travel Research, 57(2):232–242, 2018.

[72] Steffen Kallbekken and Håkon Sælen. 'nudging'hotel guests to reduce food waste as a win–win environmental measure. Economics Letters, 119(3):325–327, 2013.

[73] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. American Economic Review, 101(6):2590–2615, 2011.

[74] Yash Kanoria and Daniela Saban. Facilitating the search for partners on matching platforms. Management Science, 67(10):5990–6029, 2021.

[75] N Bora Keskin, Chenghuai Li, and Jing-Sheng Jeannette Song. The blockchain newsvendor: Value of freshness transparency and smart contracts. Management Science, 2024.

[76] Kiho Kim and Stevia Morawski. Quantifying the impact of going trayless in a university dining hall. Journal of Hunger & Environmental Nutrition, 7(4):482–486, 2012.

[77] Ismail Kirci, Dorothee Honhon, and Alp Muharremoglu. Do you lose when you sell loose? impact of unpackaged sales on profit and waste. Available at SSRN 4585591, 2023.

[78] Rhea Kundamal. The Unsustainability of Buffet Food Waste, 2023. Accessed: 2025-4-24.

[79] ChenFeng Kuo and Yahui Shih. Gender differences in the effects of education and coercion on reducing buffet plate waste. Journal of Foodservice Business Research, 19(3):223–235, 2016.

[80] Guoming Lai, Huihui Liu, Wenqiang Xiao, and Xinyi Zhao. "fulfilled by amazon": A strategic perspective of competition at the e-commerce platform. Manufacturing & Service Operations Management, 24(3):1406–1420, 2022.

[81] Guo Li, Lin Tian, and Hong Zheng. Information sharing in an online marketplace with co-opetitive sellers. Production and Operations Management, 30(10):3713–3734, 2021.

[82] Jun Li and Serguei Netessine. Higher market thickness reduces matching rate in online platforms: Evidence from a quasiexperiment. Management Science, 66(1):271–289, 2020.

[83] Jun Li, Wenjing Shen, Yi Liao, Gangshu Cai, and Xiangfeng Chen. The fulfillment service in online marketplaces. European Journal of Operational Research, 315(3):1139–1152, 2024.

[84] Yao Li, Lauren Xiaoyuan Lu, Jan A Van Mieghem, and Wenqiang Xiao. Fulfillment by amazon as a strategic lever: Anticompetitive or welfare enhancing? Available at SSRN 5376166, 2025.

[85] E Libotte, Michael Siegrist, and Tamara Bucher. The influence of plate size on meal composition. literature review and experiment. Appetite, 82:91–96, 2014.

[86] Jill Liliedahl. The Challenges with Forecasting and Restocking Inventory on Amazon, 2019. Accessed: 2025-09-12.

[87] Zekun Liu, Dennis J Zhang, and Fuqiang Zhang. Information sharing on retail platforms. Manufacturing & Service Operations Management, 23(3):606–619, 2021.

[88] Jennifer K MacCormack and Kristen A Lindquist. Feeling hangry? when hunger is conceptualized as emotion. Emotion, 19(2):301, 2019.

[89] Mashed. How All-You-Can-Eat Buffets Really Make Their Money, 2019. Accessed: 2024-09-12.

[90] K Murali and I Alev. Implications of eliminating grading in food supply chains, 2019. Working paper, Oregon State University, Corvallis.

[91] Arishekar N. Walmart Fulfillment Fees in 2025 Comprehensive Seller Guide, 2025. Accessed, 2025-09-12.

[92] Paul Nowell. 24 Amazon Seller Statistics 2025: Facts, Trends & Revenue, 2025. Accessed: 2025-09-12.

[93] Yu Nu, Elena Belavina, and Karan Girotra. Using artificial intelligence to reduce food waste. Available at SSRN 4826777, 2024.

[94] Özalp Özer, A Serdar Şimşek, Xiaoxi Zhao, Ethan Dee, and Vivian Yu. Measuring the efficacy of amazon's recommendation systems. In Tutorials in Operations Research: Smarter Decisions for a Better World, pages 224–243. INFORMS, 2024.

[95] Özalp Özer, Upender Subramanian, and Yu Wang. Information sharing, advice provision, or delegation: What leads to higher trust and trustworthiness? Management Science, 64(1):474–493, 2018.

[96] Özalp Özer and Yanchong Zheng. Establishing trust and trustworthiness for supply chain information sharing. In Handbook of information exchange in supply chain management, pages 287–312. Springer, 2016.

[97] Özalp Özer and Yanchong Zheng. Trust and trustworthiness. Handbook of Behavioral Operations, 2019.

[98] Özalp Özer, Yanchong Zheng, and Kay-Yut Chen. Trust in forecast information sharing. Management science, 57(6):1111–1137, 2011.

[99] Effie Papargyropoulou, Julia K Steinberger, Nigel Wright, Rodrigo Lozano, Rory Padfield, and Zaini Ujang. Patterns and causes of food waste in the hospitality and food service sector: Food waste prevention insights from malaysia. Sustainability, 11(21):6016, 2019.

[100] Jae-Hyuck Park, Dan Iancu, and Erica Plambeck. On the management of premade foods. Available at SSRN 4148756, 2022.

[101] Sanghyun Park and Phanish Puranam. Vicarious learning without knowledge differentials. Management Science, 70(5):2999–3019, 2024.

[102] Nicholas C Petruzzi and Maqbool Dada. Pricing and the newsvendor problem: A review with extensions. Operations Research, 47(2):183–194, 1999.

[103] Renata Soares Pinto, Renata Machado dos Santos Pinto, Felipe Fochat Silva Melo, Suzana Santos Campos, and Cláudia Marques-dos-Santos Cordovil. A simple awareness campaign to promote food waste reduction in a univ. canteen. Waste management, 76:28–38, 2018.

[104] Christopher Qian. How Do Buffets Stay Profitable?, 2022. Accessed: 2024-09-12.

[105] Xuelian Qin, Zhixue Liu, and Lin Tian. The strategic analysis of logistics service sharing in an e-commerce platform. Omega, 92:102153, 2020.

[106] Rushabh Rai. E-Commerce Fulfillment Services Market Size, Share Trends Analysis Report, 2025. Accessed: 2025-9-11.

[107] Reddit. This restaurant's food waste policy, 2019. Accessed: 2024-09-12.

[108] ReFED. Restaurant Food Waste Management Solutions, 2022. Accessed: 2024-09-12.

[109] Christian Reynolds, Liam Goucher, Tom Quested, Sarah Bromley, Sam Gillick, Victoria K Wells, David Evans, Lenny Koh, Annika Carlsson Kanyama, Cecilia Katzeff, et al. Consumption-stage food waste reduction interventions–what works and how to design better interventions. Food Policy, 83:7–27, 2019.

[110] Jennifer Robinson. The Top 5 Problems of Managing Inventory as an Amazon FBA Seller, 2023. Accessed: 2025-09-12.

[111] Barbara J Rolls, Liane S Roe, Kitti H Halverson, and Jennifer S Meengs. Using a smaller plate did not reduce energy intake at meals. Appetite, 49(3):652–660, 2007.

[112] ET Rolls. Understanding the mechanisms of food intake and obesity. Obesity reviews, 8:67–72, 2007.

[113] Jennifer K Ryan, Daewon Sun, and Xuying Zhao. Competition and coordination in online marketplaces. Production and Operations Management, 21(6):997–1014, 2012.

[114] Kim Severson. The buffet is back, stretching dollars and slinging Crab legs, 2023. Accessed: 2024-09-12.

[115] Zijun Shi, Kannan Srinivasan, and Kaifu Zhang. Design of platform reputation systems: Optimal information disclosure. Marketing Science, 42(3):500–520, 2023.

[116] Lucie Sirieix, Jan Lála, and Klára Kocmanová. Understanding the antecedents of consumers' attitudes towards doggy bags in restaurants: Concern about food waste, culture, norms and emotions. Journal of Retailing and Consumer Services, 34:153–158, 2017.

[117] Libo Sun, Guodong Lyu, Yugang Yu, and Chung-Piaw Teo. Fulfillment by amazon versus fulfillment by seller: An interpretable risk-adjusted fulfillment model. Naval Research Logistics, 67(8):627–645, 2020.

[118] Libo Sun, Guodong Lyu, Yugang Yu, and Chung Piaw Teo. Cross-border e-commerce data set: Choosing the right fulfillment option. Manufacturing & Service Operations Management, 23(5):1297–1313, 2021.

[119] Xurui Sun, Binghan Liu, Yuan Yuan, Ying Rong, Rui Pang, and Qiu Li. Neural and hormonal mechanisms of appetite regulation during eating. Frontiers in Nutrition, 12:1484827, 2025.

[120] Lin Tian, Asoo J Vakharia, Yinliang Tan, and Yifan Xu. Marketplace, reseller, or hybrid: Strategic analysis of an emerging e-commerce model. Production and Operations Management, 27(8):1595–1610, 2018.

[121] Treendly. Unveiling the Explosive Growth of Amazon FBA: A Trend that Reigns Supreme in E-Commerce, 2025. Accessed: 2025-09-12.

[122] Christina Troitino. The tragedy of the buffet commons: Reducing food waste at hotels, 2017. Accessed: 2024-09-12.

[123] Yushi Tsunoda and Yusuke Zennyo. Platform information transparency and effects on third-party suppliers and offline retailers. Production and Operations Management, 30(11):4219–4235, 2021.

[124] Marlin W Ulmer, Barrett W Thomas, Ann Melissa Campbell, and Nicholas Woyak. The restaurant meal delivery problem: Dynamic pickup and delivery with deadlines and random ready times. Transportation Science, 55(1):75–100, 2021.

[125] UnitedNations. The 17 Sustainable Development Goals (SDGs), 2015. Accessed: 2024-09-12.

[126] Walmart. WFS fees: Estimate fees for fulfillment, storage and more., 2025. Accessed: 2025-10-01.

[127] Brian Wansink and Mitsuru Shimizu. Eating behaviors and the number of buffet trips: an observational study at all-you-can-eat chinese restaurants. American Journal of Preventive Medicine, 44(4):e49–e50, 2013.

[128] Brian Wansink and Koert Van Ittersum. Portion size me: plate-size induced consumption norms and win-win solutions for reducing food intake and waste. Journal of Experimental Psychology: Applied, 19(4):320, 2013.

[129] Hanqi Wen, Izak Duenyas, and Joline Uichanco. Pooling goods of different quality: Platform design under inventory commingling. Available at SSRN 5345565, 2025.

[130] Kelly J Whitehair, Carol W Shanklin, and Laura A Brannon. Written messages improve edible food waste behaviors in a univ. dining facility. Journal of the Academy of Nutrition and Dietetics, 113(1):63–69, 2013.

[131] UN World Food Programme. World Food Programme: The Clear Plate Project. 2022. Accessed: 2024-09-12.

[132] Chi-Mei Emily Wu and Chih-Ching Teng. Reducing food waste in buffet restaurants: A corporate management approach. Foods, 12(1):162, 2022.

[133] Qi Wu and Dorothee Honhon. Don't waste that free lettuce! impact of bogof promotions on retail profit and food waste. Production and Operations Management, 32(2):501–523, 2023.

[134] Qi Xu, Zhong-miao Sun, and Xiao-qing Gao. Ordering and inventory reallocation decisions in a shared inventory platform with demand information sharing. Annals of Operations Research, 329(1):471–499, 2023.

[135] Xiaoyue Yan and Elena Belavina. Variability propagation in perish. product supply chains. Available at SSRN 4550255, 2023.

[136] Luyi Yang and Man Yu. Too good to go: Combating food waste with surprise clearance. Available at SSRN 4573386, 2023.

[137] Man Yang, Tao Zhang, and Chuan-xu Wang. The optimal e-commerce sales mode selection and information sharing strategy under demand uncertainty. Computers & Industrial Engineering, 162:107718, 2021.

[138] Yong Zha, Quan Li, Tingliang Huang, and Yugang Yu. Strategic information sharing of online platforms as resellers or marketplaces. Marketing Science, 42(4):659–678, 2023.

[139] Dennis J Zhang, Hengchen Dai, Lingxiu Dong, Qian Wu, Lifan Guo, and Xiaofei Liu. The value of pop-up stores on retailing platforms: Evidence from a field experiment with alibaba. Management Science, 65(11):5142–5151, 2019.

[140] Feng Zhu and Qihong Liu. Competing with complementors: An empirical look at Amazon. com. Strategic Management Journal, 39(10):2618–2642, 2018.