# EFFICIENT POLICIES AND MECHANISMS FOR ONLINE PLATFORMS 

by

Negin Golrezaei

A Dissertation Presented to the FACULTY OF THE GRADUATE SCHOOL UNIVERSITY OF SOUTHERN CALIFORNIA<br>In Partial Fulfillment of the Requirements for the Degree DOCTOR OF PHILOSOPHY (BUSINESS ADMINISTRATION)

To my loving husband, and my dear parents,
for their support, encouragement, and constant love.

## Acknowledgments

This thesis is a product of my close collaboration with my research advisor Prof. Hamid Nazerzadeh. Words cannot express my feelings of gratitude towards him for his constant support, encouragement, and guidance. He is the main reason I can now look back and say that I have had an incredible last five years at USC. I am also grateful to Prof. Paat Rusmevichientong and Prof. Ramandeep Randhawa, who are also my co-authors. They have always been a constant source of guidance and inspiration.

I am also grateful to Prof. Phebe Vayanos. for being on my thesis committee. I am greatly indebted to her for her encouragement and enthusiasm about my work. I also greatly appreciate the genuine concern, constant encouragement, and help of Professors Raj Rajagopalan and Greys Sosic. During the course of my stay at Marshall School of Business, I also had the good fortune of interacting with Professors Yehuda Bassok, Amy Ward, Vishal Gupta, Kimon Drakopoulos, Leon Zhu, Song-Hee Kim, and Ashok Srinivasan they have all been highly inspirational in addition to being instrumental in making me feel part of the greater USC family.

I would like to thank the collegial faculty members and friendly staff at the Marshall School of Business, who were always available when I needed help in research or life. Special thanks to my Ph.D. friends for their sympathetic ear, altruistic help, and stimulating discussions we had about our research and lives.

Finally, I want to thank my family for their constant support all through my life - no achievement in my life would have been possible without them. I am forever indebted to my husband, Sajjad Beygi, and my parents, Mohammad Hossein Golrezaei and Batool Najafi. Their love, encouragement, and support have been the core of my strength throughout the quest of attaining my Ph.D. degree. This thesis is dedicated to them. I also express a feeling of bliss for having loving siblings like Zahra, Saeed, and Mahnaz. They have always been genuinely concerned and ever encouraging.

## Contents

Acknowledgments ..... iii
List of Tables ..... viii
List of Figures ..... ix
Abstract ..... xii
1 Real-time Optimization of Personalized Assortments ..... 1
1.1 Introduction ..... 1
1.1.1 Literature Review ..... 5
1.2 Preliminaries and Problem Formulation ..... 8
1.3 Inventory-Balancing Algorithms ..... 10
1.3.1 The Tight Upper Bound on the Competitive Ratio ..... 14
1.4 Stochastic I.I.D. Arrivals ..... 16
1.4.1 Motivating IB Algorithms via Dual-Based Heuristics ..... 18
1.5 Extensions ..... 19
1.5.1 Incorporating Partial Information and Learning Customer types? ..... 19
1.5.2 Incorporating (Uncertain) Information About Arrival Pattens ..... 20
1.5.3 Beyond Substitutability ..... 22
1.6 Numerical Experiments ..... 23
1.6.1 Dataset and Simulation Setting ..... 24
1.6.2 Performance Evaluation ..... 25
1.6.3 Benefits of Personalization: The Importance of Knowing Customer Types ..... 30
1.6.4 Comparison to Myopic Policy ..... 31
1.7 Conclusion and Future Work ..... 32
2 Auctions with Dynamic Costly Information Acquisition ..... 34
2.1 Introduction ..... 34
2.2 Setting ..... 40
2.3 Direct Mechanisms ..... 41
2.4 The Efficient Mechanism ..... 44
2.4.1 What If All Agents Are Allowed to Access the Additional Information? ..... 47
2.5 Maximizing Revenue ..... 50
2.5.1 Sequential Second-Price Mechanisms ..... 50
2.5.2 Revenue-Optimal Mechanism ..... 52
2.5.3 What If All Agents Are Allowed to Access to the Additional Information? ..... 54
2.6 Who Will Be Selected? ..... 55
2.6.1 Impacts of the Cost and Variance of Second Signals on Selection Rules ..... 58
2.7 Extensions ..... 59
2.7.1 Information Acquisition as an Entry Cost ..... 59
2.7.2 Multi-unit ..... 61
2.8 Conclusion ..... 63
3 Dynamic Pricing for Heterogeneous Time-Sensitive Customers ..... 64
3.1 Model ..... 69
3.2 Direct Mechanisms and Optimality ..... 70
3.3 Optimal Mechanism with Exponential Valuation Functions ..... 72
3.4 Optimal Mechanism with Production and Holding Costs ..... 77
3.4.1 Positive Production Cost ..... 79
3.4.2 Positive Holding Cost ..... 81
3.4.3 Positive Production and Holding Costs ..... 89
3.5 General Valuation Functions ..... 90
3.6 Conclusion ..... 92
A Technical Appendix to Chapter 1 ..... 107
A. 1 ..... 107
A.1.1 Proof of Lemma 1.2.1: ..... 107
A.1.2 Proof of Theorem 1.3.3 ..... 108
A.1.3 Proof of Proposition 1.4.2 ..... 111
A.1.4 Learning the Customer Types under the Multinomial Logit Model ..... 113
A. 2 Computational Complexity of Inventory-Balancing Algorithms ..... 113
A. 3 Numerical Experiments: Appendix to Section 1.6 ..... 115
A.3.1 Known Length of the Horizon ..... 115
A.3.2 Worst-Case Performance ..... 116
A.3.3 Learning the Customer Types ..... 117
A. 4 Relegated Proofs ..... 119
A.4.1 Appendix to Section 1.3.1 ..... 120
A.4.2 Proof of Proposition A.1.2 ..... 123
A.4.3 Appendix to Section 1.5.2 ..... 126
A. 5 Asymptotic Optimality of the Dynamic Programming Policy ..... 128
B Technical Appendix to Chapter 2 ..... 132
B. 1 Sequential Weighted Second-Price Mechanism \& Proofs of Theorems 2.4.1 and 2.5.1 ..... 132
B.1.1 Proof of Theorem B.1.4 ..... 134
B.1.2 Proof of Lemma B.1.9 ..... 138
B. 2 Proof of Theorem 2.6.1 ..... 140
B. 3 Appendix to Sections 2.4.1 and 2.5.3 ..... 141
B.3.1 Proof of Theorem 2.4.2 ..... 141
B.3.2 The All-Access Mechanism and Single Crossing Conditions ..... 143
B.3.3 Proof of Theorems 2.4.4 and 2.5.5 ..... 146
B. 4 Delegated Proofs ..... 150
B.4.1 Proof from Section 2.5 ..... 150
B.4.2 Proofs from Section B. 1 ..... 152
B.4.3 Proofs from Section B.1.1 ..... 154
B.4.4 Proofs from Section B.1.2 ..... 158
B. 5 Numerical Experiments ..... 159
B.5.1 Payments in the First Round ..... 159
B.5.2 The SSP Mechanism versus the optimal Mechanism ..... 160
B.5.3 More Agents ..... 161
C Technical Appendix to Chapter 3 ..... 163
C. 1 Proof of Main Results ..... 163
C.1.1 Appendix to Section 3.2 ..... 163
C.1.2 Lower Bound on the Revenue Gain of the Dynamic Pricing Policy ..... 165
C.1.3 Proof of Theorem 3.4.4 ..... 166
C.1.4 Optimal Mechanism for a Low Holding Cost ..... 167
C.1.5 Optimal Mechanism for a Medium Holding Cost ..... 174
C.1.6 Optimal Mechanism for a High Holding Cost ..... 175
C.1.7 Proof of Theorems 3.4.1 an 3.4.5 ..... 177
C.1.8 Discussing the Assumption in Theorem 3.4.4 ..... 178
C.1.9 Proof of Theorem 3.5.1 ..... 181
C. 2 Technical Proofs ..... 186
C.2.1 Proof of Lemma C.1.6 ..... 186
C.2.2 Proofs of Lemmas in Sections C.1.4 and C.1.5 ..... 189
C.2.3 Proof of Lemmas in Section C.1.9 ..... 194

## List of Tables

1.1 Top 10 DVDs nationally and the percentage of customers in each location who purchase each DVD. ..... 2
1.2 The locations used in the simulations and the top three best-selling DVDs in each location. ..... 24
1.3 Revenue comparison when the length of the horizon is random. The standard errors of all numbers are less than $0.1 \%$. ..... 28
1.4 The average revenue for IB algorithms in multiple-type model and the improvement over the single-type model. All numbers are statistically significant and the standard errors are less than $0.1 \%$. ..... 30
1.5 Revenue Comparison. The standard errors of all numbers are less than $0.1 \%$. ..... 32
A. 1 Revenue comparison when the length of horizon is known in advance. The standard errors of all numbers are less than $0.1 \%$. ..... 115
A. 2 Worst-Case Performance Comparison when the length of horizon is unknown. ..... 117
A. 3 The average revenue for the LIB and EIB algorithms when the underlying parameters are unknown, and each algorithm uses the estimated parameters based on data collected in the previous periods ..... 118
B. 1 Revenue of the SSP mechanism (with revenue-maximizing $r$ ) as a percentage of the optimal revenue with $F=N(0.5,0.5)$. Here, the standard errors of all numbers are less than $1 \%$. ..... 161

## List of Figures

1.1 The C.C.D.F. of revenues with $\mathrm{LF}=1.8$ and $\mathrm{CV}=2$. For each algorithm, the curve shows the fractions of problem instances (out of 250 ) whose revenue exceed certain percentages of the upper bound. ..... 29
2.1 The average number of agents with the updated valuations versus the cost in the efficient and All-Access (with no reserve) mechanisms with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=$ $\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$. ..... 50
2.2 The average social welfare versus the cost in the efficient and All-Access (with no reserve) mechanisms with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$. ..... 51
2.3 Depicting the results in Theorem 2.5.5 with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$. ..... 56
2.4 The revenue of the revenue-optimal and All-Access (with revenue-maximizing reserve price) mechanisms versus the cost with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$. ..... 57
2.5 Selected agents in the optimal mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2, c_{1}=0.01, c_{2}=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$ for $i=1,2$. ..... 58
2.6 Selected agents in the efficient mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2, c_{1}=0.01, c_{2}=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$ for $i=1,2$. ..... 59
2.7 Average number of selected agents in the revenue-optimal and efficient mechanisms versus the standard deviation of second signals, $\sigma$, with $G_{i}=N\left(0, \sigma^{2}\right)$ and $c_{i}=0.05$ for $i=1,2$. . ..... 60
2.8 Average number of selected agents in the revenue-optimal and efficient mechanisms versus the cost of information, $c$, with $G_{i}=N(0,0.5)$. ..... 61
2.9 Average revenue of the OPT and OPT-E mechanisms versus the cost with $G_{i}=N(0,0.5)$, and $c_{i}=c$ for $i=1,2$. In all the figures, $n=2$ and $F=N(0.5,0.5)$. ..... 62
3.1 At time 0 , customer 1 has a higher value than customer 2 . But, her value decreases faster than customer 2 , and beyond $t_{0}$, customer 2 has a higher value.
3.2 Time of purchase versus type in the fixed and dynamic pricing policies with the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and $h=c=0$.
3.3 Payment versus type in the fixed and dynamic pricing policies with the customer type $\theta \sim$ $U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and $h=c=0$.
3.4 Utility versus type in the fixed and dynamic pricing policies with the customer type $\theta \sim$ $U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and $h=c=0$.
3.5 The cut-off $\theta_{c}$ versus the production cost $c$ for the mechanism described in Theorem 3.4.1 and the FP policy. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and the holding cost $h=0$.
3.6 The social welfare and revenue gain of DP (relative to FP ) as the production cost varies. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and the holding cost $h=0$.
3.7 The structure of the optimal mechanism as a function of the holding cost $h$. Here, the customer type $\theta \sim U(0,1), \beta=0.1, V(\theta, t)=\theta e^{-\beta \theta t}$, and the production cost $c=0$.
3.8 The social welfare and the revenue gain of the optimal mechanism (relative to the FP policy) as a percentage versus the holding cost. Here, the customer type $\theta \sim U(0,1), \beta=0.1$, $V(\theta, t)=\theta e^{-\beta \theta t}$, and the production $\operatorname{cost} c=0$.
3.9 Time of purchase versus type. Here, $H_{l}=0.004, H_{h}=0.025$, the customer type $\theta \sim$ $U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and production cost $c=0$.
3.10 Payment versus type. Here, $H_{l}=0.004, H_{h}=0.025$, the customer type $\theta \sim U(0,1)$, $V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and production cost $c=0$.
3.11 Utility versus type. Here, $H_{l}=0.004, H_{h}=0.025$, the customer type $\theta \sim U(0,1)$, $V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and production cost $c=0$.
3.12 The thresholds versus the exponent $a$ for the mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.93
3.13 The social welfare and revenue gain of DP (relative to the FP policy) in percentage as a function of the exponent $a$. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}$, $g(\theta)=\theta^{a}$, and $h=c=0$.93
3.14 Time of purchase versus type in the optimal mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.
3.15 Payment versus type in the optimal mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.
3.16 Utility versus type in the optimal mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.
A. 1 The cumulative revenue over time for $L F=1.8$ and $C V=2$ when the length of the horizon is known in advance.
B. 1 The payment of agent 1 in the first round, $t_{1}$, in the optimal mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2, c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.
B. 2 The payment of agent 1 in the first round, $t_{1}$, in the efficient mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2, c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.
B. 3 Revenue in the optimal and efficient mechanisms versus number of agents, $n$, with $c=0.05$, $F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.
B. 4 Social welfare in the optimal and efficient mechanisms versus number of agents, $n$, with $c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.

162
B. 5 The average number of selected agents of the optimal and efficient mechanisms versus number of agents, $n$, with $c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.

## Abstract

Markets have been studied for centuries. But, because of the rise of online markets in the past decade, our ability to experiment with markets and get immediate feedback has radically changed. Such opportunities exist because online platforms can collect unprecedented amounts of data. That data can be incorporated into designing optimal marketplaces, and contribute to making informed operational decisions. For instance, data enables online markets to provide personalized services by exploiting heterogeneity among their customers. My research interests lie in understanding the economic and operational properties of online markets and contribute to their design and advancement.

In this dissertation, I focus on designing and developing fast, robust algorithms and mechanisms for two types of online markets that have been significantly influenced by the availability of massive amounts of data: online retailing markets and online advertising markets.

In Chapter 1, I study the problem of optimizing personalized assortment planning in online retail stores. The online retail stores provide a platform that enables consumers to directly buy goods or services from a seller over the Internet. These platforms allow online retailers to store data from customers to get information about their behavior and needs. This gives online retailers an opportunity to use such information to personalize product-offering (aka assortment planning) based on customers' preference, interests, and other characteristics. By personalizing assortment planning, online retailers can gain substantial potential revenue improvements. In this chapter, we propose a family of simple and effective algorithms, called Inventory-Balancing, for real-time personalized assortment optimization, that does not require any forecasting. Our proposed algorithm has a strong performance guarantee. We show that our algorithms obtain at least $1-1 / e=63 \%$ of the benchmark revenue, even when there are sudden shocks in the customers' arrival patterns, either from seasonality or other non-stationarity effects. The results in this chapter are published
as the journal paper by Golrezaei et al. (2014). Furthermore, this paper got nominated for the MSOM Student Paper Competition in 2016 and the Production and Operations Management Society College of Supply Chain Management 2017 Student Paper Competition.

In Chapter 2, we study the impacts of data in two-sided online markets where two sets of independent self-interested agents interact with each other, and the decisions of each set of agents affect the outcomes of the other set of agents. In two-sided platforms such as online advertising markets, data impacts both sets of agents in possibly different ways. Now the challenge is to design selling mechanisms that effectively disclose and/or use data, and incentivize self-interested agents to act in a globally optimal manner. In this chapter, we study the mechanism design problem for a seller (web publisher) of an indivisible good in a setting where privately informed buyers (advertisers) can acquire additional information and refine their valuations for the good at a cost. For this setting, we propose optimal (revenue-maximizing) and efficient (welfare-maximizing) mechanisms that induce a right level of investment in information acquisition. We show that because information is costly, in the optimal and even the efficient mechanisms, not all the buyers would obtain the additional information. In fact, these mechanisms incentivize buyers with higher initial valuations to acquire information. The results in this chapter are published as the journal paper by Golrezaei \& Nazerzadeh (2016).

In Chapter 3, we study the problem of pricing in online retailers in the presence of strategic customers. In online retailing, time-based pricing (dynamic pricing), which is is a pricing strategy in which firms set flexible prices for products based on current market demands, has been becoming increasingly prevalent. One of the main advantages of dynamic pricing is that it helps mitigate the risk associated with demand uncertainty (see, for instance, Aviv \& Pazgal 2008 and Cachon \& Swinney 2011). In this chapter, we show that dynamic pricing can play an important role in differentiating between customers over time even in the absence of demand uncertainty. In many settings, especially in fashion and electronic gadget retail, a customer's willingness-to-pay (or valuation) for the product is time-sensitive and decreases over time. In these situations, customers are not only different in terms of their initial willingness-to-pay for these products when they are first introduced to the market, but they are also different in terms of how rapidly they lose their interest in these products. We characterize the optimal mechanism for selling durable products in such environments and show that delayed allocation and dynamic pricing can be an effective screening tool for maximizing profit of a firm. This chapter is based on a joint work with Professors Hamid Nazerzadeh and Ramandeep Randhawa.

## Chapter 1

## Real-time Optimization of Personalized Assortments

### 1.1 Introduction

The availability of real-time data on customer characteristics has encouraged companies to personalize operational decisions to each arriving customer. For instance, the product recommendations Amazon.com makes to each customer dynamically change depending on recent reviews, ratings, purchases of the customer herself, other customers with similar interests to hers, and several other factors (Amazon's Recommendation Systems 2012). Orbitz.com, as another example, has found that users of Apple Macintosh computers spend as much as $30 \%$ more per night on hotels; consequently, the company can show Mac users different and more expensive assortments of hotels and travel options, than Windows users (Mattioli 2012). Locationbased deals and coupons are offered by Groupon, Yelp, Foursquare and other Internet companies (Wortham 2012). In online advertising, the advertisements that are displayed to a user browsing a website are routinely personalized based on the user's browsing history, demographic information, and her location (Helft \& Vega 2010). Even brick-and-mortar grocery stores are starting to offer personalized real-time coupons based on each customer's purchasing history and the available products on the shelf in the aisle where each customer is currently shopping (Clifford 2012).

These examples raise a key question that motivates our work: Given the complexity of coordinating the real-time, front-end, customer-facing decision with the back-end supply chain constraints, what policies should companies use to take advantage of such data?

We answer these questions by formulating a real-time, personalized, choice-based assortment optimization problem, involving multiple products with limited inventories, and arbitrary customer types. The type of each arriving customer can be arbitrary, and it is indexed by a (possibly infinite) set $\mathcal{Z}$. Examples of types include the customer's computer (Mac vs PC), her current location, purchasing history, the average
household income in her neighborhood, the competitors' current offerings and prices, time of day, or a combination of other observable characteristics.

For an arriving customer of type $z \in \mathcal{Z}$, the company must decide, in real-time, on the assortment of products to offer. Given an assortment $S$, the customers make choices on which products to buy, if any, according to a general choice model that is specific to each customer type. Our goal is to develop a revenuemaximizing policy that determines the assortment to offer to each arriving customer, taking into account the customer type and the current inventories.

The above formulation captures the essential features of the situation faced by companies that sell services, products, or advertisement to heterogenous customer types that requires real-time decision-making, with inventory constraints.

We first observe that differentiating customer types (even just their locations) can significantly increase revenues. We consider the top ten DVDs with the highest national sales volumes during the summer of 2005, and compare their sales in two locations: Urbana-Champaign, IL and Miami, FL. Table 1.1 shows the sales rate of each DVD in each location, which is defined as the proportion of the potential customers in each location ${ }^{1}$ who purchased each DVD.

$\left.$| DVD | Title | Sales Rate |  |  |
| :---: | :--- | :---: | :---: | :---: |
|  |  |  | National | Miami | | Urbana- |
| :---: |
| Champaign | \right\rvert\,

Table 1.1: Top 10 DVDs nationally and the percentage of customers in each location who purchase each DVD.

We observe that these two locations exhibit very different purchasing behaviors, which are significantly different from the national sales pattern. Consider the sci-fi DVD "Firefly - The Complete Series", which

[^0]is the second most popular DVD nationally, with a sales rate of $6.9 \%$. None of the customers in Miami purchased this DVD, while $12 \%$ of the customers in Urbana-Champaign, almost twice the national rate, bought the series. On the other hand, none of the customers from Urbana-Champaign bought "What the Bleep Do We Know!?", while almost $10 \%$ of the customers in Miami bought it. In Section 1.6.3, we evaluate the performance of our algorithms and observe that customizing the assortments of DVDs to each customer's location leads to over $10 \%$ increase in revenues. The improvement can be as high as $21 \%$ when the size of the assortments are constrained.

As our main contribution, we propose a family of simple and effective algorithms, called InventoryBalancing, for real-time personalized assortment optimization, that do not require any forecasting: An Inventory-Balancing algorithm, maintains a "discounted-revenue index" for each product, in which the (actual) revenue is multiplied by a (virtual) discount factor that depends on the fraction of the product's remaining inventory. Upon an arrival of each customer, based on the customer's type, the algorithm offers to her the assortment that maximizes the expected discounted revenue. Each Inventory-Balancing algorithm is characterized by the penalty function that discounts the marginal revenue of each product, as the inventory level reduces. By adjusting the revenue of each product according to its remaining inventory, the algorithms hedges against the uncertainty in the types of future customers, by reducing the rate at which products with low inventory are offered. Thus, the discounted-revenue index serves as a simple mechanism that coordinates the front-end customer-facing decision with the back-end supply chain constraints. Our Inventory-Balancing algorithms offer the following benefits.

- No forecasting: A traditional approach for dynamic assortment optimization, both in the literature and in practice, is to forecast the demand over time by estimating the distribution of the number of customers of each type, and then find an optimal policy based on the forecast using re-optimization methods (Gallego \& van Ryzin 1994a, Jasin \& Kumar 2012) or dynamic programming (Bernstein et al. 2011). In our sales data, we observe a large variability in the number of customers, across time and locations. Such a large variability in the demand process often makes forecasting difficult, and not surprisingly, could lead to poor performance for the policies obtained under this approach.

An alternative approach for making a real-time decision is to solve the "off-line" assortment optimization problem repeatedly, using the most up-to-date inventory levels of each product and the latest demand forecast. This can be done by repeatedly solving a series of linear programs; see, for example, Jasin \& Kumar (2012). When the number of customers is known in advance, the re-optimization
methods work extremely well and yield nearly optimal revenue because they can effectively ration the inventory to all customers. However, when there is a significant uncertainty in the market size, the problem becomes more challenging. In this setting, our Inventory-Balancing algorithms perform very well, yielding $5 \%-11 \%$ more revenues than re-optimization methods.

- Strong performance under both non-stationary and stationary demand processes: As a performance benchmark, we compare the revenue of our algorithms to the revenue of a clairvoyant optimal solution that has complete knowledge of the sequence of the types of the customers that arrive in the future, but does not know the (random) choice for each future customer ${ }^{2}$. We prove that Inventory-Balancing algorithms with a strictly concave penalty function always obtain more than $50 \%$ of the optimal revenue; see Theorem 1.3.3 and Corollary 1.3.4. We also provide an Inventory-Balancing algorithm that obtains at least $\left(1-\frac{1}{e}\right) \approx 63 \%$ of the benchmark revenue. This implies that even when there are sudden shocks in the customers' arrival patterns, either from seasonality or other non-stationarity effects, the algorithm maintains a strong performance guarantee.

The $\left(1-\frac{1}{e}\right)$ fraction of the benchmark revenue is optimal for non-stationary stochastic arrivals in the sense that in the worst-case, no deterministic or randomized policy can achieve a higher competitive ratio; see Theorem 1.3.6.

When customer arrivals are stationary, our algorithms perform even better. We show that when the types of arriving customers are independently and identically distributed, our algorithm is guaranteed to obtain at least $75 \%$ of the benchmark; see Proposition 1.4.2. In our numerical experiments, presented in Section 1.6, our algorithms perform even better than what is predicted by the worst-case bound, obtaining revenues that are within $96 \%-99 \%$ of the benchmark.

- Simplicity, robustness, and flexibility: In contrast to the existing methods, our Inventory-Balancing algorithms are extremely simple and fast. We do not need to solve any offline assortment optimization problem, and consequently, we can compute the decision for each customer quickly. Our formulation allows for infinite customer types, and thus, our algorithms are robust to changing customer types over time. In addition, under mild assumptions, our analysis and performance guarantees continue to hold

[^1]when the choice models of the customers are learned over time and the algorithm uses estimations of the parameters of the choice models; see Proposition 1.5.1.

Our proposed algorithms are also flexible, and they can be easily combined with existing reoptimization methods while maintaining a worst-case performance guarantees; see Section 1.5.1 and Proposition 1.5.2. Our numerical experiments show that such a hybrid method brings out the advantages of all methods, especially when there is uncertainty in the number of future customers.

The key message in this chapter is that real-time optimization of personalized assortments can be done efficiently and robustly. Our proposed policies maintain a simple index for each product, which balances the nominal revenue with the value of each unit of remaining inventory. These indices are easy to implement, and they serve as a simple mechanism that coordinates between the front-end real-time decision and the back-end supply chain constraints. As the volumes of data on customer profiles and preferences continue to grow, we believe that companies will consider personalization of other operational decisions, such as pricing or shipping options. The framework and analysis in this chapter can serve as a starting point for more complex models.

### 1.1.1 Literature Review

Our work is related to the growing literature on assortment planning. We describe a brief overview of the area, to provide a context for our work. Assortment planning problems focus on the relationships among assortment offerings, customer choices, and inventory constraints. van Ryzin \& Mahajan (1999) introduced one of the first models that capture the tradeoffs between inventory costs and product variety. Mahajan \& van Ryzin (2001) followed up on this work with a study on the optimal inventory levels in the presence of stockouts and substitution behavior. Since their seminal work, researchers have considered a variety of choice models and studied how such models affect the optimal assortment and the inventory level of products we should carry. Examples include the demand substitution model (Smith \& Agrawal 2000), Lancaster choice model (Gaur \& Honhon 2006), ranked-list preferences (Honhon et al. 2010, Goyal et al. 2011), and multinomial logit models (Talluri \& van Ryzin 2004b, Gallego et al. 2004, Liu \& van Ryzin 2008a, Topaloglu 2013). Recently, Farias et al. (2013) introduced a very general class of choice models based on a distribution over permutations, and developed efficient algorithms for determining the optimal assortment. For a survey of the assortment planning literature, the reader is referred to Kök et al. (2008).

Two of the most important decisions in modeling an assortment planning problem are determining the customers' choice models and capturing the arrival process of the customers. The family of choice models considered in our work is quite general and includes most of the choice models used in practice or previously studied by researchers. What distinguishes this work from the existing literature is the fact that we do not impose any restrictions on the arrival process and most of our results hold even if an adversary chooses the sequence of customers. In the following, we briefly discuss the prevalent approaches to model the arrival process.

A common approach to model customer arrivals is to assume that arrivals follow a stochastic process. In this model, the optimal sequence of assortments can be planned by solving a multi-dimensional dynamic program. Not surprisingly, this approach suffers from the curse of dimensionality, even for stationary processes.

Recently, Bernstein et al. (2011) studied the aforementioned assortment planning model under the assumption that the type of customers (represented by multinomial logit choice models) is drawn identically and independently from a stationary distribution; i.e., I.I.D. arrivals. For two products with equal revenue, two customer types (with each type following a multinomial logit choice model), and Poisson arrivals, they provide structural properties of the optimal solution. Interestingly, they show that the optimal dynamic program may withhold products with low remaining inventory for future customers that are more interested in them. Based on this observation, they propose a heuristic that, roughly speaking, reduces the general problem with multiple products to a two-product problem by separating the products into two groups based on their inventory to demand ratio. They do not provide any performance guarantees for the heuristic.

In practice, we do not expect the distribution of customer types to remain constant over time because of seasonality effects or changing popular trends. In the context of airline revenue management, the fraction of business customers tend to increase as the departure date approaches. Prior to our work, to the extent of our knowledge, the best known performance guarantee for a heuristic with respect to a clairvoyant optimal solution in non-stationary stochastic environments was a ratio of $\frac{1}{2}$ that follows from Chan \& Farias (2009). When the arrivals are stochastic, with some adjustments, the assortment planning studied in the chapter would fit into the stochastic depletion framework proposed by Chan \& Farias (2009). They show that in their framework the competitive ratio of a myopic policy is at least $\frac{1}{2}$.

Re-optimization policies (Jasin \& Kumar 2012) are applicable to non-stationary stochastic environments. However, we are not aware of any results that provide a performance guarantee. ${ }^{3}$ The closest work to this line of research is by Ciocan \& Farias (2013), who studied re-optimization policies for a network revenue management problem where the distribution of the valuation of the customer (i.e., the distribution of the types) is constant over time, however the size of the market changes over time according to a stochastic (e.g., multi-variate Gaussian) process. They showed that a re-optimization policy that adjusts prices of the products by solving a linear program obtains about one-third of the optimal revenue; see also Chen \& Farias (2013). In contrast to dynamic pricing problems, where firms manage their profits and capacities by controlling the prices, in assortment planning models, product prices are exogenously determined and remains constant over the horizon, and the firms decide on the selection of the assortment to offer to each customer.

We choose the competitive ratio as our performance benchmark because it allows for arbitrary nonstationary, even adversarial, arrivals and it does not require any prior knowledge about the arrival patterns. This notion has been previously applied by Ball \& Queyranne (2009) to the problem of capacity allocation. Besbes \& Zeevi (2011) and Besbes \& Sauré (2012) also used similar notions of optimality when they studied revenue management problems where the demand may change dramatically because of shocks.

The problem we study here resembles some of the aspects of the Adwords problem (Mehta et al. 2007, Buchbinder et al. 2007, Goel, Mahdian, Nazerzadeh \& Saberi 2010, Azar et al. 2009), where the goal is to allocate a sequence of advertisement spaces associated to search queries to budget-constrained advertisers. Both the Adwords and personalized assortment optimization problems contain the b-matching problem as a special case (Kalyanasundaram \& Pruhs 2000). Mehta et al. (2007) proposed an algorithm that achieves an optimal-competitive ratio for the Adwords problem by taking into account both the bid and budget of the advertisers; see Buchbinder \& Naor (2007) for a survey on online algorithms and also Acimovic \& Graves (2011) for another application in the context of inventory management.

Organization: In Section 1.2, we formally define our problem. Our algorithm and main results are presented in Section 1.3. We discuss the performance of our algorithm under stationary stochastic arrivals in Section 1.4, followed by discussions of extensions of our original model in Section 1.5. We present

[^2]the numerical experiments in Section 1.6. The conclusion and direction for future work are given in Section 1.7. Finally, the proof of the competitive ratio and discussions of computational complexity are in Appendix B.1.1.

### 1.2 Preliminaries and Problem Formulation

Consider a firm that sells $n$ products, indexed by $1,2, \ldots, n$, to customers that arrive sequentially over time. The firm obtains a revenue $r_{i}>0$ for selling each unit of product $i$, which has an initial inventory of $c_{i} \in \mathbb{Z}_{+}$, with no replenishment. We denote the no-purchase option as product 0 , with $r_{0}=0$. Let $\mathcal{Z}$ denote the set of possible customer types. Once a customer arrives, her type, denoted by $z \in \mathcal{Z}$, is revealed. For instance, the type of a customer can correspond to his or her web browser, i.e., $\mathcal{Z}=\{\mathrm{Mac}, \mathrm{PC}\}^{4}$ As mentioned in the introduction about assortment personalization by Orbitz.com, the type $z=$ Mac may suggest that the user is more likely to choose expensive travel options. If we are interested in the location of each customer, then the type of the customer can correspond to his or her zip-code. ${ }^{5}$ In addition, the revelation of each customer's type can happen when the customer logs in to the websites, e.g., Amazon or eBay.

Based on the customer's type and the remaining inventory, the firm offers an assortment $S \in \mathcal{S}$, where $\mathcal{S}$ denotes the set of all feasible assortments; we assume $\{0\} \in \mathcal{S}$, i.e., the firm has the option to not offer any product. The set $\mathcal{S}$ allows us to incorporate a variety of constraints on the assortments, such as the shelf-space or size constraints.

Associated with each customer type $z \in \mathcal{Z}$ is the probability of purchasing each product under each assortment. More specifically, each customer type $z \in \mathcal{Z}$ corresponds to a general choice model which specifies the probability of purchasing each product under each assortment. We denote by $\phi_{i}^{z}(S)$ the probability that a customer of type $z$ purchases product $i$, when assortment $S$ is offered. In fact, all of our results continue to hold when each customer may purchase more than one product at a time: define $\Phi^{z}: \mathcal{S} \times \mathcal{S} \rightarrow[0,1]$, where $\Phi^{z}\left(S^{\prime}, S\right)$ is the probability that a customer of type $z$ purchases exactly the products in set $S^{\prime}$, when assortment $S$ is offered; in addition, $\Phi^{z}\left(S^{\prime}, S\right)=0$ when $S^{\prime} \nsubseteq S$ or $S \notin \mathcal{S}$ and $\sum_{S^{\prime} \subseteq S} \Phi^{z}\left(S^{\prime}, S\right)=1$.

[^3]Hence, we have $\phi_{i}^{z}(S)=\sum_{S^{\prime}: i \in S^{\prime}, S^{\prime} \subseteq S} \Phi^{z}\left(S^{\prime}, S\right)$. In the remainder of the chapter, we only use the notation $\phi_{i}^{z}(\cdot)$.

Our goal is to design an algorithm that offers an assortment to each arriving customer in order to maximize the total expected revenue. Let a vector $\left\{z_{t}\right\}_{t=1}^{T}=\left(z_{1}, z_{2}, \cdots, z_{T}\right)$ represents the sequence of types of the arriving customers, where for each $t, z_{t} \in \mathcal{Z}$ denotes the type of the customer that arrives in period $t$.

Definition 1. For any algorithm $A$ and any sequence of customer types $\left\{z_{t}\right\}_{t=1}^{T}$, we denote by $\operatorname{Rev}_{A}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ the expected revenue obtained by algorithm $A$ from the customers $\left\{z_{t}\right\}_{t=1}^{T}$, where the expectation is taken with respect to the choices made by each customer, and possibly random selections of the algorithm (if the algorithm is not deterministic).

We do not assume any arrival patterns and the algorithm does not know the sequence of the customers in advance. Therefore, we use the notion of competitive ratio, defined below, to measure the performance of an algorithm. The following lemma establishes an upper bound on the expected revenue that can be obtained by any algorithm from a sequence of customers.

Lemma 1.2.1 (Revenue Upper Bound). For any sequence of customers $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ and any algorithm $A$, $\operatorname{Rev}_{A}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ is bounded by the optimal value of the linear program Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ defined below:
maximize $\quad \sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \sum_{i=1}^{n} r_{i} \phi_{i}^{z_{t}}(S) y^{t}(S)$

SUBJECT TO: $\quad \begin{array}{rll}\sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \phi_{i}^{z_{t}}(S) y^{t}(S) & \leq c_{i} & 1 \leq i \leq n \\ \sum_{S \in \mathcal{S}} y^{t}(S) & =1 \quad 1 \leq t \leq T \\ y^{t}(S) & \geq 0 \quad 1 \leq t \leq T, \forall S \in \mathcal{S} \\ & & \\ & & \\ & & \end{array}$

In the linear program above, $y^{t}(S)$ corresponds to the probability that the set $S$ is offered to the customer of type $z_{t}$ in period $t$. With a slight abuse of notation, we denote the optimal value of the linear program above by Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ as well. The proof, given in Appendix A.1.1, follows from the fact that Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ is an upper bound on the expected revenue of the optimal clairvoyant solution that knows the sequence of the customer types in advance. Namely, we construct a feasible solution for the linear program above based on the optimal clairvoyant solution, taking into account the realizations of the customers' choice models.

By the above lemma, no algorithm without hindsight would obtain revenue equals to Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$. However, an algorithm with no knowledge of the future types might be able to obtain a fraction of the revenue of this clairvoyant optimal solution. Therefore, the competitive ratio of an algorithm is defined as follows:

Definition 2 (Competitive Ratio). An algorithm $A$ is $\alpha$-competitive if:

$$
\inf _{T \geq 1} \inf _{\left\{z_{t}\right\}_{t=1}^{T}: z_{t} \in \mathcal{Z} \forall t} \frac{\operatorname{Rev}_{A}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)}{\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)} \geq \alpha
$$

The infimum is taken over all possible sequences of customer arrivals of arbitrary lengths. In other words, the competitive ratio is defined as the worst-case ratio between the "expected revenues" of an algorithm and the optimal clairvoyant solution over a (possibly infinite) sequence of customer types, where the expectation is with respect to the realization of customers' choice models. ${ }^{6}$

One potential criticism of the notion of competitive ratio could be that it compares algorithms with a benchmark that is too strong. However, as we show in the following sections, in the context of assortment planning, it leads to simple algorithms that perform very well with respect to this benchmark. Moreover, our numerical simulations demonstrate the practical relevance of our method, and show that our algorithms outperform existing methods in the literature.

### 1.3 Inventory-Balancing Algorithms

We present a family of algorithms called Inventory-Balancing (IB), which take into account both the revenue that would be obtained from the customer and the current inventory levels, in order to decide which assortments to offer. Each Inventory-Balancing algorithm is defined with a penalty function $\Psi:[0,1] \rightarrow[0,1]$, which is an increasing function with $\Psi(0)=0$ and $\Psi(1)=1$.

Recall that $c_{i}$ is the initial inventory of product $i$. Let $I_{i}^{t}$ denote the remaining inventory of product $i$ at the end of period $t$. Note that $I_{i}^{0}=c_{i}$, and for $t \geq 1, I_{i}^{t}=\max \left\{I_{i}^{t-1}-Q_{i}^{t}, 0\right\}$, where $Q_{i}^{t}$ is a binary random

[^4]variable that is equal to 1 if the customer has chosen product $i$, and is 0 otherwise. We are now ready to describe the algorithm.

## Inventory-Balancing with a Penalty Function $\Psi$

Upon the arrival of the customer in period $t \in\{1, \ldots, T\}$, of type $z_{t}$, offer an assortment $S^{t}$ :

$$
S^{t}=\arg \max _{S \in \mathcal{S}} \sum_{i \in S} \Psi\left(I_{i}^{t-1} / c_{i}\right) r_{i} \phi_{i}^{z t}(S)
$$

The assortment $S^{t}$ can be found in polynomial time for a broad class of choice models; see Appendix A. 2 for details. In the case of ties, we choose any of the sets with the smallest number of products. We can think of $r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right)$ as the discounted revenue associated with product $i$, where the discount factor $\Psi\left(I_{i}^{t-1} / c_{i}\right)$ is determined by the penalty function and it depends on the fraction of the initial inventory that remains. As we discuss in Section 1.4.1, $\Psi\left(I_{i}^{t-1} / c_{i}\right)$ corresponds to a dual solution to $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$. Namely, for each $t$, using $\Psi\left(I_{i}^{t-1} / c_{i}\right)$, we can construct a feasible solution for the dual of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$, and the value of this feasible dual solution is within a "constant" factor of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$.

The main idea behind the algorithm is simple. Sometimes it might be better to sell a product with a lower marginal revenue but a high inventory level, than to sell a product with high marginal revenue but few remaining inventory. This is because the future customers might only be interested in the products with low (or no) inventory, and if we have already sold those products, we would lose on these profitable opportunities. ${ }^{7}$ The penalty function thus protects against the uncertainty in future customer types. The following example shows what would happen if we ignore the inventory level, and only offer assortments with the highest revenues; see Mehta et al. (2007) and Bernstein et al. (2011) for similar examples.

## Example 1.3.1. (Myopic Policy and Why Inventory Levels Matter?)

Consider the Myopic policy that does not take into account any inventory level. The policy corresponds to the following penalty function: $\Psi(x)=\mathbb{1}[x>0]$, i.e., $\Psi(x)$ is equal to 1 if the remaining inventory of the product is positive and is equal to 0 otherwise. This algorithm, at any period, offers the assortment,

[^5]among products with positive inventory, that maximizes the expected revenue. The following scenario shows that the competitive ratio of the Myopic policy is at most $\frac{1}{2}$. Suppose the length of the horizon is equal to $T$. There are two products with the following parameters, $r_{1}=1+\epsilon, r_{2}=1$, and $c_{1}=c_{2}=\frac{T}{2}$. We have two customer types. The first type arrives during periods $1, \ldots, \frac{T}{2}$, and the second type arrives in periods $\frac{T}{2}+1, \ldots, T$. For $t \in\left\{1, \ldots, \frac{T}{2}\right\}, \phi_{1}^{z_{t}}(\{1\})=\phi_{2}^{z_{t}}(\{2\})=1$, and $\phi_{i}^{z_{t}}(S)=0$ otherwise. For $t \in\left\{\frac{T}{2}+1, \ldots, T\right\}, \phi_{1}^{z_{t}}(\{1\})=1$ and $\phi_{i}^{z_{t}}(S)=0$ otherwise. In this setting, a customer of the first type is interested in both products 1 and 2 , while the second type is only interested in product 1.

The Myopic policy will allocate all the inventory of product 1 to the customers that arrive in periods $1,2, \ldots, \frac{T}{2}$, and obtains a revenue of $\frac{T}{2}(1+\epsilon)$. However, the optimal solution of Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ allocates all the inventory of product 2 to the customers that arrive in periods $1,2, \ldots, \frac{T}{2}$, and then sells product 1 to customers that arrive afterwards in period $\frac{T}{2}+1, \ldots, T$, yielding a revenue of $\frac{T}{2}(2+\epsilon)$. Note that $\frac{T(1+\epsilon)}{T(2+\epsilon)} \leq \frac{1}{2}+\epsilon$. Since $\epsilon$ can be arbitrarily small, the competitive ratio of the Myopic policy is at most $\frac{1}{2} .{ }^{8}$

The above example, though rather stylized, highlights the importance of inventory levels in assortment planning. We will show that, by discounting the revenue of each product based on its remaining inventory, the Inventory-Balancing algorithms obtain a better competitive ratio. Throughout this chapter, we impose the following mild assumption on the choice models.

Assumption 1 (Substitutability). For all $z \in \mathcal{Z}, S \in \mathcal{S}$, and $i \neq j, \phi_{i}^{z}(S) \geq \phi_{i}^{z}(S \cup\{j\})$.

The above assumption implies that adding another product to an assortment does not increase the probability of selling other products in the assortment. It is easy to verify that the above assumption encompasses all choice models that are consistent with random utility maximization ${ }^{9}$, including the multinomial logit choice model, the nested logit, and many others; see Appendix A. 2 for details. The above assumption leads to the following desirable property. The proof is in Appendix A.4.

[^6]Lemma 1.3.2. Under Assumption 1, the Inventory-Balancing algorithm never offers an assortment that includes a product with zero remaining inventory.

In Section 1.5.3, we relax Assumption 1 and extend our results to a more general setting where stockouts are allowed. We now use Lemma 1.3.2 to establish the competitive ratio. The proof is given in Appendix B.1.1.

Theorem 1.3.3 (Competitive Ratio). Let $c_{\mathrm{MIN}}=\min _{i=1, \ldots, n} c_{i}$. Suppose $\Psi$ is an increasing, concave, and twice-differentiable penalty function. The competitive ratio of Inventory-Balancing algorithm with a penalty function $\Psi$ is at least equal to $\alpha_{c_{\text {MIN }}}(\Psi)$, where

$$
\alpha_{C_{\mathrm{MIN}}}(\Psi)=\min _{x \in\left[0,1-\frac{1}{c_{\mathrm{MIN}}}\right]}\left\{\frac{1-x}{\frac{1}{c_{\mathrm{MIN}}}+1-\Psi(x)+\int_{x+\frac{1}{c_{\mathrm{MIN}}}}^{1} \Psi(y) d y}\right\}
$$

We emphasize that the competitive ratio in the above theorem only depends on $c_{\text {MIN }}$ and penalty function $\Psi$, and the ratio does not depend on the length of the horizon $T$. Hence, the above result holds when $T$ increases to infinity. It also holds for any sequence of customer types of arbitrary length, even if the sequence is chosen by an adversary.

Many of the previous results in the literature are established for an asymptotic regime where the size of initial inventories $c_{\text {MIN }}$ and the length of the horizon $T$ tend to infinity. The justification for the asymptotic analysis is that often the initial inventory of products and the number of customers are large. In this asymptotic regime, we can simplify the expression for the competitive ratio. Define:

$$
\alpha(\Psi):=\alpha_{\infty}(\Psi)=\min _{x \in[0,1]}\left\{\frac{1-x}{1-\Psi(x)+\int_{x}^{1} \Psi(y) d y}\right\} .
$$

We observe that the competitive ratio of the algorithm improves slightly as $c_{\text {MIN }}$ becomes larger. For instance, for the polynomial penalty functions $\Psi(x)=\sqrt{x}$, the competitive ratio with $c_{\text {MIN }}=2,5$, and 10 is, respectively, $0.52,0.55$, and 0.57 . This ratio approaches $\alpha(\sqrt{x})=0.60$ as $c_{\text {MIN }}$ grows.

As a corollary of Theorem 1.3.3, we can show that the competitive ratio is at least $\frac{1}{2}$ for any increasing concave function. Therefore, by taking into account the remaining inventory levels, we obtain a better performance guarantee than a Myopic policy that ignores inventory.

Corollary 1.3.4. For the Inventory-Balancing algorithm with linear penalty function (LIB), $\Psi(x)=x$, the competitive ratio $\alpha_{c_{\text {MIN }}}(x)$ is equal to $\frac{1}{2}$, for any $c_{\text {MIN }} \geq 1$. For any increasing strictly concave and differentiable penalty function, the competitive ratio is strictly greater than $\frac{1}{2}$.

The proof is given in Appendix A.4. Note that this is a worst-case performance guarantee and does not imply that the Inventory-Balancing algorithms outperform the Myopic policy on every sequence of customers. In practice, as suggested by our numerical simulations, we expect that the IB algorithms, and even the Myopic policy, to often perform better than their theoretical worst-case bounds; see Section 1.6.2.

The choice of the penalty function determines the trade-offs between the revenue from selling a product and the value of the remaining inventory. For a linear penalty function $\Psi(x)=x$, the derivative is always 1 , and a reduction in a unit of inventory has the same penalty, regardless of the inventory level. On the other hand, the derivative of the exponential penalty function $\Psi(x)=\frac{e}{e-1}\left(1-e^{-x}\right)$ is given by $\frac{e}{e-1} e^{-x}$, which decreases from 1.58 at $x=0$ to 0.58 at $x=1$. Under the exponential penalty function, consuming one unit of inventory incurs a higher penalty when the inventory is scarce. In regimes with high demand and low inventory, we would expect that the Inventory Balancing algorithm with an exponential penalty function (EIB) to be more conservative and hold back more products to hedge against future arrivals. As we show in the next section, the best competitive ratio can be obtained using an exponential penalty function $\Psi(x)=\frac{e}{e-1}\left(1-e^{-x}\right)$.

The complexity of the IB Algorithm is determined by the complexity of determining the optimal assortment for each arriving customer, corresponding to solving the combinatorial optimization problem $\max _{S \in \mathcal{S}} \sum_{i \in S} r_{i} \Psi\left(I^{t-1} / c_{i}\right) \phi_{i}^{z t}(S)$. This problem can be solved efficiently for a broad class of choice models; see Appendix A.2.

### 1.3.1 The Tight Upper Bound on the Competitive Ratio

We start this section by providing an upper bound on the competitive ratio. Then, in Theorem 1.3.6, we show that an IB algorithm with an exponential penalty function achieves this upper bound, showing that our proposed method achieves an optimal competitive ratio.

Lemma 1.3.5 (Upper Bound on the Competitive Ratio). For any number of products $n$, we can construct a non-stationary stochastic process for customer arrivals where for every deterministic algorithm (including
the optimal dynamic program), there exists a sequence of customer types $\left\{z_{t}\right\}_{t=1}^{T}$ such that the revenue of the algorithm is at most a fraction $\rho_{n}=\left(\frac{1}{n} \sum_{j=1}^{n} \min \left\{\sum_{t=1}^{j} \frac{1}{n-t+1}, 1\right\}\right)$ of Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$.

For instance, for $n=2,5$, and 20 , the upper bound $\rho_{n}$ is respectively equal to $0.75,0.69$, and 0.64 and $\rho_{n}$ approaches $\lim _{n \rightarrow \infty} \rho_{n}=1-\frac{1}{e} \approx 63 \%$ as number of products $n$ increases.

In the proof of the above lemma, we construct a stochastic process that consists of $n$ products. The per-unit revenue from each product is equal to 1 and the initial inventories are equal to $\frac{T}{n}$. Think of $T$, the length of the horizon, as a very large number (that would tend to infinity) and a multiple of $n$. The number of types is equal to $2^{n}-1$. Each type corresponds to a nonempty set $\Theta$ of products that a customer of that type equally likes; the "no-purchase" probability for all types is equal to zero. Note this is a special case of the multinomial logit choice model where the all the products have weight either 0 or 1 . The arrival process is defined as follows: customer arrives in $n$ phases of equal length, that is, the number of customers in each phase is $\frac{T}{n}$. All the customers in each phase have the same type. Customers in the first phase are interested in all the products. After that, in each phase, customers randomly lose interest in one of the products of interest in the previous phase; i.e., there are $n$ ! sequences of customer arrivals, each with equal probability.

Now, we show that the Exponential Inventory-Balancing algorithm achieves the optimal competitive ratio.

Theorem 1.3.6 (Exponential IB Achieves The Optimal Competitive Ratio). The competitive ratio of the Inventory-Balancing algorithm with exponential penalty function (EIB), $\Psi(x)=\frac{e}{e-1}\left(1-e^{-x}\right), x \in[0,1]$, approaches $\left(1-\frac{1}{e}\right)$ as $c_{\text {MIN }}$ increases to infinity, i.e., $\alpha\left(\frac{e}{e-1}\left(1-e^{-x}\right)\right)=1-\frac{1}{e}$. Moreover, no algorithm, deterministic or randomized, that does not know the sequence of customer types in advance can obtain a competitive ratio better than $1-\frac{1}{e}$.

The proof is given in Appendix A.4.1. The first part of the proof is based on Theorem 1.3.3. The second part follows from applying Yao's Lemma (Yao 1977) to Lemma 1.3.5. Yao's Lemma implies that the competitive ratio of any randomized algorithm that does not know the input sequence in advance is bounded by the competitive ratio of any deterministic algorithm that knows the distribution over the input sequence.

We note that the upper bound of $\left(1-\frac{1}{e}\right)$ applies to all deterministic algorithms, including the optimal dynamic programming (Bernstein et al. 2011) and re-optimization (Jasin \& Kumar 2012) policies. Thus, by
the theorem above, in terms of the competitive ratio, the Inventory-Balancing algorithm with an exponential penalty function is optimal for this problem. We remark that this notion of optimality does not imply that the algorithm would obtain the highest revenue from every sequence of customers.

By Theorem 1.3.3, the competitive ratio of the Exponential Inventory-Balancing algorithm with limited inventory, for instance for $c_{\text {MIN }}=5,10,20$, and 30 is, respectively, equal to $0.57,0.60,0.61$, and 0.62 . The ratio approaches 0.63 rather rapidly as $c_{\text {MIN }}$ grows. We emphasize that these ratios hold for all values of $T$, including the asymptotic regime where $T$ increases to infinity (at a possibly faster rate than $c_{\text {MIN }}$ ).

Moreover, as shown in our numerical experiments, our algorithms often perform much better than the worst-case guarantee bounds.

### 1.4 Stochastic I.I.D. Arrivals

The competitive ratio of our IB algorithm in Theorems 1.3.3 and 1.3.6 hold for any arbitrary, possibly adversarially chosen, sequence of customer types. It turns that our IB algorithm performs even better if the customer arrivals follow a stochastic process; that is, when the sequence of customers $\left\{z_{t}\right\}_{t=1}^{T}$ is generated by a stochastic process that is known in advance. In this model, the optimal sequence of assortments can be planned by solving a multi-dimensional dynamic program; for more details see Appendix A.5. Not surprisingly, this approach suffers from the curse of dimensionality, even for stationary processes; see Bernstein et al. (2011).

Under stochastic models, although a dynamic programming approach may be intractable, there is room for natural and powerful heuristics. First observe that for any algorithm $A$, the expected revenue of the algorithm, denoted by $\mathrm{E}_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{A}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]$, is well-defined where $\mathrm{E}_{\left\{z_{t}\right\}_{t=1}^{T}}$ is the expectation with respect to the sequence of customers. Recall that by definition, the expectation over customers' choices is taken into account by $\operatorname{Rev}_{A}(\cdot)$. Furthermore, we can establish an upper bound on the revenue of the algorithm. The proof is omitted due to its similarity to Lemma 1.2.1.

Lemma 1.4.1 (Revenue Upper Bound for I.I.D. Arrivals). Suppose the types of the customers is drawn, independently and identically, from a known distribution. Let $\eta^{z}$ be the expected number of customers of
type $z \in \mathcal{Z}$. The expected revenue of any algorithm $A, E_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{A}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]$, is bounded by the optimal value of the linear program Primal-S defined below: ${ }^{10}$

$$
\begin{array}{rll}
\text { MAXImize } \quad \sum_{z \in \mathcal{Z}} \sum_{S \in \mathcal{S}} \sum_{i \in S} \eta^{z} r_{i} \phi_{i}^{z}(S) y^{z}(S) & \\
\text { SUBJECT TO: } \quad \sum_{z \in \mathcal{Z}} \sum_{S \in \mathcal{S}} \eta^{z} \phi_{i}^{z}(S) y^{z}(S) & \leq c_{i} & 1 \leq i \leq n,  \tag{Primal-S}\\
\sum_{S \in \mathcal{S}} y^{z}(S) & =\eta_{z} & \forall z \in \mathcal{Z}, \\
y^{z}(S) & \geq 0 & \forall z \in \mathcal{Z}, S \in \mathcal{S}
\end{array}
$$

In the above linear program, $y^{z}(S)$ is the probability of offering the set $S$ to a customer of type $z$. As before, we also denote the optimal solution of the above linear program by Primal-S. Note that Lemma 1.2.1 provides a stronger upper bound since it holds for every customer sequence, while the above upper bound holds only in expectation. Theorem 1.3 .3 provides a bound on the performance of our algorithms with respect to the upper bound in Lemma 1.2.1. When we have I.I.D. arrivals and we use Primal-S as the benchmark, as stated by the proposition below, we obtain an even stronger performance guarantee for our algorithms.

Proposition 1.4.2 (Improved Performance Guarantee in the I.I.D. Arrival Model). Suppose in every period $t$, the type of the arriving customer is drawn independently and identically from a common distribution over the set of types $\mathcal{Z}$. In the asymptotic regime where $c_{\text {MIN }}$ and $T$ increase to infinity with $\frac{T}{c_{\text {MIN }}}=k$ for some positive integer $k$, then with high probability, the Inventory-Balancing algorithms with linear (LIB) and exponential (EIB) penalty functions satisfy the following inequalities:

$$
\lim _{T, c_{\mathrm{MIN}} \rightarrow \infty} \frac{E_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{L I B}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]}{\text { Primal-S }} \geq 0.72 \quad \text { and } \quad \lim _{T, c_{\mathrm{MIN}} \rightarrow \infty} \frac{E_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{E I B}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]}{\operatorname{Primal-S}} \geq 0.75
$$

where the expectations in $E_{\left\{z_{t}\right\}_{t=1}^{T}}[\cdot]$ is taken with respect to the sequence of arriving customers.

The proof is given in Appendix A.4.2. The basic idea is to construct a (Factor Revealing) linear program denoted by FRLP. With high probability, every solution obtained by the Inventory-Balancing algorithm corresponds to a feasible solution of FRLP and the objective corresponds to the ratio of the expected revenue

[^7]$\mathrm{E}_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{E I B}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]$ of the IB algorithm and Primal-S. FRLP is parameterized by a discretization parameter $\epsilon$. For each $\epsilon$, we can solve the linear program to determine the lower bound on the competitive ratio.

### 1.4.1 Motivating IB Algorithms via Dual-Based Heuristics

In this section, we provide a motivation and intuition behind our Inventory-Balancing algorithm. The discounted revenue index in our IB algorithm are proxies (approximation) for the dual variables. To see this, consider the following policy for I.I.D. arrival model.

1. Observe the type of the first $\epsilon T$ customers. ${ }^{11}$
2. Solve the dual of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ for the first $\epsilon T$ customer:

$$
\begin{array}{ll}
\text { MINIMIZE } & \sum_{t=1}^{\epsilon T} \lambda^{t}+\sum_{i=1}^{n} \theta_{i} \epsilon c_{i} \\
\text { SUBJECT TO: } & \lambda^{t} \geq \sum_{i \in S}\left(r_{i}-\theta_{i}\right) \phi_{i}^{z_{t}}(S)  \tag{1.1}\\
& \theta_{i} \geq 0 \\
& 1 \leq t \leq \epsilon T, S \in \mathcal{S}, \\
& 1 \leq i \leq n .
\end{array}
$$

Let $\theta_{i}^{*}(\epsilon T), i=1,2, \ldots, n$, be the solution of the linear program above.
3. For each subsequent customer of type $z \in \mathcal{Z}$, we offer an assortment $S^{t}$ :

$$
S^{t}=\arg \max _{S \in \mathcal{S}} \sum_{i \in S}\left(r_{i}-\theta_{i}^{*}(\epsilon T)\right) \phi_{i}^{z}(S)
$$

Note that the algorithm does not need to know the distribution in advance. Following from the results of Devenur \& Hayes (2009), Agrawal et al. (2009), Feldman et al. (2010), Jaillet \& Lu (2012), we can show that this algorithm is asymptotically optimal for I.I.D. stochastic arrivals. ${ }^{12}$ Since the proof is similar to the existing literature, we omit the details.

[^8]Note that the selection rule of the above heuristic is similar to our algorithms by replacing $r_{i}-\theta_{i}^{*}$ with $r_{i} \times \Psi\left(I_{i}^{t-1} / c_{i}\right)$. In our Inventory-Balancing algorithms, however, we do not assume any patterns for the arrival of future customers. Thus, we can think of the discounted revenue index $r_{i} \times \Psi\left(I_{i}^{t-1} / c_{i}\right)$ as the "estimate" of the dual parameters based on the current inventory levels. This index value does not require any forecasting, and by choosing an appropriate penalty function $\Psi$, we can obtain an optimal competitive ratio, as shown by Theorem 1.3.3. See Feldman et al. (2010) who apply similar ideas to the online allocation of display advertisement.

### 1.5 Extensions

In this section, we discuss how our policies can be extended to more general settings and incorporate additional information about the customers' choice models or arrival patterns.

### 1.5.1 Incorporating Partial Information and Learning Customer types?

So far, we assumed that in each period, the algorithm knows the customer choice models. Namely, $\phi_{i}^{z}(S)$ of the customer of type $z$ is the exact value of the selection probability. But, this may not always be the case. The firm may learn the choice model associated with each customer type over time. For instance, in our numerical simulations, we associate the type of each customer to his or her location. We consider the case where the firm estimates the parameters of an MNL model for each location by learning from purchases of the previous customers from that location.

Suppose the arriving customer in period $t$ is of type $z$. We denote by $\bar{\phi}_{i}^{z}(S)$ the true selection probability of product $i$, when assortment $S$ is offered to a customer of type $z$. In this environment, $\phi_{i}^{z t}(S)$ represents the current estimation of the selection probabilities. ${ }^{13}$ These estimates can be obtained using partial information or historical data. We do not require specifics about how these estimations are made. But a good example would be when the customer types are drawn from a stationary distribution and the parameters of the choice model are learned from observing customers' choices. Under standard assumptions, we expect that the estimated selection probabilities would converge to the true selection probabilities, see Appendix A.1.4 for an example.

[^9]The following proposition provides a lower bound on the competitive ratio when we have estimation errors.

Proposition 1.5.1 (Competitive Ratio with Estimation Errors). For each $t$, let $\epsilon^{t}=\max _{i, S}\left|\phi_{i}^{z_{t}}(S)-\bar{\phi}_{i}^{z_{t}}(S)\right|$ be the random variable corresponding to the maximum estimation error in period $t$. Suppose the InventoryBalancing algorithm sells at least one unit of each product. Then, the competitive ratio of the InventoryBalancing algorithm is at least equal to

$$
\min _{x \in\left[0,1-\frac{1}{c_{\mathrm{MIN}}}\right]}\left\{\frac{1-x}{\frac{1}{c_{\mathrm{MIN}}}+1-\Psi(x)+\int_{x+\frac{1}{c_{\mathrm{MIN}}}}^{1} \Psi(y) d y+\frac{2}{c_{\mathrm{MIN}}} E\left[\sum_{t=1}^{T} \epsilon^{t}\right]}\right\}
$$

The proof is given in Appendix A.4.3. Note that when there is no estimation error, i.e., $\phi_{i}^{z_{t}}(S)=\bar{\phi}_{i}^{z_{t}}(S)$, then the above expression is the same as the competitive ratio in Theorem 1.3.3. Furthermore, the only assumption made on the estimations errors is that the algorithm should sell at least one unit of each product. This assumption is made mainly for technical reasons to rule out the situations that the estimations are so far off that the algorithm never sells the products sold by the optimal solution (that knows the true estimations). We expect this condition to be satisfied when the estimation errors are small or if they vanish over time as the mechanism gathers more data about each type.

In Appendix A.1.4, we present an example that demonstrates how learning customer types can be incorporated in our framework. Furthermore, in Appendix A.3.3, using numerical simulations, we evaluate the performance of the IB algorithms when the selection probability for each customer type is unknown and must be estimated from data collected in earlier periods.

### 1.5.2 Incorporating (Uncertain) Information About Arrival Pattens

The Inventory-Balancing algorithms do not rely on any forecast of the future customer arrivals; however, if such forecast exists, it could potentially be used to improve the performance of the algorithms. Consider a heuristic $\mathcal{L}$, for example, the linear program re-optimization, that relies on the distribution (e.g., the estimated number) of the customers of each type. This heuristic would perform well if the estimations are accurate, however, it performs poorly when the estimate turned out to be inaccurate or there is a high degree of uncertainty; see Section 1.6. We propose a family of algorithms called the Hybrid algorithm that combines
the solution of such heuristics and IB algorithms. These algorithms incorporate additional information about the arrival sequence and at the same time maintain a reasonable competitive ratio in unpredictable scenarios; see Mahdian et al. $(2007,2012)$.

The Hybrid algorithm, given below, is parameterized by a number $\gamma \geq 1$. This parameter controls the extent to which one would rely on heuristic $\mathcal{L}$.

## The HYbrid Algorithm with parameter $\gamma$

Upon the arrival of the customer in period $t \in\{1, \ldots, T\}$, of type $z_{t}$ :

- Let $S_{\mathcal{L}}^{t}$ be the set that heuristic $\mathcal{L}$ recommends in period $t$.
- Offer the assortment $S_{\mathcal{L}}^{t}$ if:

$$
\gamma\left(\sum_{i \in S_{\mathcal{L}}^{t}} \Psi\left(I_{i}^{t-1} / c_{i}\right) r_{i} \phi_{i}^{z_{t}}(S)\right) \geq \max _{S \in \mathcal{S}}\left\{\sum_{i \in S} \Psi\left(I_{i}^{t-1} / c_{i}\right) r_{i} \phi_{i}^{z_{t}}(S)\right\}
$$

- Otherwise, offer an assortment $S^{t} \in \arg \max _{S \in \mathcal{S}} \sum_{i \in S} \Psi\left(I_{i}^{t-1} / c_{i}\right) r_{i} \phi_{i}^{z t}(S)$.

The next proposition provides a lower bound on the competitive ratio of the Hybrid algorithm.

Proposition 1.5.2 (Competitive Ratio of the Hybrid Algorithm). Suppose $\Psi$ is an increasing, concave, and twice-differentiable penalty function. As $c_{\text {MIN }} \rightarrow \infty$, the competitive ratio of the Hybrid algorithm with a penalty function $\Psi$ and parameter $\gamma$, and for any heuristic $\mathcal{L}$, is at least equal to $\alpha_{\infty}^{\gamma}(\Psi)$, where

$$
\alpha_{\infty}^{\gamma}(\Psi)=\min _{x \in[0,1]}\left\{\frac{1-x}{\gamma(1-\Psi(x))+\int_{x}^{1} \Psi(y) d y}\right\}
$$

For example, for the exponential penalty function, $\alpha_{\infty}^{\gamma}\left(\frac{e}{e-1}\left(1-e^{-x}\right)\right)$, for $\gamma=1.5$ and $\gamma=2$, is approximately equal to 0.48 and 0.39 , respectively. The proof is very similar to proof of Theorem 1.3.3 and is omitted. The main idea is to assign $\lambda^{t}=\gamma\left(\sum_{i \in S^{t}} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) \phi_{i}^{z_{t}}\left(S^{t}\right)\right)$. Intuitively, we are
extending the feasible region of the dual problem which allows the algorithm to follow the heuristics on the recommendations that are considered "safe".

Our simulation results in Section 1.6.2 and Appendix A.3.1, we consider a Hybrid algorithm that combines the EIB algorithm (Theorem 1.3.6) and the linear program re-optimization heuristic. We show that the Hybrid algorithm outperforms the Inventory-Balancing algorithm, when the number of customers is known in advance by the re-optimization policies. On the other hand, when the number of customers is uncertain, the Hybrid algorithm outperforms the re-optimization methods.

### 1.5.3 Beyond Substitutability

In this section, we explain how we can relax Assumption 1. Recall that, by Lemma 1.3.2, Assumption 1 implies that our algorithm does not benefit from showing a product with no remaining inventory. But, sometimes the assumption may not hold. For instance, when the dissimilarity parameters in the nested logit model are larger than 1 (Davis et al. 2011, Bhat 2002), or there are externalities among the products ${ }^{14}$.

More specifically, we assume that the choice model satisfies the following property: Suppose a customer is offered a set $S$ and then she chooses product $i \in S$. Then, the customer buys product $i$ if it has positive inventory, otherwise leaves without making any purchase. Under this choice model, we can use the Inventory-Balancing algorithm exactly in the same way as before. The inventory level of the products that are stockout remains at 0 even if the product is shown to the customer.

In this model, the optimal revenue can be upper-bounded by the linear program below.

MAXIMIZE $\quad \sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \sum_{i=1}^{n} r_{i} \phi_{i}^{z_{t}}(S) y^{t}(S)-\sum_{i} r_{i} w_{i}$

SUBJECT TO: $\quad \sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \phi_{i}^{z_{t}}(S) y^{t}(S)-w_{i} \leq c_{i} \quad 1 \leq i \leq n$, $\sum_{S \in \mathcal{S}} y^{t}(S)=1 \quad 1 \leq t \leq T$,
$y^{t}(S) \geq 0 \quad 1 \leq t \leq T, \quad S \in \mathcal{S}$,
$w_{i} \geq 0 \quad 1 \leq i \leq n$.

[^10]This linear program is the same as the linear program $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ given in Section 3.1, with a new set of variables: $w_{i}$ denotes the amount of times that a customer selects product $i$ after its inventory hits 0 . In this case, the product will not be allocated to the customer.

We now argue that the algorithm obtains the same competitive ratio as before. The argument is based on the following observation. The dual of the linear program above is as follows:

$$
\begin{array}{clll}
\text { MINIMIZE } & & \\
& \sum_{t=1}^{T} \lambda^{t}+\sum_{i=1}^{n} \theta_{i} c_{i} & \\
\text { SUBJECT TO: } & \lambda^{t} & \geq \sum_{i=1}^{n} \phi_{i}^{z_{t}}(S)\left(r_{i}-\theta_{i}\right) & \\
& & 1 \leq t \leq T, S \in \mathcal{S}, \\
& \theta_{i} & \leq r_{i} & 1 \leq i \leq n, \\
& \theta_{i} & \geq 0 & 1 \leq i \leq n .
\end{array}
$$

Note that compared to the previous dual in Section B.1.1, the above linear program has a new set of constraints: $\theta_{i} \leq r_{i}$. However, these constraints are satisfied by our construction of the feasible solution in the proof of Theorem 1.3.3. Therefore, the ratio of the primal and dual solutions and the competitive ratio of the algorithm are the same as are described in Theorem 1.3.3.

### 1.6 Numerical Experiments

In this section, we numerically evaluate our Inventory-Balancing (IB) and Hybrid algorithms, and compare them with existing methods in the literature. The first IB algorithm is the Linear Inventory-Balancing (LIB) with a linear penalty function $\Psi(x)=x$. The second algorithm is the Exponential Inventory-Balancing (EIB) with penalty function $\Psi(x)=\frac{e}{e-1}\left(1-e^{-x}\right)$.

In the next section, we describe the dataset and the simulation setting. Then, we compare the revenues and running times of our IB and Hybrid algorithms with the Myopic policy and existing heuristics based on resolving linear programs. Section 1.6 .3 shows the benefits of differentiating customers by their types. Finally, in Section 1.6.4, we present a synthetic simulation to compare our algorithm with the Myopic policy.

| Location | Top Three DVDs |
| :--- | :--- |
| Jersey City, NJ | Lost; Sin City; The Simpsons |
| Manhattan NY | Lost; Sin City; Bob Dylan - No Direction Home |
| Orlando, FL | Lost; The Muppet Show: Season One; Firefly |
| Miami, FL | What the Bleep Do We Know!?; Sin City; Lost |
| San Jose, CA | Star Wars: Episode III; Lost; The Complete Thin Man Collection |
| Beverly Hills, CA | The Simpsons; L'Auberge Espagnole; Sin City |
| Inglewood, CA | Lost; Sin City; Final Fantasy VII - Advent Children |
| Dallas, TX | My Big Fat Greek Wedding; Lost; Shark Tale |
| Urbana Champaign, IL | Firefly; Lost; The Simpsons |
| Charlotte, NC | Lost; Curb Your Enthusiasm; Family Guy Presents Stewie Griffin |

Table 1.2: The locations used in the simulations and the top three best-selling DVDs in each location.

### 1.6.1 Dataset and Simulation Setting

We use the DVD sales data from a large online retailer, and consider DVDs that are sold during a four-month period from June 1, 2005 through September 30, 2005. During this period, the retailer sold over 5.7 million DVDs in the United States, spanning across 55,875 DVD titles.

We consider the location of each customer as her type, and for our analysis, we choose 10 different locations that would reflect diverse purchasing patterns of the customers. Table 1.2 shows the list of locations and the top three DVDs with the highest sale volumes in each location ${ }^{15}$. In our experiments, we consider a total of 73 DVD titles ( $n=73$ ), obtained by taking the union of the top 20 DVDs in each location and removing duplicates. For each DVD $i=1,2, \ldots, 73$, we set the revenue $r_{i}$ to be the average selling price of the DVD during this period. The DVD prices range from $\$ 9-\$ 81$, with more than $50 \%$ of the DVDs are priced less than $\$ 20$, and the average and standard deviation of the DVD selling prices are $\$ 25.2$ and $\$ 12.9$, respectively. We assume that all DVDs have the same initial inventory, with no replenishment. Finally, we assume a multinomial logit choice model for each customer type. For each type $z \in \mathcal{Z}$, we estimate the preference weight parameters $\left(v_{0}^{z}, v_{1}^{z}, \ldots, v_{n}^{z}\right) \in \mathbb{R}_{n}^{+}$by computing the maximum likelihood estimate based on the sales data in each location.

In the simulations, each problem class is characterized by two parameters: the loading factor (LF) and the coefficient of variation (CV) associated with the proportions of customer of each type. Note that the

[^11]coefficient of variation is the ratio of standard deviation to the mean; that is, for a given fixed mean, larger CV means that there is a larger uncertainty in the type of an arriving customer. The loading factor $(L F)$ is the ratio between the (expected) total number of customers and the total number of units. For a given loading factor, we can determine the (expected) total number of customers. The number of customers of each type is generated by multiplying the total number of customers by a 10 -dimensional Dirichlet random variable $\boldsymbol{\beta}=\left(\beta_{1}, \ldots, \beta_{10}\right)$, where $\beta_{i}$ represents the proportion of customers type $i^{16}$. Thus, on average, each type is equally likely ${ }^{17}$. Once the number of customers of each type is generated, the order of arrivals is random. We further assume that a single customer arrives in each period.

### 1.6.2 Performance Evaluation

In this section, we compare the performance of our two IB algorithms to the Hybrid algorithm, the Myopic policy, and the LP-based heuristics.

The Myopic policy offers the assortment with the maximum expected revenue among all assortments that include only products with positive remaining inventory. ${ }^{18}$

The first LP-based algorithm, called the LP One-Shot (LPO), is a heuristic based on the solution of the linear program Primal-S where the expected number of each customer type $\eta^{z}$ is equal to $0.1 \times \mathrm{E}[T]{ }^{19}$. This follows because each of the 10 customer types is, on average, equally likely, and $\mathrm{E}[T]$ is the expected number of customers. Although the linear program (LP) Primal-S has exponentially many variables, we can solve it efficiently using techniques from Gallego et al. (2004), Liu \& van Ryzin (2008a), Topaloglu (2013). Under the LPO, we solve the linear program Primal-S exactly once and denote the solution by $\left\{\bar{y}^{z}(S): z \in \mathcal{Z}, S \in \mathcal{S}\right\}$. LPO policy is constructed as follows: upon the arrival of a customer of type $z$, offer her assortment $S$ with probability $\bar{y}^{z}(S) / \sum_{S^{\prime} \in \mathcal{S}} \bar{y}^{z}\left(S^{\prime}\right)$.

[^12]We also consider an adaptive variation of this policy, called ALPO, that excludes any product with zero remaining inventory from an offered set. Note that ALPO is an adaptive policy because the assortment that we offer to each customer may change over time, depending on the product availability.

The second LP-based algorithm is called LP Resolving (LPR). Under this algorithm, the linear program Primal-S is resolved periodically every $h$ periods, with the up-to-date inventory levels and forecasts of the proportion of customers of each type. We denote this heuristic by LPR ${ }^{h}$. In this heuristic, at the beginning of a resolving period $t, \eta^{z}$ is replaced by $\eta^{z}(t)$. We define $\eta^{z}(t)$ as the empirically estimated expected number of customers of type $z$ that would arrive between periods $t$ and $T, \eta^{z}(t)$ is equal to the fraction of customers of type $z$ that arrived between time 1 and $t-1$ multiplied by $\frac{\mathrm{E}[T-t+1 \mid T \geq t]}{t-1}$, where $\mathrm{E}[T-t+1 \mid T \geq t]$ is the expected number of future customers given that $t$ customers have already arrived ${ }^{20}$. Also, we set $\eta^{z}(1)=0.1$. As the system evolves, the realized number of customers of each type will differ from the expected value, and $\eta^{z}(t)$ incorporates these differences into the estimates. The LPR algorithm also updates the inventory level to reflect the current remaining inventory of each product.

The Hybrid algorithm combines IB algorithms with LP based heuristics, see Section 1.5.2. ${ }^{21}$ In particular, we assume that Hybrid algorithm incorporates the solutions of EIB algorithm and the LPR ${ }^{500}$ heuristic with parameter $\gamma=1.5$ and 2. For $\gamma=1.5$ and $\gamma=2$, the competitive ratio of the Hybrid algorithm is at least equal to 0.48 and 0.39 , respectively. Recall that parameter $\gamma$ should be greater than 1 and the larger $\gamma$, the more the Hybrid algorithm applies the LPR heuristic. We denote the Hybrid algorithm with parameter $\gamma$ by $\mathcal{H}_{\gamma}$.

We evaluate the aforementioned algorithms by comparing the revenue with the upper-bound given by clairvoyant optimal solution which knows the arrival sequence of all the customers in advance; see Lemma 1.2.1 for more details.

Simulation Parameters: We consider 9 problem classes, corresponding with loading factors 1.4, 1.6, and 1.8 , and coefficients of variation of $2.0,1.0$, and 0.1 . We choose the loading factor greater than 1 because we are interested in cases where the demand exceeds supply. Each problem class consists of 250 problem instances. In each problem instance, we set the initial inventory levels to 100 , i.e., $c_{i}=100, i=1,2, \ldots, 73$,

[^13]and the length of the horizon is randomly and uniformly chosen from an interval $[\underline{T}, \bar{T}]^{22}$. In table 1.3, for each problem class, we present the average revenue of each algorithm as a percentage of the upper bound, averaged over 250 instances in each problem class. We present the results for periodic LPR with $h=50$ and $h=500$.

As expected, the upper bound increases as the loading factor increases and the variability decreases. Both the LIB and EIB algorithms surpass all other policies. In addition, LPR policies cannot obtain more than $92 \%$ of the upper bound, even if we increase the frequency of resolving ${ }^{23}$. In all cases, the $\mathcal{H}_{1.5}$ algorithm performs better than the $\mathcal{H}_{2}$ algorithm because the larger $\gamma$, the more the Hybrid algorithm relies on the LPR heuristic. Note that the Hybrid algorithms get $5 \%-12 \%$ more revenue than the LP resolving heuristics. This highlights the benefits of combining the solution of IB policies and LP based heuristics. As we will discuss later, we also consider the worst-case performance of all policies; see Figure 1.1 and Appendix A.3.1.

In most problem classes, the LPO algorithm has the lowest revenue and its performance decreases by increasing CV and loading factor. The main reason for the poor performance of the LPO algorithm is the lack of strategy adjustment. The LPO algorithm solves the linear program exactly once and does not update its decision based on the current inventory level (it just checks whether the product is available, in case of ALPO), nor does it consider the proportion of each customer type. When the uncertainty in type of arriving customers and the number of customers is high, this lack of adjustment can result in significant reduction in the performance of the LPO algorithm, e.g., for $\mathrm{CV}=0.1$ and loading factor of 1.4, the ALPO algorithm obtains $93.5 \%$ of the optimal solution while for $C V=2$, it gets $81.4 \%$ of that. Note that the ALPO algorithm that adjusts its offered set based on the product availability (offering only those with positive inventories) yields $1 \%-8 \%$ increase in the revenue compared with the non-adaptive LPO algorithm. This also emphasizes the need for a real-time algorithm that adjusts its strategy based on the remaining inventory of products and customers' choices.

[^14]| Problem <br> Class |  | Upper <br> Bound (in \$1000) | Average Revenue Under Different Policies (as \% of the Upper Bound) |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Inventory-Balancing | Myopic <br> Policy | One-shot LP |  | LP Resolving |  | Hybrid |  |
| LF | CV |  |  | EIB | LIB | LPO | ALPO | $\mathrm{LPR}^{500}$ | LPR ${ }^{50}$ | $\mathcal{H}_{1.5}$ | $\mathcal{H}_{2}$ |
| 1.4 | 2.0 |  | 165 | 97.0 | 96.9 | 96.4 | 73.6 | 81.4 | 91.3 | 91.3 | 96.5 | 95.8 |
|  | 1.0 | 169 | 96.8 | 96.9 | 95.8 | 84.3 | 88.4 | 89.8 | 90.0 | 96.6 | 96.5 |
|  | 0.1 | 172 | 97.5 | 97.5 | 96.4 | 92.8 | 93.5 | 90.9 | 90.8 | 96.8 | 96.7 |
| 1.6 | 2.0 | 170 | 97.3 | 97.3 | 96.8 | 66.7 | 73.6 | 88.5 | 88.6 | 96.8 | 95.9 |
|  | 1.0 | 174 | 97.5 | 97.6 | 96.5 | 82.5 | 86.2 | 88.1 | 88.1 | 97.2 | 96.2 |
|  | 0.1 | 178 | 98.4 | 98.5 | 97.3 | 91.8 | 92.4 | 89.5 | 89.7 | 98.3 | 97.7 |
| 1.8 | 2.0 | 174 | 98.0 | 97.9 | 97.2 | 62.9 | 68.8 | 86.9 | 87.0 | 97.9 | 97.2 |
|  | 1.0 | 178 | 98.0 | 98.1 | 97.1 | 78.2 | 81.9 | 86.5 | 86.6 | 98.1 | 97.9 |
|  | 0.1 | 178 | 97.8 | 97.9 | 96.7 | 85.1 | 85.2 | 86.1 | 86.3 | 97.8 | 97.5 |

Table 1.3: Revenue comparison when the length of the horizon is random. The standard errors of all numbers are less than $0.1 \%$.

Explanation of LPR performance: In appendix A.3.1, we show that when the horizon length $T$ - corresponding to the total number of customers - is known in advance, the LPR heuristic performs very well and beats all other policies. We also observe that the LPR revenue increases linearly over time, which shows that the LPR method exploits the known length of the horizon to effectively ration the inventory to all customers. However, when $T$ is random, and the length of the horizon is less than the expectation, the LPR heuristic may end up with rather large left-over inventory (of more profitable products). On the other hand, if $T$ is larger than the expectation, then the LPR heuristic may run out of some the products, and customers would face stock out.

We note that, in all problem classes, Myopic policy performs well because there is only a small chance that a customer will not purchase a product; that is, most of the preference weight parameters of the MNL model are positive for all customer types. Thus, if the Myopic policy runs out of the most profitable products quickly by offering them to customers, then it can still collect high revenues because future customers continue to like the rest of the available products. To demonstrate this point, in Section 1.6.4, we conduct a simulation on synthetic data where each customer type is only interested in a subset of products, and the preference weights are small for most products. We observe that in this setting, our IB algorithms can beat the Myopic policy up to $8 \%$.

Running Time: Although the $\mathrm{LPR}^{50}$ and $\mathrm{LPR}^{500}$ heuristics do not solve the linear program very often, they are still twenty five and four times slower than the Inventory-Balancing algorithms, respectively. We


Figure 1.1: The C.C.D.F. of revenues with $\mathrm{LF}=1.8$ and $\mathrm{CV}=2$. For each algorithm, the curve shows the fractions of problem instances (out of 250 ) whose revenue exceed certain percentages of the upper bound.
note that, in the special case of the multinomial logit choice model, the linear program Primal-S with $O\left(2^{n}\right)$ variables and $O(|\mathcal{Z}|)$ can be reduced to an equivalent LP with just $O(n|\mathcal{Z}|)$ variables and $O(n|\mathcal{Z}|)$ constraints Topaloglu (2013). However, for the general model where this reduction is not possible, the LPR heuristic will be even slower.

Distribution of revenues across problem instances: To get more insight into the performance of different policies, in Figure 1.1, we depict the Complementary Cumulative Distribution Function (C.C.D.F.) of revenue, as percentage of the optimal clairvoyant revenues, across 250 problems instances in the problem class with $\mathrm{LF}=1.8$ and $\mathrm{CV}=2$, for Myopic, LIB, EIB, $\mathrm{LPR}^{500}$, LPO, ALPO, and $\mathcal{H}_{1.5}$ algorithms. In Figure 1.1, the curve for each algorithm shows, for $\alpha \in[90,100]$, the fraction of problem instances whose revenues are at least $\alpha \%$ of the upper bound. It can be seen that both IB polices almost dominate all other policies stochastically. ${ }^{24}$ In addition Myopic and Hybrid policies stochastically dominate the LPR ${ }^{500}$ heuristic, which demonstrates the weak performance of LPR heuristics, in the worse-case sense, when there is uncertainty in the number of customers. We observe that for $80 \%$ of problem instances, EIB and LIB algorithms obtain more than $95.8 \%$ of the optimal clairvoyant revenue. However, the LPR ${ }^{500}$ heuristic gets more than $95.8 \%$ of the upper bound only for about $33 \%$ of all instances.

[^15]
### 1.6.3 Benefits of Personalization: The Importance of Knowing Customer Types

To quantify the benefits of personalization, we compare the revenue in two cases: multiple-type versus single-type. We consider the same 10 locations as in section 1.6.1. The total number of DVDs is 73 , corresponding to the union of top 20 DVDs with the highest sales volumes in each location. The initial inventory of each DVD is set at 30 . The setting for the multiple-type case is exactly the same as Section 1.6.1, where we estimate the parameters $\left(v_{0}^{z}, v_{1}^{z}, \ldots, v_{n}^{z}\right)$ for each type $z$, and customize the decision to the type of each arriving customer. On the other hand, in the single-type case, we estimate a single parameter vector $\left(v_{0}, v_{1}, \ldots, v_{n}\right)$, and use it to make decision for all customers. In our simulation, when we compute the revenue, we assume that each customer makes a decision based on her own type. Since the multipletype model is more accurate, we expect that the multiple-type setting will yield higher revenue, but the key question is the magnitude of the improvements.

| Loading <br> Factor <br> (LF) | Assort- <br> ment <br> Size <br> (C) | $\begin{gathered} \text { Upper } \\ \text { Bound } \\ \text { (in } \$ 100 \text { ) } \\ \hline \end{gathered}$ | Revenue in the <br> Multiple-type Case <br> (as \% of the Upper Bound) |  | Improvements Over the Single-type Case (as \% of the Upper Bound) |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | LIB | EIB | LIB | EIB |
| 1.2 | $\begin{aligned} & 10 \\ & 20 \\ & \hline \end{aligned}$ | 526 | $\begin{array}{r} 82.3 \\ 90.8 \\ \hline \end{array}$ | $\begin{aligned} & 82.5 \\ & 90.7 \\ & \hline \end{aligned}$ | $\begin{aligned} & 17.4 \\ & 11.6 \\ & \hline \end{aligned}$ | $\begin{gathered} 15.4 \\ 9.8 \\ \hline \end{gathered}$ |
| 1.4 | $\begin{aligned} & 10 \\ & 20 \end{aligned}$ | 546 | $\begin{aligned} & 88.0 \\ & 95.9 \\ & \hline \end{aligned}$ | $\begin{aligned} & 88.0 \\ & 95.6 \end{aligned}$ | $\begin{aligned} & 19.3 \\ & 11.1 \\ & \hline \end{aligned}$ | $\begin{aligned} & 17.5 \\ & 10.0 \end{aligned}$ |
| 1.6 | $\begin{aligned} & 10 \\ & 20 \end{aligned}$ | 549 | $\begin{aligned} & 95.3 \\ & 98.9 \\ & \hline \end{aligned}$ | $\begin{aligned} & 94.9 \\ & 98.7 \end{aligned}$ | $\begin{gathered} 21.3 \\ 6.5 \end{gathered}$ | $\begin{gathered} 19.2 \\ 5.6 \end{gathered}$ |

Table 1.4: The average revenue for IB algorithms in multiple-type model and the improvement over the single-type model. All numbers are statistically significant and the standard errors are less than $0.1 \%$.

We impose a constraint $C$ on the size of the assortment that we can offer to each customer, and we consider $C \in\{10,20\}$. Columns 4 and 5 in Table 1.4 show the revenues of the IB algorithms in the multipletype case, for different values of the size constraints. The corresponding difference between the revenues in the multiple-type and single-type case is reported in columns 6 and 7 . We consider the loading factors of $1.2,1.4$, and 1.6 , but set the coefficient of variation at 0.2 because the results for the other coefficients of variations are similar. ${ }^{25}$ As expected, the revenue from the multiple-type case exceeds the single-type,

[^16]regardless of the algorithms we use. However, the benefits of customization is very significant when the assortment size is small. For example, for loading factor of 1.6 and $C=10$, using the LIB algorithm in the multiple-type case yields over $21 \%$ improvement. As the size of the assortment increases, however, the benefits of customization decreases.

### 1.6.4 Comparison to Myopic Policy

Theoretically, it is easy to show that our algorithm would perform significantly better than a Myopic policy. Using the experiment on real data (for choice models) in Section 1.6.2, we also observe improvement of $1 \%$ $-2 \%$ over the Myopic policies. In this section, we run a synthetic simulation in which our algorithm can get up to $7.4 \%$ more revenue compare to Myopic policy. In this simulation there are 73 DVDs and 10 types. As in our simulations, for each DVD $i=1,2, \ldots, 73$, the revenue $r_{i}$ is the average selling price of the DVD obtained from the dataset and the initial inventory level $c_{i}$ is set to 30 . We also assume that customers purchase according to MNL model. However, in this simulation we do not obtain the preference weight parameters, $\left(v_{0}^{z}, v_{1}^{z}, \ldots, v_{n}^{z}\right)$ for each type $z \in \mathcal{Z}$, from the dataset. Rather, we order the products based on their prices so that product 1 has the highest price and product 73 has the lowest price. We assume that for each of the first 9 customer types ( $z \in\{1,2, \ldots, 9\}$ ), the customer type $z$ is only interested in the products indexed by $A^{z}=\{1,2, \ldots, 7 z\}$, with $v_{i}^{z}=1$ for $i \in A^{z}$, and $v_{i}^{z}=0.001$ for $i \notin A^{z}$. Customers of type $z=10$ are interested in all product equally likely. We further assume that $v_{0}^{z}=1$ for all 10 types. The way we construct the preference weight parameters reflects the fact that some customers are only interested in recently released DVDs that are of course more expensive while some customers are less sensitive to this issue. The arrival process is exactly the same as previous sections namely the number of customers of each type is computed by multiplying the total number of customers by a 10 -dimensional Dirichlet random variable. After generating the number of customers of each type, we uniformly permute them to determine the order of arrival.

Table 1.5 compares different policies in this setting when the length of the horizon is drawn from the uniform distribution with $\bar{T}-\underline{T}=\mathrm{E}[T]$ for coefficients of variations 0.5 and 1.0 and loading factor of 1.2, 1.4, and 1.6. We observe that the both EIB and LIB algorithms beat the Myopic policy by $4.5 \%-7.4 \%$. In addition, IB policies surpass all other policies as well.

| Problem <br> Class |  | Upper <br> Bound (in \$100) | Avg. Revenue Under Different Policies (as \% of the Upper Bound) |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Inventory-Balancing | Myopic Policy | One-shot LP |  | $\begin{gathered} \text { LP Resolving } \\ \hline \text { LPR }^{500} \end{gathered}$ |
| LF | CV |  |  | EIB | LIB |  | LPO | ALPO |
| 1.2 | 1.0 |  | 499 | 95.5 | 96.0 | 90.1 | 65.2 | 66.6 | 90.5 |
|  | 0.5 | 516 | 94.9 | 95.5 | 88.1 | 80.8 | 82.2 | 90.8 |
| 1.4 | 1.0 | 514 | 96.1 | 96.6 | 90.8 | 63.9 | 65.3 | 89.8 |
|  | 0.5 | 530 | 95.6 | 96.2 | 89.5 | 79.4 | 80.7 | 90.7 |
| 1.6 | 1.0 | 524 | 96.8 | 97.3 | 92.3 | 71.3 | 72.7 | 91.3 |
|  | 0.5 | 538 | 96.5 | 97.0 | 90.9 | 88.7 | 89.8 | 92.7 |

Table 1.5: Revenue Comparison. The standard errors of all numbers are less than $0.1 \%$.

### 1.7 Conclusion and Future Work

Motivated by the availability of instantaneous data on customer characteristics, we formulated a real-time, personalized, choice-based assortment optimization problem, with arbitrary customer types. Our proposed Inventory-Balancing algorithms are simple, intuitive, and effective. We establish the competitive ratio for our algorithms, and prove that it is the best possible for this problem.

The managerial insight from our work is that companies can increase their revenue by personalizing the operational decisions to each customer, using the real-time information about the customer's characteristics. This process requires coordination between front-end customer-facing decisions with backend supply chain constraints. We demonstrate that such a coordination can be achieved using simple index-based algorithms that can be easily implemented. As the volume and speed in which real-time data become available, the opportunity in this area will continue to grow, and we believe this work can serve as a starting point for more complex models.

Our proposed Inventory-Balancing algorithms maintain an index for each product, in which the marginal revenue is discounted by a penalty based on the product's remaining inventory. The index serves as a simple mechanism for coordinating between the fast-moving customer-facing decision and the back-end operational constraints. In addition to inventory, it would be interesting to consider other supply chain constraints. For example, in the network revenue management, each product corresponds to an itinerary and it uses a common pool of resources, corresponding to seats on flights. When resources are shared among products, coordinating and valuing the benefit of each additional unit of resource becomes challenging. One way to extend our index-based framework is to consider an index that depends on the inventory of all resources simultaneously.

Another exciting direction is to explicitly model the mechanism for learning the choice model of each customer. Recent advances by Farias et al. (2013) have allowed one to model and estimate a very rich class of choice models, and it would be interesting to understand how we can incorporate the learning mechanism with the real-time decision-making process considered in this chapter. Finally, it would be interesting to consider personalization of other operational decisions, such as pricing, warranty services, or shipping options. We believe that the framework in this work can serve as a starting point for analyzing the personalization of these decisions.

## Chapter 2

## Auctions with Dynamic Costly Information

## Acquisition

### 2.1 Introduction

In the traditional auction theory, it is often assumed that bidders have full information about their valuations for the items that are sold at the auction, and the challenge for the auctioneer is to design a mechanism that elicits the preferences and valuations of the bidders. However, this assumption may not always hold, for instance, in the context of online advertising. We explain this below.

The Internet has provided an unprecedented platform for advertisers to reach their consumers and customize their advertisements at an individual level. For instance, major retailers and department stores, such as Macy's or Nordstrom, target users who have previously visited their websites by showing them ads featuring products viewed by the users. These ads usually include a picture of the product(s), often with a promotional price or a discount linked to the retailer's website. This level of targeting has been made possible by using HTTP cookies that allow advertisers to track the users that visit their websites. A cookie contains the information that browsers communicate to the websites visited by the users (see RFC6265 (2011)).

A growing trend in online advertising is the emergence of companies that provide information about Internet users; we will refer to these as information providers. ${ }^{1}$ The information about the users is usually gathered via third-party cookies, which are often installed by a website that a user may not have visited. Using this technology, it is possible to track a user across different websites in order to build a browsing history. The browsing history of the user provides valuable information for advertisers, including the user's interests and intentions. For example, consider a user who has searched a website for flights to Hawaii. Later, when the user visits other websites, the advertisers can follow her and show her deals for vacation

[^17]packages to Hawaii (cf. Helft \& Vega (2010)). Some companies, such as Epsilon and Acxiom, take this one step further by merging the information gathered online with offline marketing databases (cf. Singer (2012)). The information providers, who have a detailed and up-to-date history for each user, sell their data to the advertisers, for instance, via cookie-matching services. The advertisers use this data to update their valuations and target their potential customers.

In many other environments, bidders can obtain additional information about their valuations at a cost. For instance, in the sale of complex financial or business assets, to make a better investment decision, the bidders invest heavily in the due diligence process to uncover the value of the assets (Vallen \& Bullinger 1999). In timber auctions, bidders need to examine the volume and composition of wood on tracts (Roberts \& Sweeting 2010, Athey \& Levin 1999). Similarly, in oil and gas auctions, bidders can conduct seismic studies to better assess the likelihood of finding oil and gas (Hendricks et al. 2003).

In all of the examples above, the bidders incur the cost of information. But, the seller in fact indirectly bears these costs. The intuition is that these costs will affect the bidders' (ex-ante) willingness to pay, and this in turn impacts their participation and bidding. Therefore, the seller should ensure that bidders do not over-invest or under-invest in information acquisition. On one hand, to avoid under-investment, the seller should motivate some of the bidders to obtain information. On the other hand, to avoid over-investment, the seller may would like to restrict some of the bidders from accessing to costly information if he is able to control access to information.

In the aforementioned applications, the seller can exert such a control. For instance, in the sale of complex financial assets, the seller may control how much detail is disclosed to the potential buyers. In timber (oil and gas) auctions, the bidders can only obtain the additional information if the auctioneer allows them to examine the tracts (fields). In the context of online advertising, the seller (online publisher) can release the identity of the user to only a subset of the advertisers. The identity of the user can be revealed to advertisers by disclosing HTTP cookies (Kristol 2001). If the publisher releases these cookies to an advertiser, the advertiser can subsequently take the cookies to the aforementioned information provider and obtain (purchase) information. On the flip side, if the publisher does not disclose these cookies to some advertisers, then these advertisers cannot acquire any information about the user. ${ }^{2}$

[^18]In this work, we answer the following question: If the seller controls access to (additional) information, how can he incentivize a right set of bidders to invest in information acquisition?

## Our Contributions

We present a model to study costly information acquisition in auctions. Our model consists of an auctioneer that sells an item to a set of agents. The agents have an initial private valuation for the item and can obtain additional information at a cost. However, "access" to this additional information is controlled by the auctioneer, and the mechanism may grant access only to a subset (or none) of the agents.

We present a two-stage efficient mechanism in our setting in Section 2.4. A mechanism is efficient if it maximizes the sum of the social welfare of the auctioneer and the agents, taking into account the cost incurred to obtain the additional information. When there is no such cost, the efficient mechanism allows all agents to obtain the additional information. However, when the information is costly, the efficient mechanism grants access to the additional information only to a subset of the agents.

The efficient mechanism works as follows: In the first stage, agents bid in an initial round of bidding. Then, based on their initial bids, the auctioneer selects a subset of agents and grants them access to obtain additional information. Each selected agent can acquire the additional information by incurring a cost. The selected agents then update their valuations and bid in the second round of bidding. The second stage corresponds to the second-price auction with no reserve. ${ }^{3}$

In order to increase the revenue, the seller may want to set a reserve price in the second stage. In Section 2.5.1, we extend our analysis to mechanisms where the item is allocated via a second-price auction with a reserve. These mechanisms are appealing from a practical perspective because they allocate the item using the second-price auctions which is the prevalent mechanism used in ad exchanges (Muthukrishnan 2009).

We further present a revenue-optimal mechanism in our setting. It turns out that the allocation stage of the revenue-optimal mechanism is a bit more complicated than a second-price auction. In order to optimize the revenue, the mechanism selects a set of agents so that it maximizes the "virtual revenue" minus the cost

[^19]of information. The item is allocated via a weighted second-price auction where the weights favor the agents with higher initial bids.

In the above mechanisms, the auctioneer controls access to information. To study the impacts of such a control, in Sections 2.4.1 and 2.5.3, we investigate a mechanism called All-Access, where the auctioneer does not control access to the additional information. In this mechanism, agents obtain additional information if they choose to incur the cost. That is, the mechanism leaves the decision on obtaining the costly information to the agents. The item is allocated via a standard second-price auction (with a reserve). We show that the mechanism always admits a pure strategy Nash equilibrium. The equilibria, however, might not be efficient/revenue-optimal. We observe that under the All-Access mechanism, the agents tend to overinvest in information acquisition when the cost is low. Similarly, when the cost is high, the agents tend to under-invest in information acquisition.

We numerically compare the above mechanisms. Interestingly, on average, the revenue-optimal mechanism allows fewer agents to obtain the additional information compared to the efficient mechanism. It is well established that the revenue-optimal mechanism distorts the allocation and creates inefficiencies in order to extract more revenue from agents with higher valuations. We observe that the revenue-optimal mechanism distorts the revelation of information in addition to the allocation; see Section 2.6. In addition, we observe that the revenue-optimal mechanism can significantly increase the revenue compared with the aforementioned All-Access mechanism; see Section 2.5.3.

Our proposed mechanisms are flexible and can be generalized to a setting where multiple units of the item are sold; see Section 2.7.2. Furthermore, they can be extended to environments where the cost of information can be seen as an entry cost, and all agents who may receive the item with non-zero probability must invest in obtaining information; see Section 2.7.1.

## Related Work

In this section, we briefly discuss the literature related to our work.

Dynamic Mechanism Design: Our work belongs to the growing body of research on mechanism design; see Bergemann \& Said (2011b) for a survey. In particular, our work is closely related to that of Ëso \& Szentes (2007) and generalizes their model to a setting where information acquisition is costly. In the
absence of this cost, Ëso \& Szentes (2007) show that the revenue-optimal mechanism grants all agents access to the additional information. In contrast, we show that when obtaining the additional information is costly, the auctioneer, even in the efficient mechanism, may not allow all bidders to acquire the additional information. The selection rule of our mechanism determines the set of agents who could access (and obtain) the additional information. From a technical perspective, as we discuss later in Appendix B.1.1, this makes our proof a bit challenging. See Kakade et al. (2013), Pavan et al. (2014), Battaglini \& Lamba (2012), Boleslavsky \& Said (2013), and Lobel \& Xiao (2013) for recent results on designing optimal dynamic mechanisms.

Costly Information Acquisition: Most previous work on information acquisition considers settings where the bidders do not have any private information prior to entering the auction. In such a setting, where the auctioneer controls the bidder's access to information, Crémer et al. (2009) show that the auctioneer can extract all the surplus by imposing an admission fee; see also Pancs (2013). Information acquisition has also been studied in the principle-agent context (Cremer \& Khalil 1992, Szalay 2009) and reverse auctions (Beil et al. 2015).

Shi (2012) studies costly information acquisition in a setting where bidders do not have any private information prior to entry and can decide on how much to invest in order to obtain information. He shows that the optimal mechanism takes the form of standard auctions (e.g., second-price) with a reserve price. In contract, in our setting the bidders are privately informed before they decide on the obtaining additional information.

Ye (2007) and Quint \& Hendricks (2012) study indicative bidding auctions that are commonly used in selling financial assets. The auction works as follows: bidders submit non-binding bids to indicate their interest in the assets. The auctioneer then selects some of the bidders that have higher valuations to proceed to the second round, which involves a costly due diligence process and final bidding. Ye (2007) shows that in the indicative bidding, efficient entry of the bidders is not guaranteed and the most qualified bidders might not be selected by the auctioneer. Note that in contrast to our work, the number of selected bidders is predetermined. In addition, only selected bidders who invest in obtaining the additional information may participate in the allocation stage of the mechanism.

One of the closest works in the literature to ours is that by Lu \& Ye (2014), who study the design of a twostage revenue-maximizing mechanism when acquiring information is costly. Similar to indicative bidding
auctions, they assume that obtaining the costly additional information is necessary for agents to participate in the second stage. Under this assumption, as the initial valuations of the agents increase, fewer agents will be allowed to acquire information. Specifically, the selection rule of the mechanism is "monotone" in initial valuations.

In contrast, we observe in Section 2.6 that the selection rule of our proposed mechanisms is nonmonotone. The reason is that in our model all agents-including those who do not update their valuationsparticipate in the second stage and have a chance to receive the item; see Section 7 for details. The nonmonotonicity of the allocation rule makes our proofs more complicated. Furthermore, they assume that the seller can observe who obtains the additional information, which may not be a realistic assumption in many practical contexts. In contrast, in our setting the seller does not observe who obtains the additional information. Therefore, our mechanisms should be designed in a way that all selected agents willingly acquire the costly information. This requirement makes the mechanism design problem more challenging.

In addition, we observe that the seller earns (significantly) higher revenue in our setting. We provide a numerical example in Section 2.7.1. This is quite intuitive; When the cost of information is high as in our setting, the seller can allocate the item even when no agent invests in information. Another reason that the seller can obtain higher revenue in our setting is that he can allocate the item to one of the agents that did not obtain information when the updated valuations were all low.

In Section 2.7.1, we discussed how our mechanisms can be extended to the setting studied by Lu and Ye (2014). In particular, we extend their revenue-maximizing mechanism to a setting where the cost of information is not the same across agents. Note that in Lu and Ye's paper, all agents incur the same cost when they obtain information. In addition, we presented a two-stage efficient mechanism in their setting.

Another related work is that by Hatfield et al. (2015). They focus on efficient mechanisms where bidders can invest in costly information acquisition to determine their valuations. They show that bidders make efficient investment choices when the utility of an agent is equal to his marginal contribution to the social welfare; see also Bergemann \& Välimäki (2002). In contrast, we observer that for the All-Access mechanism, there might exist an equilibrium where agents do not make efficient investment decisions; see Example 2.4.3 in Section 2.4.1.

For a discussion on settings where the computation of valuation is costly, see Larson \& Sandholm (2001).

Information Disclosure: Our work also contributes to the vast literature on information disclosure. In the following, we briefly discuss this literature, focusing on the works motivated by applications in online advertising.

Recently, several papers have studied the effect of sharing cookies and targeting in advertising. Abraham et al. (2011) show that in a common value setting when some advertisers are able to better utilize information obtained from cookies, asymmetry of information can sometimes lead to low revenue in this market; see also Syrgkanis et al. (2013). Several recent papers, such as Ghosh et al. (2007), Rayo \& Segal (2010), Bergemann \& Bonatti (2011), Emek et al. (2012), Hummel \& McAfee (2012), Bergemann \& Bonatti (2013), and Bhawalkar et al. (2014), analyze the role of providing more (targeting) information in the context of online advertising and show that more information may reduce the revenue. Our proposed mechanisms can control the access to information in order to maximize the revenue.

Information disclosure has been studied in other applications. For instance, Jing (2011) studies customer learning for new durable goods. In his model, the seller invests in informing customers before releasing the goods to the market. In addition, see Lewis (2011).

The remainder of this chapter is organized as follows: In Section 2.2, we formally define our model. Direct mechanisms are defined in Section 2.3. We present our efficient mechanism followed by the AllAccess mechanism in Section 2.4. Section 2.5 discusses revenue maximization and presents the revenueoptimal mechanism. We discuss the selection rule of our mechanisms in Section 2.6. Finally, Section 2.7 explores some of the extensions of our mechanisms.

### 2.2 Setting

We consider a setting with a seller of one (indivisible) item and $n$ agents. The initial valuation of each agent $i$ for the item is denoted by $v_{i, 0} \in[\underline{v}, \bar{v}]$, which is drawn independently from distribution $F$, with probability distribution function (p.d.f.) $f$. Distribution $F$ is known to the seller and all the agents. However, $v_{i, 0}$ is known only to agent $i$.

The seller may allow some of the agents to obtain (additional) information about their valuations. Suppose agent $i$ is one of the agents to whom the seller has "granted" access to the additional information. In this
case, agent $i$ may decide to incur cost $c_{i}$ and obtain signal $\delta_{i}$ about his valuation where $\delta_{i}$ is drawn independently (of $v_{i, 0}$ and other agents' second signals) from distribution $G_{i}$. The distributions $G_{i}, i=1,2, \ldots, n$, are publicly known, but the second signals are private information. Without loss of generality, we assume $\mathrm{E}\left[\delta_{i}\right]=0$. Note that if $\mathrm{E}\left[\delta_{i}\right]=\Delta>0$, we can add $\Delta$ to $v_{i, 0}$ and then subtract $\Delta$ from $\delta_{i}$.

If agent $i$ obtains second signal $\delta_{i}$, his updated final valuation, denoted by $v_{i, 1}$, would be equal to $v_{i, 0}+\delta_{i}$. For the agents who did not learn their second signals, either because the seller denied them the access or by their own choice, let $v_{i, 1}=v_{i, 0}$.

As an example, suppose an advertiser values male users at $\$ 0$ and female users at $\$ 6$. Assume that each user has the same chance of being male as of being female. Thus, when the user's gender is unknown, his expected value, that is, his initial valuation, is $\frac{6+0}{2}=3$. By revealing the gender, the valuation of the advertiser will change; with probability $\frac{1}{2}$, his valuation is increased by $\$ 3$, and with probability $\frac{1}{2}$, it is decreased by the same amount. That is, the second signal is either 3 or -3 with equal probability.

Throughout this chapter, we denote the vector of the initial and final valuations of all agents by $v_{0}$ and $v_{1}$. Also, $v_{-i, 0}$ and $v_{-i, 1}$, respectively denote the vector of the initial and final valuations of all agents except for agent $i$.

The agents are risk neutral. The utility of an agent $i$ who has received the item is equal to his valuation, $v_{i, 1}$, minus his payment to the mechanism and the (possible) cost of information acquisition. We will specify utility of the agents more precisely in the next section.

### 2.3 Direct Mechanisms

In this section, we consider direct revelation mechanisms (Myerson 1986) where agents report their valuations in two rounds. First, they report their initial valuations to the mechanism. Then the mechanism decides on the set of agents that will have access to information. Those agents report their updated valuations to the mechanism in the second round and finally, the mechanism allocates the item.

More precisely, any direct mechanism $\mathcal{M}$ is defined by a tuple ( $s, q, p$ ), which respectively represents its selection, allocation, and payment rules. The seller announces the mechanism to the agents and commits to $(s, q, p)$. Following are the stages of the mechanism:

1. Initial Bidding: Agents report in the first round. The initial report (bid) of agent $i$ is denoted by $b_{i, 0}$. Throughout this chapter, we will use "reporting" and "bidding" interchangeably.
2. Selection: Based on the initial reports, the mechanism selects a set of agents that we call selected agents. The mechanism grants access to each selected agent $i$ to acquire additional information (signal $\delta_{i}$ ) and in return charges them price $t_{i}$. More precisely, selection rule $s: \mathbb{R}^{n} \rightarrow(\{0,1\} \times \mathbb{R})^{n}$ maps the initial bids to a pair $\left(s_{i}, t_{i}\right)$ for each agent $i$. If agent $i$ is selected, $s_{i}\left(b_{0}\right)$ is equal to 1 . Otherwise, $s_{i}$ is equal to 0 . Each selected agent pays amount $t_{i}$ to the mechanism to access his signal (the agent would still need to incur an additional cost $c_{i}$ to learn the signal). To simplify the presentation, we assume $t_{i}$ is equal to 0 for non-selected agents. ${ }^{4}$
3. Obtaining Information: Each selected agent $i$ decides on whether to incur cost $c_{i}$ and learn $\delta_{i}$. We define $e_{i}$ to denote the decision variable for agent $i$ for incurring $\operatorname{cost} c_{i}$ and updating his valuation. $e_{i}$ is equal to 1 if the selected agent $i$ learns $\delta_{i}$. For non-selected agents, $e_{i}$ is defined to be 0 . Neither the mechanism nor other agents can observe decision $e_{i}$ of an agent $i$; it is only known to that agent.
4. Final Bidding: In the final round of reporting, only selected agents get a chance to update their reports. For any selected agent $i, b_{i, 1}$ denotes the updated (and final) report of agent $i$. For all other agents, let $b_{j, 1}=b_{j, 0}$.
5. Allocation and Payments: Based on the initial and final bids of all agents (both selected and unselected), the seller decides to whom to allocate the item, allocation rule $q:(\mathbb{R} \times \mathbb{R})^{n} \rightarrow \mathbb{R}^{+}$, and how much to charge each agent, payment rule $p:(\mathbb{R} \times \mathbb{R})^{n} \rightarrow \mathbb{R}$. Namely, given all the bids and decision variables, $q_{i}\left(b_{0}, b_{1}\right)$ is the allocation probability, and $p_{i}\left(b_{0}, b_{1}\right)$ is the payment of agent $i$.

The utility of agent $i$ participating in the mechanism is equal to $q_{i} v_{i, 1}-p_{i}-t_{i}-e_{i} c_{i}$ (more precisely, $\left.q_{i}\left(b_{0}, b_{1}\right) v_{i, 1}-p_{i}\left(b_{0}, b_{1}\right)-t_{i}\left(b_{0}\right)-e_{i} c_{i}\right)$. Each agent chooses a best-response strategy to deal with the mechanism and strategies of the other agents in order to maximize his expected utility, where the expectation is taken with respect to the second signals of the (selected) agents. More formally, the best response strategy of each agent $i$ can be described with the following mappings: $b_{i, 0}: \mathbb{R} \rightarrow \mathbb{R}, e_{i}: \mathbb{R}^{3} \rightarrow\{0,1\}$, and $b_{i, 1}: \mathbb{R}^{6} \rightarrow \mathbb{R}$. With slight abuse of notation, we denote the decision variables and functions with the

[^20]same notation. Function $b_{i, 0}$ maps $v_{i, 0}$, the initial valuation of the agent, to the bid in the first round $b_{i, 0}$. $e_{i}$ is a function of the initial valuation, $v_{i, 0}$, and initial bid, $b_{i, 0}$, of the agent and his payment $t_{i}$. Finally, $b_{i, 1}$ is a function of the whole history of agent $i$ (i.e., $<v_{i, 0}, b_{i, 0}, s_{i}, t_{i}, e_{i}, v_{i, 1}>$ ) and determines the final bid. Given the strategy of the other agents, agent $i$ optimizes over tuple $\left(b_{i, 0}, e_{i}, b_{i, 1}\right)$ to obtain his best (utility-maximizing) strategy.

The truthful strategy for agent $i$ consists of i) reporting truthfully in the first round ( $b_{i, 0}=v_{i, 0}$; ; ii) obtaining additional information if selected ( $e_{i}=1$ if $s_{i}=1$ ); and iii) reporting truthfully in the final round $\left(b_{i, 1}=v_{i, 1}\right)$.

A dynamic direct mechanism is incentive compatible (IC) if for every agent and every truthful history, truth-telling is a best response given that all other agents report truthfully. More precisely, a mechanism is IC if
$\mathrm{E}_{\text {TRUTHFUL }}\left[q_{i}\left(v_{0}, v_{1}\right) v_{i, 1}-p_{i}\left(v_{0}, v_{1}\right)-t_{i}\left(v_{0}\right)-e_{i} c_{i}\right]$
$=\max _{b_{i, 0}, e_{i}, b_{i, 1}}\left\{\mathrm{E}\left[q_{i}\left(\left(b_{i, 0}, v_{-i, 0}\right),\left(b_{i, 1}, v_{-i, 1}\right)\right) v_{i, 1}-p_{i}\left(\left(b_{i, 0}, v_{-i, 0}\right),\left(b_{i, 1}, v_{-i, 1}\right)\right)-t_{i}\left(\left(b_{i, 0}, v_{-i, 0}\right)\right)-e_{i} c_{i}\right]\right\}$,
where the expectations are taken assuming other agents are truthful; that is, agents report truthfully in both rounds, i.e., $b_{-i, 0}=v_{-i, 0}$ and $b_{-i, 1}=v_{-i, 1}$, and obtain information if selected, that is, $e_{j}=1$ if $s_{j}=1$ for $j \neq i$. In addition, in the l.h.s., the expectation is taken under the truthful strategy of agent $i$.

We show that our proposed mechanisms satisfy stronger incentive compatibility properties. Namely, selected agents always bid truthfully in the final round even if they have deviated from the truthful strategy in the past. In addition, each selected agent prefers to obtain additional information even if they observe other agents' initial valuations. Currently, we assume that the agent only observes $s_{i}$ and $t_{i}$, which may contain information about other agents' valuations.

We can now define the participation constraints for the mechanism. An IC mechanism is individually rational (IR) if for each agent $i$, the expected utility under the truthful strategy is non-negative, that is,

$$
\mathrm{E}_{\text {TRUTHFUL }}\left[q_{i}\left(v_{0}, v_{1}\right) v_{i, 1}-p_{i}\left(v_{0}, v_{1}\right)-t_{i}\left(v_{0}\right)-e_{i} c_{i}\right] \geq 0 .
$$

### 2.4 The Efficient Mechanism

The social welfare of a mechanism is defined as the sum of the utility of the agents and the seller minus the cost incurred to obtain additional information. Note that an IC and IR mechanism is efficient if it obtains the maximum social welfare equal to:

$$
\max _{S \subseteq\{1, \cdots, n\}}\left\{\mathrm{E}_{S}\left[\sum_{i=1}^{n} q_{i} v_{i, 1}\right]-\sum_{i \in S} c_{i}\right\}=\max _{S \subseteq\{1, \cdots, n\}}\left\{\Omega\left(v_{0}, S\right)\right\},
$$

where $\mathrm{E}_{S}$ denotes the expectation with respect to the realizations of the second signals when all agents in set $S$ (and only those agents) obtain the additional information. In addition, $\Omega\left(v_{0}, S\right)$, defined below, denotes the maximum social welfare that can be obtained when set $S$ of agents acquires information.

$$
\begin{equation*}
\Omega\left(v_{0}, S\right)=\mathrm{E}_{S}\left[\max \left\{\max _{j \in S}\left\{v_{j, 0}+\delta_{j}\right\}, \max _{j \notin S}\left\{v_{j, 0}\right\}, 0\right\}\right]-\sum_{j \in S} c_{j} . \tag{2.1}
\end{equation*}
$$

To gain intuition, let us consider the following scenario. Assume all the agents bid truthfully, agents in set $S$ obtain information, and subsequently, each agent $j$ who updates his valuation incurs cost $c_{j}$. The total cost of information is equal to $\sum_{j \in S} c_{j}$. The mechanism maximizes the social welfare by allocating the item to the agent with the highest non-negative final bid, that is, agent $i^{\star} \in$ $\arg \max \left\{\max _{j \in S}\left\{v_{j, 0}+\delta_{j}\right\}, \max _{j \notin S}\left\{v_{j, 0}\right\}\right\}$, where $\max _{j \in S}\left\{v_{j, 0}+\delta_{j}\right\}$ and $\max _{j \notin S}\left\{v_{j, 0}\right\}$ are, respectively, the maximum updated bids of agents who obtain information and the maximum bid of agents who do not. The item is allocated to $i^{\star}$ if his (final) bid is positive; otherwise, there would be no allocation.

Throughout this chapter, we assume that the maximum expected social welfare that can be obtained is bounded; that is, $\mathrm{E}\left[\max _{S \subseteq\{1,2, \ldots, n\}}\left\{\Omega\left(v_{0}, S\right)\right\}\right]<\infty$, where the expectation is with respect to initial valuations $v_{0}$.

We now present an efficient mechanism.
$\mathcal{M}^{\text {EFF }}$ Mechanism: The selection, allocation, and payment rules are defined as follows:

- Selection: Select a set of agents such that granting them access to information maximizes the social welfare of the seller and agents, taking into account the cost of information. Specifically, select the
following set of agents $\mathcal{S}\left(b_{0}\right)=\arg \max _{S \subseteq\{1, \cdots, n\}}\left\{\Omega\left(b_{0}, S\right)\right\}$, where $\Omega\left(b_{0}, S\right)$ is defined in Eq. (2.1).

The payment of selected agent $i$ is equal to

$$
\begin{equation*}
t_{i}\left(b_{0}\right)=-c_{i}+\mathrm{E}\left[q_{i} v_{i, 1}-p_{i} \mid v_{0}=b_{0}\right]-\int_{\underline{v}}^{b_{i, 0}} \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=b_{-i, 0}\right] d z, \tag{2.2}
\end{equation*}
$$

where the expectations are with respect to the second signals. Notation $\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right]$ denotes the expected probability of the allocation of agent $i$, where the initial valuations of agent $i$ and other agents are, respectively, equal to $z$ and $x_{-i}$, assuming all the agents, including agent $i$, are truthful. Note that the first term in the payment implies that the seller subsidizes the cost of information, $c_{i}$, for each selected agent $i$. In lemmas B.1.6 and B.1.9 and their proofs, we discuss how this payment is calculated using the Envelope Theorem and show that it incentivizes the agents to be truthful.

Recall that for non-selected agents, $t_{i}$ is equal to 0 . See Appendix B.5.1 for examples that depict initial payments.

- Allocation and Payments: Allocate the item to the agent with the highest non-negative bid at a price equal to the second highest bid or a reserve. More precisely, consider an agent $i^{\star} \in \operatorname{argmax}_{i}\left\{b_{i, 1}\right\}$. If $b_{i^{\star}, 1} \geq 0$, agent $i^{\star}$ receives the item and pays $p_{i^{\star}}=\max \left\{\max _{i \neq i^{\star}}\left\{b_{i, 1}\right\}, r\right\}$, where $r: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a function of initial bids and will be defined later in Eq. (2.3).

Let agent $\ell \in \arg \max _{j \notin \mathcal{S}\left(b_{0}\right)}\left\{b_{j, 0}\right\}$ be an unselected agent with the highest initial bid. The reserve price $r$ is simply zero when $b_{\ell, 0}<0$ or all agents are selected; otherwise $r$ is the solution of the equation below

$$
\begin{equation*}
\int_{\max \{r, 0\}}^{b_{\ell, 0}} \operatorname{Pr}\left[z \geq \max _{j \in \mathcal{S}\left(b_{0}\right)}\left\{b_{j, 0}+\delta_{j}\right\}\right] d z=\int_{\underline{v}}^{b_{\ell, 0}} \mathrm{E}\left[q_{\ell} \mid v_{\ell, 0}=z, v_{-\ell, 0}=b_{-\ell, 0}\right] d z \tag{2.3}
\end{equation*}
$$

Lemma B.1.1 and Corollary B.1.2 in Online Appendix, Section B.1, show that there exists an $r \in$ $\left[0, b_{\ell, 0}\right]$ that satisfies the above equation. If there are multiple solutions to the above equation, we choose the largest one.

Observe that because $r \in\left[0, b_{\ell, 0}\right]$, the reserve price does not change the allocation or the payment of the selected agents. Specifically, if a selected agent $i$ wins the item, he pays $\max \left\{\max _{j \neq i}\left\{b_{j, 1}\right\}, r\right\}$,
which is identical to $\max \left\{\max _{j \neq i}\left\{b_{j, 1}\right\}, 0\right\}$. Then, one can describe the payment of the mechanism as follows: If agent $i^{\star}$ was a selected agent, then he pays $\max \left\{\max _{j \neq i^{\star}}\left\{b_{j, 1}\right\}, 0\right\}$. If $i^{\star}$ was not a selected agent, then he pays max $\left\{\max _{j \neq i \star}\left\{b_{j, 1}\right\}, r\right\}$.

In fact, by introducing the reserve price $r$, the mechanism charges agent $\ell$ differently in the second round to incentivize him to bid truthfully in the first round. Note that the initial bid of agent $\ell$ can change the set of selected agents. In addition, similar to all selected agents, agent $\ell$ has a chance to win the item. However, unlike the selected agents, agent $\ell$ was not charged in the first round.

We now present the main result of this section.
Theorem 2.4.1 (Efficient Mechanism). Mechanism $\mathcal{M}^{\mathrm{EFF}}$ is individually rational, incentive compatible, and efficient.

From its construction, it is not difficult to see that if mechanism $\mathcal{M}^{\text {EFF }}$ is IC, then it is also efficient. We prove the incentive compatibility of the mechanism in Appendix B.1. Observe that selected agents bid truthfully in the second round because the item is allocated using a second-price auction. We then show that any selected agent that bids truthfully in the first round obtains information. To this aim, we show that the mechanism's selection rule aligns with the agent's incentive. Specifically, the marginal change in the utility of a selected agent from not obtaining information is equal to the change in the social welfare. The challenging part of the proof is showing that agents bid truthfully in the first round because agents' bids in the first round determine the set of selected agents endogenously.

The selection rule ensures that a right set of agents invest in information acquisition. Note that as more agents obtain information, there is a higher chance that an agent has a high valuation for the item, which could increase the social welfare. On the other hand, the increase in the highest valuation comes at the cost of information acquisition. Thus, there is a trade-off here, and the selection rule aims to avoid over or underinvestment in information acquisition. In fact, as we show in Section 2.4.1, over or under-investment in information acquisition may not be avoided if the seller cannot control access to information via a selection rule.

We also note that the selection rule of the efficient mechanism chooses a set of agents and allows them to update their valuations simultaneously. Alternatively, one can consider an "adaptive" selection rule that discloses information step-by-step (cf. McAfee \& McMillan (1988)). During the selection stage, at each
step, the mechanism selects one of the agents to obtain information. Then, based on the report of that agent, the mechanism makes a decision on obtaining more information or proceeding to the allocation stage. In this work, we consider a two-stage information disclosure, which could be more appealing from a practical perspective. The sequential search can be time-consuming and complex. For instance, in the example from online advertising, the mechanism should be executed in milliseconds, and sequential information acquisition may not be feasible.

### 2.4.1 What If All Agents Are Allowed to Access the Additional Information?

Intuitively, the ability of the seller to control access to the additional information could impact the social welfare and the utilities of the agents. To formalize this intuition, in this section, we present a mechanism called "All-Access." In this mechanism, the seller allows all agents to obtain information, if they wish, i.e., $s_{i}=1$ and $t_{i}=0$ for $i=1,2, \ldots, n$. We observe that when the seller leaves the decision on acquiring costly information to the agents, the agents' decisions can create inefficiency. Without the seller's control, the agents may over-invest or under-invest in information acquisition. In addition, we observe that the agents tend to invest in information acquisition when their initial valuations are not too high or too low. This is in contrast with the efficient mechanism that incentivizes the agents with higher initial valuations to invest in information acquisition. We provide examples in Section 2.6.

The All- Access mechanism works as follows: First, agents observe their initial valuations. Next, each agent $i$ makes a decision on incurring cost $c_{i}$ and obtaining his second signal. The item is allocated via a second-price auction with no reserve price. Similar to our original setting, the investment decision of an agent is only known to him. In addition, initial valuations and second signals are private information and only their distributions are publicly known.

In this mechanism, agents bid only once after they decide on refining their valuations. To be consistence with the notation of our original setting, we denote the bid of an agent $i$ in the second-price auction by $b_{i, 1}$; the initial bid of an agent $i, b_{i, 0}$, is set to be zero. The item is allocated to an agent with the highest non-negative bid, that is, $i^{\star} \in \arg \max _{i}\left\{b_{i, 1}\right\}$ if $b_{i^{\star}, 1} \geq 0$; in case of ties, the item is allocated at random to one of the agents. Agent $i^{\star}$ pays the second highest bid $p_{i^{\star}}=\max \left\{\max _{i \neq i^{\star}}\left\{b_{i, 1}\right\}, 0\right\}$. If agent $i^{\star}$ has obtained additional information, his utility will be equal to $\left(v_{i^{\star}, 0}+\delta_{i^{\star}}\right)-p_{i^{\star}}-c_{i^{\star}}$. Otherwise, his utility is equal to $v_{i^{\star}, 0}-p_{i^{\star}}$. For any agent $i \neq i^{\star}$ that does not receive the item, $p_{i}=0$.

We now consider the Nash Equilibrium (NE) of the All-Access mechanism. We assume that all the agents bid truthfully $\left(b_{i, 1}=v_{i, 1}\right)$ in the auction because bidding truthfully is a weakly dominant strategy for any agent $i$, independent of his and other agents' decisions on obtaining information. Therefore, to characterize the Nash Equilibrium, we focus on the decision of the agents on information acquisition. We define $\tilde{e}_{i}\left(v_{i, 0}\right)=e_{i}\left(v_{i, 0}, b_{i, 0}=0, t_{i}=0\right)$ as the investment strategy of agent $i$ with initial valuation $v_{i, 0}$ in the All-Access mechanism.

The next theorem shows that the All-Access mechanism always admits a pure strategy Nash equilibrium (NE).

Theorem 2.4.2 (Equilibria of the All-Access Mechanism). The All-Access mechanism induces a game of incomplete information among agents where strategy of the agents are defined by $\tilde{e}=<\tilde{e}_{1}, \tilde{e}_{2}, \ldots, \tilde{e}_{n}>$. For this game, there exists a pure strategy Nash equilibrium such that no agent $i$ can gain by changing his investment strategy $\tilde{e}_{i}(\cdot)$ if the investment strategies of the other agents remain unchanged.

All the proofs of this section are provided in Appendix B.2.
To gain insight into the All-Access mechanism, in the rest of this section, we consider the following example.

Example 2.4.3. Assume that there are two agents that participate in the All-Access mechanism with no reserve. The cost of obtaining second signals for both agents is the same $c_{1}=c_{2}$. The initial valuation of agents is drawn from a uniform distribution over $[0,1]$, i.e., $\operatorname{Uniform}(0,1)$, and the second signals are drawn from $\operatorname{Uniform}(-1,1)$.

Although the setting in Example 2.4 .3 is seemingly simple, it highlights challenges in characterizing equilibrium of the All-Access mechanism. To characterize equilibrium of a game of incomplete information, the single crossing conditions are often used (Athey 2001). However, for the setting in Example 2.4.3, we show that the single crossing conditions do not hold; see Proposition B.3.1 in Appendix B.3.2. In the light of this observation, in the next theorem, we present the equilibria of the All-Access mechanism for a wide range of the cost.

Theorem 2.4.4. Consider the All-Access mechanism with no reserve price and the setting in Example 2.4.3. Then,

- when cost $c \leq \frac{7}{96}$, there exists an equilibrium in which both agents always obtain the additional information, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=1$ for $i=1,2$ and $v_{i, 0} \in[0,1]$, and
- when cost $c \geq \frac{7}{48}$, there exists an equilibrium in which none of the agents obtain the additional information, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=0$, for $i=1,2$ and $v_{i, 0} \in[0,1]$.

Theorem 2.4.4 characterizes the equilibrium of the All-Access mechanism when the cost is less than $\frac{7}{96} \approx 0.072$ and greater than $\frac{7}{48} \approx 0.145$. In order to compare the All-Access and the efficient mechanisms, we numerically (using an iterative procedure) find the equilibrium of the All-Access mechanism when the cost is within $\left(\frac{7}{96}, \frac{7}{48}\right)$; see Online Appendix, Section B.3.1, for details.

Our numerical studies show that there exist equilibria in which agents follow interval-based investment decisions. An equilibrium with interval-based investment decisions can be defined with four parameters $a_{1}, a_{2}, \beta_{1}$, and $\beta_{2}$ where $a_{1} \leq \beta_{1} \in[0,1]$ and $a_{2} \leq \beta_{2} \in[0,1]$. In the equilibrium, each agent $i$ obtains information only when his initial valuation $v_{i, 0}$ lies in $\left[a_{i}, \beta_{i}\right]$. We note that the range $\left[a_{1}, \beta_{1}\right]$ is not necessarily equal to $\left[a_{2}, \beta_{2}\right]$. That is, there exist equilibria in which the investment decisions of the agents are not symmetric. ${ }^{5}$

Under an interval investment strategy, agents do not obtain information when their initial valuations are too low or too high. The intuition is that an agent with high initial valuation does not have incentive to invest in information acquisition as he already has a high chance of winning the item without incurring the cost of information. On the other hand, an agent with low initial valuation is not willing to acquire costly information because he has a slim chance of winning the item.

Figures 2.1 and 2.2 compare the All-Access mechanism with the efficient mechanism. Observe that the equilibria of the All-Access mechanism may not be efficient as agents in the All-Access mechanism tend to over-invest in information acquisition when the cost is low, and they tend to under-invest when the cost is high. Yet another source of inefficiency comes from that fact that agents follow interval-based investment strategies and as a result, they may not acquire information when their initial valuations are high. This is in contrast with the efficient mechanism that motivates agents with higher initial valuations to acquire information; see Section 2.6.

[^21]
### 2.5 Maximizing Revenue

In the previous sections, we presented a welfare-maximizing (efficient) mechanism. Our goal here is to design a revenue-optimal mechanism. We start with a heuristic called "Sequential Second-Price (SSP) Mechanism." In this mechanism, the allocation and payments are determined via a second-price auction that makes the mechanism appealing from a practical perspective. Furthermore, this class of mechanisms is motivated in part by the structure of the efficient mechanism.

Despite the desirable properties of the SSP mechanism, it is not able to achieve the maximum revenue. Therefore, in Section 2.5.2, we present a revenue-optimal mechanism. The allocation rule of this mechanism favors agents with higher initial valuations and extracts more revenue from those agents in the first round. Both these mechanisms control agents' access to the additional information. In Section 2.5.3, we investigate the impacts of such a control by re-visiting the All-Access mechanism.

### 2.5.1 Sequential Second-Price Mechanisms

The second-price auctions and their variations are prevalent in online advertising and are used by Google and other major platforms. In this section, we present a class of mechanisms, called Sequential Second-Price (SSP), which extends the second-price auction to our setting with dynamic information acquisition.

The SSP mechanisms are similar to the efficient mechanism. The main difference is that the SSP mechanism sets a lower bound (reserve price) $r$ on the final bid of the agents. That is, it allocates the item to the


Figure 2.1: The average number of agents with the updated valuations versus the cost in the efficient and All-Access (with no reserve) mechanisms with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$.
agent with the highest bid as long as his bid is greater than or equal to $r$. In fact, the mechanism selects agents in set $\mathcal{S}_{r}$, where

$$
\mathcal{S}_{r}\left(b_{0}\right) \in \underset{S \subseteq\{1, \cdots, n\}}{\arg \max }\left\{\Omega_{r}\left(b_{0}, S\right)\right\}
$$

Here,

$$
\Omega_{r}\left(b_{0}, S\right)=\underset{S \subseteq\{1, \cdots, n\}}{\arg \max }\left\{\mathrm{E}_{S}\left[\max \left\{\max _{i \in S}\left\{b_{i, 0}+\delta_{i}\right\}, \max _{i \notin S}\left\{b_{i, 0}\right\}, r\right\}\right]-\sum_{i \in S} c_{i}\right\}
$$

Each selected agent $i$ pays $t_{i}\left(b_{0}\right)$ in the first round, where $t_{i}\left(b_{0}\right)$ is given in Eq. (2.2).

The parameter $r$ in the SSP mechanism can be optimized to maximize the revenue of the seller. In Appendix B.5.2, we compare the revenue of the SSP mechanism with the optimal reserve and the revenueoptimal mechanism. In our examples, the SSP mechanism yields more than $84 \%$ of the optimal revenue.

In practice, the SSP mechanism can be implemented via private auctions (Google AdX Documentation 2015) and pre-negotiated contracts that grant advertisers access to additional information in advance and advertisers bid for the impressions over time. For instance, consider a set of advertisers that are willing to display their ads on a specific website over a period of time. Their initial valuations, which depend on the contents of the website and advertisers' products and services, remain constant over time. However, advertisers' final valuation may vary over time because it also depends on the demographic or behavioral attributes of the user(s). In this case, there is no need for advertisers to report their initial valuations for every impression; they only need to report their updated valuations.


Figure 2.2: The average social welfare versus the cost in the efficient and All-Access (with no reserve) mechanisms with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$.

### 2.5.2 Revenue-Optimal Mechanism

For revenue maximization, without loss of generality, using the revelation principle (cf. Myerson (1986)), we focus on IC and IR mechanisms.

Definition 3 (Optimality). An incentive compatible and individually rational mechanism is optimal if it maximizes the revenue, equal to $E\left[\sum_{i=1}^{n}\left(t_{i}+p_{i}\right)\right]$, among all incentive compatible and individually rational mechanisms. ${ }^{6}$

Let $\alpha\left(v_{i, 0}\right)=-\frac{\left(1-F\left(v_{i, 0}\right)\right)}{f\left(v_{i, 0}\right)}$ be the negative of the inverse hazard rate associated with distribution $F$. We make the following assumption about $\alpha(\cdot)$.

Assumption 2 (Monotone Hazard Rate). Distribution F, with p.d.f. $f$, has a monotone hazard rate; that is, $\alpha(\cdot)$ is non-decreasing in $v_{i, 0}$. Furthermore, assume that $\alpha(\cdot)$ is differentiable and $\sup _{v_{i, 0} \in[\underline{v}, \bar{v}]}\left\{\alpha^{\prime}\left(v_{i, 0}\right)\right\}<$ $\infty$.

The above assumption is standard in the revenue-optimal mechanism design and ensures that the virtual valuations of the agents are increasing in their initial valuations (Myerson 1981). We now present a revenueoptimal mechanism.
$\mathcal{M}^{\text {Opt }}$ Mechanism: The selection, allocation, and payment rules are defined as follows:

- Selection: Select the following set of agents:
$\mathcal{S}_{\text {OPT }}\left(b_{0}\right) \in \underset{S \subseteq\{1,2, \ldots, n\}}{\arg \max }\left\{\mathrm{E}_{S}\left[\max \left\{\max _{i \in S}\left\{b_{i, 0}+\alpha\left(b_{i, 0}\right)+\delta_{i}\right\}, \max _{i \notin S}\left\{b_{i, 0}+\alpha\left(b_{i, 0}\right)\right\}, 0\right\}\right]-\sum_{i \in S} c_{i}\right\}$,
and $t_{i}$ is defined the same as before; see Eq. (2.2).
- Allocation and Payments: Allocate the item to the agent with the highest non-negative weighted bid. More precisely, consider an agent $i^{\star} \in \operatorname{argmax}_{i}\left\{b_{i, 1}+\alpha\left(b_{i, 0}\right)\right\}$. If $b_{i^{\star}, 1}+\alpha\left(b_{i^{\star}, 0}\right) \geq 0$, then the item is allocated to agent $i^{\star}$ at price $p_{i^{\star}}=\max \left\{\max _{i \neq i^{\star}}\left\{b_{i, 1}+\alpha\left(b_{i, 0}\right)\right\}, r\right\}-\alpha\left(b_{i^{\star}, 0}\right)$ where $r: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a function of initial bids and is defined below.

[^22]With slight abuse of notations, let $\ell \in \arg \max _{j \in \mathcal{S}_{\text {ort }}\left(b_{0}\right)}\left\{b_{j, 0}+\alpha\left(b_{j, 0}\right)\right\}$ be an unselected agent with the highest weighted bid. Then, if $b_{\ell, 0}+\alpha\left(b_{\ell, 0}\right)<0$ or all agents are selected, $r=0$. Otherwise, $r$ solves the following equation

$$
\begin{equation*}
\int_{\max \{r, 0\}}^{b_{\ell, 0}+\alpha\left(b_{\ell, 0}\right)} \operatorname{Pr}\left[z \geq \max _{j \in \mathcal{S}_{\text {Opr }}\left(b_{0}\right)}\left\{b_{j, 0}+\delta_{j}+\alpha\left(b_{j, 0}\right)\right\}\right] d z=\int_{\underline{v}}^{b_{\ell, 0}} \mathrm{E}\left[q_{\ell} \mid v_{\ell, 0}=z, v_{-\ell, 0}=b_{-\ell, 0}\right] d z \tag{2.5}
\end{equation*}
$$

Lemma B.1.1 and Corollary B.1.3 in Appendix B. 1 show that there exists an $r \in\left[0, b_{\ell, 0}+\alpha\left(b_{\ell, 0}\right)\right]$ that solves the above equation. Note that similar to the efficient mechanism, $r$ does not change the allocation or payment for selected agents.

Mechanism $\mathcal{M}^{\text {Opt }}$ is built upon the ideas of virtual value formulation of Myerson (1981) for static revenue-maximizing mechanism design. It allocates the item to the agent that has the highest (final) virtual valuation $v_{i, 1}+\alpha\left(v_{i, 0}\right)$. The mechanism maximizes the virtual value of the winner minus the cost of information acquisition.

The following theorem establishes the optimality of $\mathcal{M}^{\text {Opr }}$.
Theorem 2.5.1 (Revenue-Optimal Mechanism). Suppose Assumption 2 holds. Mechanism $\mathcal{M}^{\text {Opt }}$ described above is incentive compatible, individually rational, and optimal.

In Appendix B.1, we show that mechanism $\mathcal{M}^{\text {OPT }}$ is IC and IR. To complete the proof of Theorem 2.5.1, in the following we show that mechanism $\mathcal{M}^{\text {Opt }}$ is revenue-optimal. First, via the following lemma, we show that the selection rule aims to optimize the revenue by maximizing the expected "virtual revenue."

Lemma 2.5.2 (Revenue of $\left.\mathcal{M}^{\text {Opr }}\right)$. If all the agents follow the truthful strategy, then the expected revenue of $\mathcal{M}^{\text {opt }}$ is equal to

$$
\begin{equation*}
\mathrm{E}\left[\max _{S \subseteq\{1,2, \ldots, n\}}\left\{\mathrm{E}_{S}\left[\max \left\{\max _{i \in S}\left\{v_{i, 0}+\alpha\left(v_{i, 0}\right)+\delta_{i}\right\}, \max _{i \notin S}\left\{v_{i, 0}+\alpha\left(v_{i, 0}\right)\right\}, 0\right\}\right]-\sum_{i \in S} c_{i}\right\}\right], \tag{2.6}
\end{equation*}
$$

where the inner expectation is with respect to the second signals and the outer expectation is with respect to the initial valuations.

The proof is given in Appendix B.4.1. We now provide an upper bound on the revenue of any IC mechanisms that matches the revenue of mechanism $\mathcal{M}^{\text {Opt }}$.

Lemma 2.5.3 (Upper Bound). The expected revenue of the seller is at most equal to Eq. (2.6).

The upper bound is established using a closely related problem with fewer constraints called the relaxed problem, cf. Ëso \& Szentes (2007), Kakade et al. (2013), and Pavan et al. (2014). In the relaxed problem, the mechanism, on the behalf of an agent, can decide to obtain information, and then both the agent and the mechanism learn his second signal. Because any mechanism that is IC in the original setting would also be IC in the relaxed setting, the revenue of the optimal relaxed mechanism provides an upper bound for the revenue of the revenue-optimal mechanism in the original setting; see Appendix B.4.1 for details.

The revenue-optimal mechanism is a generalization of the handicap mechanism of Ëso \& Szentes (2007) and matches the mechanism when information acquisition is costless (i.e., $c_{i}=0$ ) where the seller allows all agents access to information. We show that when the information is costly, the seller grants access to additional information only to a subset of agents. The intuition is that the seller indirectly bears the cost of information because the costs will affect the agents' (ex-ante) willingness to pay.

An important distinction between our mechanism and the handicap mechanism is the selection rule that grants access to additional information to a right set of agents. A technical challenge that we need to address is that the selection rule depends on the initial bids which, as we discuss in Section 2.6 is "non-monotone" in the initial valuations of the agents; also see Lemma B.1.9 for more details.

### 2.5.3 What If All Agents Are Allowed to Access to the Additional Information?

Here, we re-visit the All-Access mechanism to study the impacts of controlling access to the additional information on the revenue-optimal mechanism. Specifically, we will introduce a reserve price $r$ to the AllAccess mechanism to ensure that the seller has a degree of freedom to extract more revenue from the agents. In this mechanism, privately informed agents decide on updating their valuations and then participate in a second price auction with reserve price $r$. One can easily extend Theorem 2.4.2 to show that the All-Access mechanism with a reserve price always admits a pure strategy Nash equilibrium.

In the following, we will re-examine Example 2.4.3 when agents participate in the All-Access mechanism with reserve price $r$.

Example 2.5.4. Assume that there are two agents that participate in the All-Access mechanism with reserve price $r \geq 0$. The cost of obtaining second signals for both agents is the same, $\delta_{i} \sim \operatorname{Uniform}(-1,1)$, and $v_{i, 0} \sim \operatorname{Uniform}(0,1)$ for $i=1,2$.

The following theorem, which generalizes Theorem 2.4.4, sheds light on the equilibria of the All-Access mechanism with reserve price $r$ for a wide range of the cost.

Theorem 2.5.5. Consider the All-Access mechanism with reserve price $r \in[0,1]$ and the setting in Example 2.5.4. Then,

- when cost $c \leq \min \left\{\frac{4 r^{3}-3 r^{2}-6 r+5}{48}, \frac{8 r^{3}+6 r^{2}+7}{96}\right\}$, there exists an equilibrium in which both agents always obtain information, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=1, i=1,2$, for $v_{i, 0} \in[0,1]$, and
- when $c \geq \frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}$ and $r \leq \sqrt{2}-1$, or $c \geq \frac{3 r-r^{3}+1}{12}$ and $r \in(\sqrt{2}-1,1]$, there exists an equilibrium in which none of the agents obtain information, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=0, i=1,2$, for $v_{i, 0} \in[0,1]$.

The proof is given in Appendix B.3.3. Figure 2.3 depicts the results of Theorem 2.5.5.
Next, in Figure ??, we compare the All-Access mechanism (with revenue-maximizing reserve price) with mechanism $\mathcal{M}^{\text {Opt }}$ in terms of their collected revenue for different values of the cost. As expected, the All-Access mechanism fails to obtain the maximum revenue. Surprisingly, the revenue of the All-Access mechanism is not monotone in the cost. For instance, the revenue of the All-Access mechanism when the cost $c$ is 0.18 is less than when the cost is 0.2 . This follows from agent's investment strategies. For $c=0.2$, none of the agents acquire information while for $c=0.18$, agents update their valuations when their initial valuations are close to 0.5 . For the latter case, it is likely that the updated valuations of the agents fall below the reserve price considering the fact that the revenue-maximizing reserve price is almost 0.5 and second signals are drawn from Uniform $(-1,1)$. This will reduce the seller's revenue. In fact, the seller prefers that none of the agents obtain information.

### 2.6 Who Will Be Selected?

Here, we discuss the selection rule of our mechanisms in more detail. We first present an example that shows that the selection rule may not be monotone in initial valuations. We then show that under certain symmetry
assumptions, the selection rule favors agents with higher initial valuations. Finally, we demonstrate how the selection rule reacts to increase in the cost and variance of the second signals.

Figures 2.5 and 2.6 depict the selected agents in the revenue-optimal and efficient mechanisms, respectively, for all realizations of $v_{1,0}$ and $v_{2,0}$ in the range of $[-1.5,2.5]$ with $n=2, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$ for $i=1,2$. The cost of information for the first and second agents is, respectively, 0.01 and 0.05 , that is, $c_{1}=0.01$ and $c_{2}=0.05$. The $x$-axis is the initial valuation of the second agent, and the $y$-axis is the initial valuation of the first agent. The areas in the figures are divided into several regions. In the white and green regions, the number of selected agents is zero and two, respectively. In the purple regions, only agent 1 whose cost of information is lower is selected while in the yellow regions, only agent 2 is selected.

Non-monotonicity of Selection Rule: Observe that the number of selected agents does not always increase as we move along one of the axes. Furthermore, when the initial valuation of a selected agent increases and the initial valuation of the other agent remains the same, he will not necessarily be selected. For instance, the efficient mechanism selects both agents when $v_{1,0}=v_{2,0}=1$ but does not select any agents when $v_{1,0}=2.4$ and $v_{2,0}=1$. The reason that the selection rule is not monotone is that all agents, including those who did not update their valuations, participate in the second round and have a chance to win the item; see Section 2.7.1. The selection rule will remain non-monotone even if the cost of information is the same across agents.

Also observe that the selection rule of both mechanisms favors agents with higher initial valuations and lower costs, as agents with high initial valuations are more likely to win the item in the second round.


Figure 2.3: Depicting the results in Theorem 2.5 .5 with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=$ $\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$.

We can formalize this intuition under certain symmetric and independence assumptions.

Theorem 2.6.1. Suppose for each agent $i$, the second signal $\delta_{i}$ is drawn, independently of other agents' signals, from distribution $G_{i}=\mathcal{G}$. In addition, assume that distribution $\mathcal{G}$ is symmetric and $c_{i}=c$ for $i=1,2, \ldots, n$. If an agent is selected in the revenue-optimal or efficient mechanisms, then all agents with higher initial bids will also be selected.

A distribution $\mathcal{G}$ is symmetric if $\mathcal{G}(-y)=1-\mathcal{G}(y)$. For instance, normal and uniform distributions satisfy the assumption. The proof is presented in Appendix B.2. The intuition is that the seller's objective is an increasing function of the maximum valuation of the agents. When agents are symmetric in terms of the cost and distribution of second signals, the seller would rather select agents with higher initial valuations. These are the agents who have a greater chance of receiving the item in the second round. We note that when the distribution of the valuations are asymmetric, it is likely that the mechanism selects an agent with lower initial valuation but higher variance of the second signal.

Theorem 2.6.1 provides a simple way to find the selected agents. One can sort the agents according to initial valuations (bids) in descending order and evaluate the value of the selection rule's objective function for each of the $n+1$ subsets $\emptyset,\{1\},\{1,2\}, \cdots$, and $\{1,2, \cdots, n\}$ and then select the subset that maximizes the objective. See Guha et al. (2006) and Goel, Guha \& Munagala (2010) for optimization problems similar to our selection rule in more general settings.


Figure 2.4: The revenue of the revenue-optimal and All-Access (with revenue-maximizing reserve price) mechanisms versus the cost with $n=2, F=\operatorname{Uniform}(0,1)$, and $G_{i}=\operatorname{Uniform}(-1,1)$ for $i=1,2$. Here, $c_{1}=c_{2}=c$.

### 2.6.1 Impacts of the Cost and Variance of Second Signals on Selection Rules

To get more insight about the selection rule, we numerically study how the expected number of selected agents changes as the cost and variance of second signals increase. Figure 2.7 illustrates the impact of the variance of second signals, $\sigma^{2}$, on the average number of selected agents with $n=2, F=N(0.5,0.5)$, $G_{i}=N\left(0, \sigma^{2}\right)$, and $c_{i}=0.05$ for $i=1,2^{7}$ When the second signals are more "uncertain," the average number of selected agents in the revenue-optimal (OPT) and efficient (EFF) mechanisms increases. The intuition is that for larger variance, the seller anticipates seeing larger second signals and selects more agents.

Figure 2.8 depicts the impacts of the cost of information, $c_{i}=c$, on the average number of selected agents when $n=2, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$. Both revenue-optimal and efficient mechanisms react to an increase in the cost of information by restricting the number of agents that can obtain information. Interestingly, to obtain higher revenue, the revenue-optimal mechanism selects fewer agents. This implies that the revenue-optimal mechanism distorts the revelation of information to extract more revenue from the agents.

[^23]

Figure 2.5: Selected agents in the optimal mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2$, $c_{1}=0.01, c_{2}=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$ for $i=1,2$.

### 2.7 Extensions

In this section, we discuss some of the extensions of our mechanisms. In particular, we extend our mechanism to a setting where agents are"extremely risk-averse" in a sense that they do not engage in the auction without obtaining additional information. Moreover, we generalize our mechanisms to a setting with multiple units for sale.

### 2.7.1 Information Acquisition as an Entry Cost

In our model, the seller may allocate the item to an agent who has not obtained additional information. However, in some applications, such as the sale of high-valued assets, buyers might face a significant risk if they purchase the item without gathering enough information. For these applications, one can interpret the cost of obtaining information as an entry cost. That is, buyers must invest in information to be considered in the allocation round.


Figure 2.6: Selected agents in the efficient mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2$, $c_{1}=0.01, c_{2}=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$ for $i=1,2$.

Our proposed mechanisms can be extended to this setting by excluding unselected agents from the allocation stage. More precisely, an efficient mechanism in this setting selects the following set of agents

$$
\mathcal{S}_{\mathrm{EFF-E}}\left(b_{0}\right)=\underset{S \subseteq\{1,2, \ldots, n\}}{\arg \max }\left\{\mathrm{E}_{S}\left[\max \left\{\max _{j \in S}\left\{b_{j, 0}+\delta_{j}\right\}, 0\right\}\right]-\sum_{j \in S} c_{j}\right\},
$$

and similarly, a revenue-optimal mechanism selects the following set of agents

$$
\mathcal{S}_{\text {OPT-E }}\left(b_{0}\right)=\underset{S \subseteq\{1,2, \ldots, n\}}{\arg \max }\left\{\mathrm{E}_{S}\left[\max \left\{\max _{j \in S}\left\{b_{j, 0}+\alpha\left(b_{j, 0}\right)+\delta_{j}\right\}, 0\right\}\right]-\sum_{j \in S} c_{j}\right\} .
$$

Note that when $S=\emptyset$, the mechanisms do not allocate the item. We refer to the corresponding efficient and revenue-optimal mechanisms as EFF-E and OPT-E, respectively.

We note that mechanism OPT-E is a generalization of a mechanism proposed by $\mathrm{Lu} \& \mathrm{Ye}$ (2014), who study the problem of designing a revenue-maximizing mechanism with an entry cost when the cost of information is the same across agents. They show that the selection rule in this setting is monotone, that is, the number of selected agents decreases as the initial valuation of an agent increases. Using monotonicity in the selection rule, $\mathrm{Lu} \& \mathrm{Ye}(2014)$ show that their proposed mechanism is incentive compatible. Note that, as we saw in Section 2.6, the selection rule will not be monotone if the item can be allocated to an unselected agent. ${ }^{8}$

[^24]

Figure 2.7: Average number of selected agents in the revenue-optimal and efficient mechanisms versus the standard deviation of second signals, $\sigma$, with $G_{i}=N\left(0, \sigma^{2}\right)$ and $c_{i}=0.05$ for $i=1,2$.

To get more insight, we numerically compare mechanisms $\mathcal{M}^{\text {OPT }}$ and OPT-E. We assume that the cost of information is the same across agents, the number of agents $n=2, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$ for $i=1,2$. Figure 2.9 compares the revenue of the mechanisms. For any value of the cost, mechanism $\mathcal{M}^{\text {Opt }}$ yields more revenue than OPT-E. The revenue of the mechanism OPT-E approaches zero as acquiring information becomes more costly. However, mechanism $\mathcal{M}^{\text {opt }}$ is more robust to the cost of information because, when the cost is high, it can allocate the item without allowing any agents to update their valuations.

### 2.7.2 Multi-unit

In this section, we discuss how our mechanisms can be extended to a setting where $m \geq 1$ units of the item are sold to $n>m$ agents. Specifically, we assume that each agent $i$ needs at most one unit of the item, and his initial valuation for the item is $v_{i, 0} \in[\underline{v}, \bar{v}]$. Similar to our original setting, $v_{i, 0}$ is agent $i$ 's private information and is drawn (independently) from distribution $F$.

We start with an efficient mechanism. The mechanism allocates the item to $m$ agents with the highest (non-negative) bids. In case the number of agents with a positive bid is less than $m$, the mechanism allocates the item to agents with positive bids. The efficient mechanism selects the following set of agents

$$
\begin{equation*}
\max _{S \subseteq\{1,2, \ldots, n\}}\left\{\mathrm{E}_{S}\left[\max _{A \subseteq\{1,2, \ldots, n\},|A| \leq m} \sum_{i \in A} b_{i, 1}\right]-\sum_{i \in S} c_{i}\right\}, \tag{2.7}
\end{equation*}
$$



Figure 2.8: Average number of selected agents in the revenue-optimal and efficient mechanisms versus the cost of information, $c$, with $G_{i}=N(0,0.5)$.
where $A$ is the set of agents that the item will be allocated to, $|A|$ is the cardinality of set $A$, and the condition $|A| \leq m$ implies that we cannot sell more than $m$ units. In addition, $b_{i, 1}=b_{i, 0}+\delta_{i}$ if $i \in S$, and $b_{i, 0}$ otherwise. It is easy to see that the selection rule of the efficient mechanism selects a welfare-maximizing set of agents provided that it is incentive compatible.

The initial payment of the selected agents is the same as before and is given in Eq. (2.2). In addition, only agents that receive the item will pay in the second round. Specifically, any selected agent that wins the item pays $\max \left\{b^{(m+1)}, 0\right\}$, where with slight abuse of notations, $b^{(m+1)}$ is the $(m+1)^{\text {th }}$ highest final bid. Furthermore, any unselected agent $j$ that receives the item pays $\max \left\{b^{(m+1)}, r_{j}\right\}$, where $r_{j}$ solves the following equation

$$
\begin{equation*}
\int_{\max \left\{r_{j}, 0\right\}}^{b_{j, 0}} \operatorname{Pr}\left[z \geq b_{-j}^{(m)}\right] d z=\int_{\underline{v}}^{b_{j, 0}} \mathrm{E}\left[q_{j} \mid v_{j, 0}=z, v_{-j, 0}=b_{-j, 0}\right] d z, \tag{2.8}
\end{equation*}
$$

where $b_{-j}^{(m)}$ is the $m^{\text {th }}$ highest final bid among all agents except for agent $j$. The term inside the integral of the 1.h.s., $\operatorname{Pr}\left[z \geq b_{-j}^{(m)}\right]$, is the cumulative distribution function of random variable $b_{-j}^{(m)}$ at point $z$. Using a similar argument in Lemma B.1.1 in Appendix B.1, one can show that there exists $r_{j} \in\left[0, b_{j, 0}\right]$ that solves the above equation.


Figure 2.9: Average revenue of the OPT and OPT-E mechanisms versus the cost with $G_{i}=N(0,0.5)$, and $c_{i}=c$ for $i=1,2$. In all the figures, $n=2$ and $F=N(0.5,0.5)$.

Next, we will present a revenue-optimal mechanism in this setting. Similar to mechanism $\mathcal{M}^{\text {OpT }}$, the mechanism allocates the item to $m$ agents with the highest (non-negative) weighted bid, where the weighted bid of an agent $i$ is $b_{i, 1}+\alpha\left(b_{i, 0}\right)$. The mechanism allows the following set of agents

$$
\begin{equation*}
\max _{S \subseteq\{1,2, \ldots, n\}}\left\{\mathrm{E}_{S}\left[\max _{A \subseteq\{1,2, \ldots, n\},|A| \leq m} \sum_{i \in A}\left(b_{i, 1}+\alpha\left(b_{i, 0}\right)\right)\right]-\sum_{i \in S} c_{i}\right\} \tag{2.9}
\end{equation*}
$$

to acquire information. The initial payment is given in Eq. (2.2). In the second round, any selected agent that receives the item has to pay the $(m+1)^{\text {th }}$ highest weighted bid if the $(m+1)^{\text {th }}$ highest weighted bid is positive and is zero otherwise. In addition, any unselected agent $j$ that wins the item has to pay a maximum of $(m+1)^{\text {th }}$ highest weighted bid and $r_{j}$. Here, $r_{j}$ solves Eq. (2.8), where $b_{-j}^{(m)}$ should be replaced by the $m^{\text {th }}$ highest final weighted bid among all agents except for agent $j$. For any agents that do not receive the item, their payment in the second round is zero.

### 2.8 Conclusion

Information structure plays a crucial role in the outcome of auctions. This role becomes even more important when information acquisition is costly. We observe that in such environments, agents may over or under-invest in information. We also presented efficient and revenue-optimal mechanisms that shows how auctioneer should control the access to information via a selection rule and prices.

In the previous section, we discussed some of the extensions of our mechanism. An important direction for future research is to extend the results to settings with adaptive selection rules where information is disclosed sequentially over time, and the mechanism makes "selection decisions" based on updated reported valuations of the agents (McAfee \& McMillan 1988). Another important direction is studying environments where the agents are risk-averse.

## Chapter 3

## Dynamic Pricing for Heterogeneous <br> Time-Sensitive Customers

Dynamic pricing is increasingly prevalent in many industries. One of the main advantages of dynamic pricing is that it helps mitigate the risk associated with demand uncertainty (see, for instance, Aviv \& Pazgal 2008 an Cachon \& Swinney 2011). In this paper, we show that dynamic pricing (DP) can play an important role in differentiating between customers over time even in the absence of demand uncertainty. In many settings, especially in fashion and electronic retail, a customer's willingness to pay (or valuation) for a product is time-sensitive and decreases over time. In these situations, customers are not only different in terms of their initial willingness to pay for these items when they are first introduced to the market, but they are also different in terms of how rapidly they lose interest in these products. Thus, we may have customers who initially value the product at a high level, but as time progresses, they lose interest in the product completely. We may also have customers who initially value the product at a low level, but still remain interested in the product as time progresses. That is, the willingness to pay of the lower valuation customers diminishes at a lower rate relative to that of the higher valuation customers. This phenomenon is illustrated in Figure 3.1.

In this paper, we show that when a firm sells to customers who have heterogeneously decreasing valuations, the firm can achieve significant benefits by incorporating dynamic pricing even in the absence of demand uncertainty. This is not the case if customers were homogeneous in their valuation decay rate. In that case, in the absence of demand uncertainty, the firm's optimal pricing strategy would be to post a fixed price, and dynamic pricing would have no benefit. When customer valuations decrease at different rates, the ranking of customers (in terms of their valuations) changes over time (as in Figure 3.1). This allows a firm to generate more revenue by revising its initial price to target customers who currently have higher valuations even though they initially had lower valuations.


Figure 3.1: At time 0 , customer 1 has a higher value than customer 2 . But, her value decreases faster than customer 2 , and beyond $t_{0}$, customer 2 has a higher value.

Formally, we characterize a revenue-optimal selling mechanism for a firm with customers who have heterogeneous valuations that decrease in a heterogeneous fashion. We assume that the firm knows the total demand and the customer valuation distribution, but does not know the precise valuation of each individual customer. We also include the firms' production costs and product holding costs. That is, the firm incurs a constant cost to produce each product and does so at time 0 , and it incurs a cost to hold these products in inventory until the sale. In our setting, the firm commits to a price trajectory, and the customers are strategic in selecting the best time to purchase so as to maximize their individual net utility. We assume that customers with higher initial valuation also have a higher rate of valuation decrease. To the best of our knowledge, this setting has not been studied in the literature.

We next describe the main characteristics of this optimal mechanism ignoring the production and holding costs. The optimal mechanism consists of the firm posting a series of decreasing prices, which essentially divides the customers into three groups based on their initial valuations. The first group comprises all customers with initial valuations above a threshold (high type customers) who purchase the product immediately. The second group consists of all customers with initial valuations below a smaller second threshold (low type customers). For these customers, the posted prices are designed in a way to extract their entire surplus. Finally, the third group consists of customers with valuations between the two thresholds (medium type customers) who do not purchase the product immediately but purchase before the low valuation customers, and obtain a positive net utility, or surplus.

The low type customers in our mechanism play an important role in contrast to what occurs in fixed pricing. In a fixed pricing policy, all customers with valuations above the price would immediately purchase,
and those with valuations below the price would not purchase. However, in our optimal mechanism, the low type customers purchase the product after some delay and the firm is able to extract their entire surplus. In the absence of production and holding costs, the firm sells the product to all customers in this fashion. Selling to all the customers not only can increase social welfare but also can generate significant additional revenue. For instance, we show that the firm can increase its revenue by approximately $23 \%$ by employing the optimal mechanism (relative to fixed pricing) when the initial valuation distribution is uniform and the valuation decay rates are proportional to initial valuations; in fact, more than three quarters of this increase is obtained by selling to the low type customers.

We next investigate the impact of production and holding costs on the optimal selling mechanism. Both of these costs motivate the firm to reduce the length of the selling period, but in different ways. The presence of production costs motivates the firm to sell fewer units by targeting customers with higher initial valuation. Interestingly, we find that the optimal mechanism here can be obtained from the baseline optimal mechanism (obtained without these costs) in a simple manner. The optimal mechanism prices in such a way that customers with initial valuations above a certain cut-off purchase the item at the same time as in the baseline setting (but at a higher price), while those with initial valuations lower than the cut-off do not purchase the product at all. That is, the production costs introduce a cut-off in customers' initial valuation.

Holding costs motivate the firm to price in a manner so that customers are incentivized to make their purchases earlier (than the baseline case). We find that depending on the holding cost, there are three types of optimal mechanisms. If the holding cost is larger than a threshold, then the firm determines it is too expensive to carry the product and simply posts a fixed price so that all customers who purchase the product do so immediately. If the holding cost is moderate (below the previous threshold and above another lower threshold), then the firm benefits from DP but cannot extract the entire surplus of customers with low initial valuations. If the holding cost is below the lower threshold, then the structure of the optimal mechanism is similar to that of the baseline optimal mechanism. There are three distinct groups of customers and the firm can extract the entire surplus of customers with low valuations. Overall, the value of DP decreases with increasing holding and production costs.

Finally, we summarize our main technical contribution. In our setting, one of the hurdles in characterizing the optimal selling mechanism is the lack of consistent customer ranking based on customer types. As a
result, satisfying the individual rationality and incentive compatibility constraints is challenging. ${ }^{1}$ Note that when there is a consistent ranking of customers, individual rationality constraints are binding for the lowest customer type that the firm would like to sell the product to, and the mechanism is incentive compatible if the allocation rule is monotone in the type of the customers. In contrast, in our setting, the individual rationality constraint is binding for a group of customers with low initial valuation. Furthermore, the monotonicity of the time of purchase in the initial valuation does not guarantee that the mechanism is incentive compatible. To characterize the optimal mechanism, we first establish necessary and sufficient conditions to thus have an incentive compatible mechanism. One of these conditions resembles the traditional envelope condition (Myerson 1981), which ensures that the mechanism is locally incentive compatible. The other condition, called interval condition, ensures that the mechanism is globally incentive compatible. We first relax the problem by ignoring the interval condition and characterizing a revenue-optimal mechanism that satisfies the individual rationality constraints and envelope condition. Then, by establishing several additional properties of this mechanism, we show that the mechanism indeed satisfies the interval condition, and thus is optimal.

## Related Work

Our work is related to the growing literature on pricing mechanisms for customers who strategically time their purchases. There is also an extensive literature on dynamic pricing with myopic customers (see for example Lazear (1984), Wang (1993), Gallego \& Van Ryzin (1994b), Feng \& Gallego (1995), Bitran \& Mondschein (1997), Federgruen \& Heching (1999), and Talluri \& Van Ryzin (2004a)). We do not provide a summary of this line of literature here, but we refer the reader to excellent surveys by Bitran \& Caldentey (2003), Chan et al. (2004), and Shen \& Su (2007).

Coase (1972) is one of the first papers to study pricing for strategic customers. Coase conjectured that when a firm sells a durable good to patient and strategic customers and cannot commit to a sequence of posted prices, then the prices would converge to the production cost. Later Stokey (1979), Gul et al. (1986), and Besanko \& Winston (1990) found that with commitment, posting a decreasing sequence of prices is revenue-optimal. In particular, Stokey (1979) showed that when production cost declines over time, posting

[^25]a decreasing sequence of prices results in higher revenue for the firm. However, when the production cost is zero, DP is not beneficial. In contrast, we show that with heterogeneous decay rates, DP can improve revenue even if the production cost is zero.

In the context of revenue management, several works show that DP can increase firm revenue when demand is uncertain (Su 2007, Aviv \& Pazgal 2008, Elmaghraby et al. 2008, Araman \& Caldentey 2009, Cachon \& Swinney 2011, Aviv et al. 2015, and Yu et al. 2015). Specifically, Aviv \& Pazgal (2008) studied a model in which a firm sells a limited inventory of a product in two periods to an unknown number of strategic customers who are heterogeneous in their valuations and time of arrival. They showed that when the level of heterogeneity in customers' valuation increases, the benefit of customer segmentation using pricing decreases. Conversely, in this work we show that as the level of heterogeneity in customers' decay rates increases, the firm can better differentiate customers and generate more revenue.

One important factor that differentiates our work from the aforementioned research is that in our work, demand uncertainty is not a key driver of DP. That is, even in the absence of demand uncertainty, DP increases revenue significantly. Furthermore, in the aforementioned papers, the firm uses the customers' fear of rationing to extract more revenue from strategic customers (see also Liu \& Van Ryzin $2008 b$ and Bansal \& Maglaras 2009), but in our work, customers do not face such a risk. In fact, when the holding and production costs are zero, all the customers purchase the item.

Other works examined intertemporal pricing with new consumers arriving in every period. Conlisk et al. (1984) and Besbes \& Lobel (2015) showed that when customers arrive over time, the firm's optimal strategy is to use a cyclic pricing policy. Borgs et al. 2014 studied how to set prices to extract revenue while guaranteeing service availability to all paying customers arriving and departing at different times. Chen \& Farias (2015) proposed a robust pricing policy for a setting where customers arrive over time and the distribution of waiting cost and valuation of customer is unknown. See Board \& Skrzypacz (2016) and Garrett (2011) for other papers that study pricing with heterogeneous arrivals. In these papers, the firm can gain from DP since it can differentiate customers based on their arrival times. In contrast, in our work, all of the customers are in the market when the sales starts and they strategically optimize their time of purchase. That is, we attempt to isolate and capture the impact of heterogeneity of valuation decay rates on the optimal DP policy, absent any other considerations.

Our work also relates to the growing body of research on dynamic mechanism design; see Bergemann \& Said (2011a) for a survey. There, the firm offers a direct mechanism that allocates the items over time as a
function of customers reports of their private valuations. See Akan et al. (2009), Kakade et al. (2013), Pavan et al. (2014), Battaglini \& Lamba (2012), Boleslavsky \& Said (2013), Golrezaei \& Nazerzadeh (2016), and Lobel \& Xiao (2013) for recent results on designing optimal dynamic mechanisms. In particular, Akan et al. (2009) studied a setting where customers are heterogeneous in their valuation distribution and in how fast they learn their true value. They show that when high type customers (such as business travelers) learn their valuation slower, relative to low type customers (such as leisure travelers), in the optimal mechanism, the firm sequentially screens customers by offering them a menu of expiring refund contracts.

The remainder of this paper is organized as follows: In Section 3.1, we formally define our model. Section 3.2 presents the direct mechanisms. Section 3.3 develops our optimal mechanism when production and holding costs are zero. In Section 3.4, we study the optimal mechanism with positive production and holding costs. We discuss the impact of the degree of customer heterogeneity on the optimal mechanism in Section 3.5. We conclude our paper in Section 3.6.

### 3.1 Model

We consider a firm that sells multiple units of an item (product) to a mass of customers over a time horizon of length $T=\infty$. We assume that the firm produces and stores all units just prior to the start of the sale period. The cost for producing each unit is $c$, and the holding cost to store each unit is $h$ per unit time.

The firm's goal is to implement a selling mechanism to maximize his revenue. At time 0 , he declares and commits to a price trajectory $\mathfrak{p}(t), t \geq 0$. Given the pre-announced prices, customers decide whether and when to purchase the item. Each customer is assumed to be infinitesimal and demands a single unit of the item. The valuation of a customer at time $t$ is $V(\theta, t)$ where $V: \mathbb{R}^{2} \rightarrow \mathbb{R}$ and $\theta$ is the customer type. Valuation function $V$ is known to the firm and customers. However, the customer type is the customer's private information, and these types are independently drawn from a known distribution $F$ with probability density function (p.d.f) $f$, where $F:[\underline{\Theta}, \bar{\Theta}] \rightarrow[0,1]$ and $\underline{\Theta} \geq 0$. The negative inverse hazard rate associated with distribution $F$ is denoted by $\alpha:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}$, and is defined as $\alpha(x)=-\frac{1-F(x)}{f(x)}$. Throughout the paper, we make the following assumption, which implies $\alpha(\cdot)$ is non-decreasing. This assumption is standard in the optimal mechanism design literature (c.f. Myerson (1981)).

## Assumption 3 (Monotone Hazard Rate). The type distribution $F$ has a non-decreasing hazard rate.

We assume that all the customers are present in the market at time 0 and exit after making a purchase. That is, customers can make a purchase at any time $t \geq 0$. Furthermore, they are fully strategic about whether and when they purchase the item from the firm. Specifically, each customer either does not purchase the item, or purchases a unit in the period in which her utility gets maximized. Customers are risk neutral, and utility of a customer with type $\theta$ that purchases the item at time $t$ at price $p$ is $V(\theta, t)-p$. Furthermore, all customers are present in the entire time horizon. Then, given prices $\mathbf{p}=\left\{p_{t}: t \geq 0\right\}$, the customer with type $\theta$ purchases a unit of the item at time $\mathfrak{t}^{*}(\theta):=\arg \max _{\tau \geq 0}\left\{V(\theta, \tau)-p_{\tau}\right\}$ if $V\left(\theta, \mathfrak{t}^{*}(\theta)\right)-p_{\mathfrak{t}^{*}(\theta)} \geq 0$, and she does not purchase otherwise. Here, $p_{t}$ is the price for the item at time $t$.

Here, we consider a deterministic baseline model where the firm knows the total mass of customers. The assumption of deterministic demand is justified when the number of customers is large and fairly predictable. This modeling choice allows us to study the impact of strategic customers and decay in customer valuation, but it deliberately removes the element of uncertainty from the model. That is, we seek to understand if the firm gains from DP when there is no demand uncertainty.

### 3.2 Direct Mechanisms and Optimality

To characterize a revenue-maximizing (optimal) selling mechanism, by the revelation principle, we focus on direct incentive-compatible and individually rational mechanisms where customers first report their type and then the mechanism determines their payment and time of allocation.

More precisely, any direct mechanism $\mathcal{M}$ consists of a pair $(\mathfrak{t}, \mathfrak{p})$, where $\mathfrak{p}:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}$ is a transfer scheme and $\mathfrak{t}:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}$ is an allocation rule. That is, $\mathfrak{p}(\theta)$ and $\mathfrak{t}(\theta)$ are respectively the price for a unit of the item and time of purchase for a customer with type $\theta .{ }^{2}$ When the customer does not purchase the item $\mathfrak{t}(\theta)=\infty$.

We start by defining incentive compatibility and individual rationality. Let $u(\theta, \hat{\theta})$ be the expected utility of a customer with type $\theta$ when she reports $\hat{\theta}$. That is,

$$
u(\theta, \hat{\theta})=V(\theta, \mathfrak{t}(\hat{\theta}))-\mathfrak{p}(\hat{\theta})
$$

[^26]Then, mechanism $\mathcal{M}$ is incentive compatible (IC) if for each customer with type $\theta \in[\underline{\Theta}, \bar{\Theta}]$, truthfulness is a best response, that is, $u(\theta, \hat{\theta}) \leq u(\theta, \theta)$. Roughly speaking, in an incentive-compatible mechanism, no customer wants to deviate from the truthful strategy.

We can now define the individual rationality constraints for the mechanism. An incentive compatible mechanism is individually rational (IR) if for each customer with type $\theta$, her utility under the truthful strategy is non-negative, i.e., for any $\theta \in[\underline{\Theta}, \bar{\Theta}]$, we have $u(\theta, \theta) \geq 0$.

The following lemma presents the necessary and sufficient conditions under which a mechanism is IC.

Lemma 3.2.1 (Necessary and Sufficient Conditions for IC). The mechanism $\mathcal{M}$ with allocation rule $\mathfrak{t}(\cdot)$ is incentive compatible if and only if both conditions stated below are satisfied:

- Envelope Condition: For any $\theta, \hat{\theta} \in[\underline{\Theta}, \bar{\Theta}]$,

$$
\begin{equation*}
u(\theta, \theta)-u(\hat{\theta}, \hat{\theta})=\int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(z)) d z \tag{3.1}
\end{equation*}
$$

where $\partial_{1} V(\theta, t)=\frac{\partial V(\theta, t)}{\partial \theta}$.

- Interval Condition: For any $\hat{\theta}<\theta$,

$$
\begin{equation*}
\int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(\hat{\theta})) d z \leq \int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(z)) d z \leq \int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(\theta)) d z, \tag{3.2}
\end{equation*}
$$

All the proofs are presented in the Online Appendix.

Lemma 3.2.1 is analogous to the characterization of incentive compatibility in standard static settings, where an envelope condition and monotonicity of allocation rule are used to characterize incentive compatibility (see Myerson (1981)). The envelope condition above is a standard one, but the interval conditions replace the monotonicity conditions. The interval conditions compare the utility obtained by the truthful strategy (middle term in Eq. (3.2)) with untruthful strategies.

We are now ready to characterize the firm revenue under any IC mechanism. Note that the expected revenue of a mechanism $\mathcal{M}$ from selling one unit of the item is the customer's payment minus the production and holding costs, that is, $\mathrm{E}[\mathfrak{p}(\theta)-c-h \mathfrak{t}(\theta)]$, where the expectation is with respect to the type of customer $\theta$. Then, the total revenue of the mechanism is the market size times the expected revenue from selling
one unit of the item. Considering that the market size is constant, that is, demand is deterministic, the total revenue of the mechanism is maximized if we maximize the expected revenue from selling one unit of the item.

An incentive-compatible and individually rational mechanism is optimal if it maximizes the expected revenue among all incentive-compatible and individually rational mechanisms.

The following lemma characterizes firm revenue in any incentive-compatible mechanism $\mathcal{M}$.
Lemma 3.2.2 (Revenue of IC Mechanisms). In any incentive-compatible mechanism, the expected firm revenue from selling one unit of the item is given by

$$
\begin{equation*}
\mathrm{E}\left[V(\theta, \mathfrak{t}(\theta))+\alpha(\theta) \partial_{1} V(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta)-c-u(\underline{\Theta}, \underline{\Theta})\right], \tag{3.3}
\end{equation*}
$$

where the expectation is taken with respect to the type of customer $\theta$.

Lemma 3.2.2 suggests that in order to optimize revenue, the optimal mechanism should maximize virtual revenue, that is, $\mathrm{E}\left[\left(V(\theta, \mathfrak{t}(\theta))+\alpha(\theta) \partial_{1} V(\theta, \mathfrak{t}(\theta))\right)-u(\underline{\Theta}, \underline{\Theta})\right]$, and pick a transfer scheme that makes it both IC and IR. Throughout the paper, we refer to $V(\theta, \mathfrak{t}(\theta))+\alpha(\theta) \partial_{1} V(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta)-c$ as virtual value/revenue of a customer with type $\theta$ at time $t$.

The results in this section hold for any generic customer valuation function $V(\theta, t)$. In the next section, we specialize our generic customer valuation function to model the main scenario we are interested in. In this scenario, customers lose interest in the product over time and customers who initially have higher interest in the product lose interest at a faster rate as well. Specifically, we focus on the case of the exponential valuation function $V(\theta, t)=\theta e^{-\beta \theta t}$, which has these properties. We discuss the general valuations function $V$ that satisfies these properties in Section 3.5.

### 3.3 Optimal Mechanism with Exponential Valuation Functions

Here, we present the optimal mechanism when the valuation function is exponential such that $V(\theta, t)=$ $\theta e^{-\beta t \theta}$. Under exponential valuation functions, one cannot define a persistent ranking for customers over time; see Figure 3.1. That a persistent ranking for customers does not exist makes the problem of designing
an optimal mechanism challenging. Furthermore, it allows the firm to extract more revenue by revisiting its prices over time.

We begin by presenting an optimal mechanism for the case when both production and holding costs are zero. Then, we devote Section 3.4 to the study of the optimal mechanism with positive production and holding costs.

By Lemma 3.2.2, an optimal mechanism should maximize virtual revenue subject to IC and IR constraints. Note that when $V(\theta, t)=\theta e^{-\beta \theta t}$ and $c=h=0$, the expected virtual revenue is given by

$$
\begin{align*}
& \mathrm{E}\left[V(\theta, \mathfrak{t}(\theta))+\alpha(\theta) \partial_{1} V(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta}, \underline{\Theta})\right] \\
& =\mathrm{E}\left[e^{-\beta \theta t(\theta)}(\theta+\alpha(\theta)(1-\beta \theta \mathfrak{t}(\theta)))-u(\underline{\Theta}, \underline{\Theta})\right]:=\mathrm{E}[R(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta}, \underline{\Theta})], \tag{3.4}
\end{align*}
$$

where $R(\theta, t)=e^{-\beta \theta t}(\theta+\alpha(\theta)(1-\beta \theta t))$ is the virtual value of a customer of type $\theta$ at time $t$. Note that initial virtual value, i.e., $R(\theta, 0)=\theta+\alpha(\theta)$, is equal to the virtual value in a standard static setting (c.f. Myerson 1981).

To characterize the optimal mechanism, we need to solve the following optimization problem.

$$
\begin{gathered}
\max _{\{t(\theta), \mathfrak{p}(\theta): \theta \in[\underline{\Theta}, \widehat{\Theta}]\}, u(\underline{\Theta}, \underline{\Theta}) \geq 0}\{\mathrm{E}[\max \{R(\theta, t(\theta))-u(\underline{\Theta}, \underline{\Theta}), 0\}]\} \\
\text { s.t. IC and IR constraints }
\end{gathered}
$$

Solving the above optimization problem is rather involved because we are maximizing over the allocation and payment functions $\mathfrak{t}(\cdot)$ and $\mathfrak{p}(\cdot)$. For this reason, we characterize the optimal solution of the above equation in two steps. In the first step, we relax the problem by ignoring the interval conditions. Recall that by Lemma 3.2.1, satisfying the IC constraints is equivalent to satisfying the envelope and interval conditions.

That is, in the first step, we only focus on satisfying the envelope conditions and IR constraints. In particular, we consider the following relaxed problem:

$$
\begin{align*}
& \max _{\{\mathfrak{t}(\theta): \theta \in[\underline{\Theta}, \bar{\Theta}]\}, u(\underline{\Theta}, \underline{\Theta}) \geq 0} \mathrm{E}[\max \{R(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta}, \underline{\Theta}), 0\}] \\
& \text { s.t. } u(\theta, \theta)=u(\underline{\Theta}, \underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta z \mathfrak{t}(z)}(1-\beta z \mathfrak{t}(z)) d z \geq 0 \text { for } \theta \in[\underline{\Theta}, \bar{\Theta}] \tag{IR}
\end{align*}
$$

Note that the equation $u(\theta, \theta)=u(\underline{\Theta}, \underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta z \mathfrak{t}(z)}(1-\beta z \mathfrak{t}(z)) d z$ follows from the envelope conditions. In the second step, we show that the solution to the relaxed problem satisfies the envelop conditions. This implies that the optimal solution of Problem relaxed is also an optimal solution of Problem OPT. More details about the proof are presented in Section 3.4.

The following lemma characterizes the optimal solution of the relaxed problem.

Lemma 3.3.1 (Optimal Solution of Problem relaxed). Given that Assumption 3 holds, in an optimal solution of Problem RELAXED, $u(\underline{\Theta}, \underline{\Theta})=0$ and the allocation rule, denoted by $\mathfrak{t}_{R}$, is given by

$$
\mathfrak{t}_{R}(\theta)=\left\{\begin{array}{lll}
0 & \text { if } \theta \geq \theta_{H} & \text { High-Type } ;  \tag{3.5}\\
\frac{\theta+2 \alpha(\theta)}{\beta \theta \alpha(\theta)} & \text { if } \theta \in\left[\theta_{L}, \theta_{H}\right] & \text { Medium-Type } ; \\
\frac{1}{\beta \theta} & \text { if } \theta \in\left[\underline{\Theta}, \theta_{L}\right] ; & \text { Low-Type }
\end{array}\right.
$$

where $\theta_{H}$ solves $\theta_{H}+2 \alpha\left(\theta_{H}\right)=0$ and $\theta_{L}$ solves $\theta_{L}+\alpha\left(\theta_{L}\right)=0$.

The following theorem shows that the optimal solution of the relaxed problem, given in Lemma 3.3.1, fulfills the interval conditions. This implies that a mechanism with time of allocation $\mathfrak{t}_{R}(\cdot)$ is revenueoptimal.

Theorem 3.3.2 (Optimal Mechanism). If Assumption 3 holds, the customer valuation function is given by $V(\theta, t)=\theta e^{-\beta \theta t}$, and the production and holding costs are zero, then in the optimal mechanism a customer of type $\theta \in[\underline{\Theta}, \bar{\Theta}]$ purchases one unit of the item at time $\mathfrak{t}_{R}(\theta)$, given in Eq. (3.6), and at price $\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{R}(\theta)\right)-\int_{\underline{\Theta}}^{\theta} e^{-\beta \mathfrak{t}_{R}(z) z}\left(1-\beta \mathfrak{t}_{R}(z) z\right) d z$.

Note that $\beta \mathfrak{t}_{R}(\theta)$ does not depend on $\beta$. Thus, if $\beta$ gets doubled, the purchase time of all customers decreases by a factor of 2 . Then, considering the fact that $u(\theta, \theta), \mathfrak{p}(\theta)$, and the revenue of the firm only depends on $\beta \mathfrak{t}_{R}(\theta)$, we can conclude that $u(\theta, \theta)$ and $\mathfrak{p}(\theta)$ stays the same when $\beta$ changes.

In the following, we will discuss the main insights of Theorem 3.3.2. First we note that the firm sells the item to all customers. In addition, the purchase time of customers $\mathfrak{t}_{R}(\theta)$ is decreasing in customer type, that is, customers with lower initial valuation purchase the item later than customers with higher initial valuation.

Theorem 3.3.2 shows that the optimal mechanism divides the customers into three groups: high-type, medium-type, and low-type. The high-type customers who have high initial valuation ( $\theta \geq \theta_{H}$ ) purchase the item immediately. The low-type customers who have low initial valuation $\theta \leq \theta_{L}$ ) delay their time of purchase. We note that the purchase time of high- and low-type customers does not depend on the distribution of the customer type, $F$, directly. The time of purchase of these customers depend on $F$ only through thresholds $\theta_{H}$ and $\theta_{L}$.

Observe that low-type customers with type $\theta$ purchase the item at time $\frac{1}{\beta \theta}$ and more importantly, get zero utility. To understand why, note that these customers pay

$$
\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{R}(\theta)\right)-\int_{\underline{\Theta}}^{\theta} e^{-\beta \mathfrak{t}_{R}(z) z}\left(1-\beta \mathfrak{t}_{R}(z) z\right) d z=V\left(\theta, \mathfrak{t}_{R}(\theta)\right),
$$

where the second equality holds because $\mathfrak{t}_{R}(z)=\frac{1}{\beta z}$ for any $z \leq \theta_{L}$. Note that this is in contrast with the traditional static mechanism design. In the static mechanism design, customers whose "type" is high enough get the product and enjoy a positive surplus, whereas low type customers do not get the product at all. In fact, there is typically one customer type on the boundary that does gets zero utility after purchasing the product. However, in our setting, we show that there exists a group of customers who purchase the item and obtain zero utility.

Note that low-type customers have negative virtual values at time zero, that is, $R(\theta, 0)=\theta+\alpha(\theta) \leq 0$ for any $\theta \leq \theta_{L}$. However, the virtual value of these customers eventually becomes positive over time, as their valuations do not decay very quickly. Therefore, the firm gains from selling the item to them later in time.

The medium-type customers who have medium initial valuation $\theta \in\left[\theta_{L}, \theta_{H}\right]$ do not purchase the item immediately. However, unlike the low-type customers, these customers enjoy a positive utility.

The optimal mechanism presented in Theorem 3.3.2 highlights the fact that the firm benefits from heterogeneity in valuation decay rate by adopting DP and delaying allocation. In fact, the extra revenue that the firm makes comes partly from the low-type customers from whom the firm extracts their entire surplus.

Comparison to a Model with Homogenous Valuation Decay Rate: When customer valuation decay rate is homogeneous, that is, is $V(\theta, t)=\theta e^{-\beta t}$, the optimal mechanism posts a fixed price $\theta_{L}$. Then, customers with a type greater than $\theta_{L}$ purchase the item at time zero. Note that here the low-type customers do not purchase the item at all. The reason is that, with homogeneous decay rates, these customers have negative virtual value throughout the time horizon. Thus, the firm is not willing to sell the item to these customers. In addition, we observe that medium-type customers, who delay their time of purchase under heterogeneous decay rates, get the item immediately with homogeneous valuation decay rates.

To shed more lights on Theorem 3.3.2, we study the following example.

Example 3.3.3 (Optimal Mechanism). Assume that $\theta$ follows the uniform distribution in the range of $[0,1]$; that is, $\theta \sim U(0,1)$. Then, the time of purchase in the optimal mechanism is given by

$$
\mathfrak{t}_{R}(\theta)= \begin{cases}0 & \text { if } \theta \geq \frac{2}{3}  \tag{3.6}\\ \frac{3 \theta-2}{\beta \theta(\theta-1)} & \text { if } \theta \in\left[\frac{1}{2}, \frac{2}{3}\right] \\ \frac{1}{\beta \theta} & \text { if } \theta \in\left[0, \frac{1}{2}\right]\end{cases}
$$

The allocation rule, $\mathfrak{t}_{R}(\theta)$, the payment rule, $\mathfrak{p}(\theta)$, and the utility of customers, $u(\theta, \theta)$, are depicted in Figures 3.2, 3.3, and 3.4 when $\beta=0.1$. This figure compares the DP policy of Theorem 3.3.2 to the FP policy. Note that the revenue-maximizing FP policy, posts a price of $\theta_{L}=\frac{1}{2}$ at time zero.

We observe that under FP, there is a one-time sale where only customers with a type greater than $\frac{1}{2}$ purchase the item immediately. But, under DP, customers with type $\theta \geq \theta_{H}=\frac{2}{3}$ purchase the item at time zero, and the rest of the customers delay their purchase. Note that the purchase time under DP, $\mathfrak{t}_{R}(\theta)$, is decreasing in the type of customer.

In Figure 3.3, we observe that the payment of customers under DP increases as their type grows. Hightype customers pay more when the firm uses DP rather than FP. However, the medium-type customers have a lower payment under DP. This group of customers delay their purchase and their valuation at the purchase time is not as high as their initial valuation. Therefore, the firm needs to reduce their payments.

Furthermore, the DP policy enables the firm to extract revenue from low-type customers. The firm increases its revenue by $18 \%$ by only selling the item to this group of customers. Note that under dynamic pricing the firm earns an expected revenue of 0.31 from each customer. However, it only earns an expected revenue of 0.25 from each customer by using $F P$. This implies that by adopting $D P$, the firm increases its revenue by more than $23 \%$.

Figure 3.4 shows that under both dynamic and fixed pricing policies the utility of customers, $u(\theta, \theta)$, is an increasing function of $\theta$. But, the customers earn higher utility under the FP policy. However, as we show in Example 3.4.3, the social welfare of the customers and the firm under DP is higher than that under $F P$.

Example 3.3.3 shows that the DP policy earns significantly more revenue than the FP policy. Motivated by that, we present a lower bound on the revenue gain of DP over FP in Appendix C.1.2. We drive the lower bound on the revenue gain by characterizing the extra revenue the firm extracts from the low-type customers. The bound suggests that for the setting in Example 3.3.3, the DP policy earns at most $18 \%$ more revenue than FP. Another interpretation of this result is, if the firm ignores the heterogeneity in decay rates and follows the optimal mechanism under the homogenous model, it will suffer from at least $18 \%$ revenue loss.

### 3.4 Optimal Mechanism with Production and Holding Costs

In this section, we present the optimal mechanism when the production and holding costs are positive. These costs have significant impact on the optimal selling mechanism. In particular, as we show in Section 3.4.1


Figure 3.2: Time of purchase versus type in the fixed and dynamic pricing policies with the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and $h=c=0$.
that when the firm only faces production costs, it ends the sale sooner, compared to when the production cost is zero. In particular, the production cost introduces a cut-off such that customers whose type $\theta$ is greater than the cut-off purchase at time $\mathfrak{t}_{R}(\theta)$, and other customers do not purchase the item at all. That is, a positive production cost does not change the time of allocation of customers who purchase the item.

Similarly, a positive holding cost motivates the firm to end the sale sooner; see Section 3.4.2. However, with a positive holding cost, the firm incentivizes customers to purchase the item earlier, as carrying the items in inventory is costly. This is in contrast with the optimal mechanism with a positive production cost. There, the purchase time for all the customers that make a purchase remain the same, whereas with a positive holding cost, customers are incentivized to purchase the item sooner.

When the firm faces both production and holding costs, the optimal mechanism is obtained by introducing a cut-off to the optimal mechanism with the same holding cost and zero production cost, called the baseline mechanism. Namely, customers whose initial valuation is above the cut-off purchase the item at the same time as the baseline mechanism, and customers whose initial valuation is lower than the cut-off do not purchase the item at all. In the optimal mechanism, all customers who do not purchase the item in the baseline mechanism do not purchase the item. But, customers with initial valuation lower than the cut-off and who purchase the item in the baseline mechanism no longer purchase the item anymore, as the firm increases the prices.

The structure of the optimal mechanism is intuitive. One can think about the holding and production costs as the dynamic and fixed costs of carrying inventory, respectively. The dynamic cost determines how the firm allocates the items to the customers over time, and the production cost determines whether the firm is willing to sell the items to the customers.


Figure 3.3: Payment versus type in the fixed and dynamic pricing policies with the customer type $\theta \sim$ $U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and $h=c=0$.

To capture the impact of each cost separately in the next section, we focus on the case of a positive production cost. Then, we study the impact of the holding cost in Section 3.4.2. Finally, in Section 3.4.3, we present the optimal mechanism with positive production and holding costs.

### 3.4.1 Positive Production Cost

Here, we present an optimal mechanism when the production $\operatorname{cost}$ is $c \geq 0$, and the holding cost $h$ is zero. Defining $\theta_{c}$ as the smallest value that solves $R\left(\theta_{c}, \mathfrak{t}_{R}\left(\theta_{c}\right)\right)=c$, we present the main result of this section. Here, $R(\theta, t)$ and $\mathfrak{t}_{R}(\cdot)$ are defined in Equations (B.14) and (3.6), respectively.

Theorem 3.4.1 (Positive Production Cost). Given that Assumption 3 holds, the customer valuation function is given by $V(\theta, t)=\theta e^{-\beta \theta t}$, the production cost $c \in[0, \bar{\Theta}]$, and the holding cost $h=0$, the optimal mechanism only sells to customers with type $\theta \geq \theta_{c}$ at time $\mathfrak{t}_{R}(\theta)$, given in Eq. (3.6), and at price $\mathfrak{p}(\theta)=$ $V\left(\theta, \mathfrak{t}_{R}(\theta)\right)-\int_{\theta_{c}}^{\theta} e^{-\beta \mathfrak{t}_{R}(z) z}\left(1-\beta \mathfrak{t}_{R}(z) z\right) d z$. Furthermore, $\mathfrak{p}(\theta)=\infty$ for $\theta<\theta_{c}$.

In Theorem 3.4.1, we assume that the production cost $c$ is less than $\bar{\Theta}$, as the firm has no incentive to produce and sell the items when the production cost is greater than the maximum valuation of customers, that is, $\bar{\Theta}$.

The main idea of the proof is to show that the virtual value of a customer with type $\theta$ at time $\mathfrak{t}_{R}(\theta)$, that is, $R\left(\theta, \mathfrak{t}_{R}(\theta)\right)$, is increasing in $\theta$. Then, provided that $R\left(\theta_{c}, \mathfrak{t}_{R}\left(\theta_{c}\right)\right)-c=0$, we have $R\left(\theta, \mathfrak{t}_{R}(\theta)\right)-c<0$ for any $\theta<\theta_{c}$. This implies that the seller would rather not sell the item to customers with type $\theta<\theta_{c}$.

Theorem 3.4.1 suggests that the production cost will not change the allocation time of customers with type $\theta \geq \theta_{c}$; rather, it only changes the payment rule such that the lower-type customers are not willing to


Figure 3.4: Utility versus type in the fixed and dynamic pricing policies with the customer type $\theta \sim U(0,1)$, $V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and $h=c=0$.
purchase the item. In other words, the payment rule is designed to enforce a cut-off, $\theta_{c}$, in the allocation rule. For more insight into Theorem 3.4.1, we revisit Example 3.3.3.

Example 3.4.2 (Revisiting Example 3.3.3: Positive Production Costs). Consider the same setting in Example 3.3.3. Figure 3.5 illustrates the cut-off $\theta_{c}$ as a function of the production cost. Observe that when the production cost is greater than 0.33 , the firm only sells the product to the high-type customers with $\theta \geq \theta_{c}$. This implies that posting a fixed price is optimal when the production cost is high enough. However, when the production cost is between 0.18 and 0.33 , then the firm sells the item to all the high-type customers and some medium-type customers, that is, those with type $\theta \in\left[\theta_{c}, \frac{2}{3}\right]$. Thus, with medium production costs, DP increases firm revenue. Finally, when the production cost is less than 0.18, the firm considers selling the item to all high-type and medium-type customers and some of low-type customers with $\theta \in\left[\theta_{c}, \frac{1}{2}\right]$. As a result, with a low production cost, the firm has an opportunity to extract the entire surplus of some of the low-type customers.

Figure 3.5 compares the threshold $\theta_{c}$ with that in the FP policy. Note that in the FP policy, the threshold $\theta_{c, f}$ solves $R\left(\theta_{c, f}, t=0\right)=c$. We observe that the threshold is smaller than in the DP policy, which suggests that the DP policy sells to more customers than the FP policy.

In Figure 3.6, we further compare the DP and FP policies in terms of their revenue and social welfare. Note that with allocation time $\mathfrak{t}_{R}(\cdot)$ and cut-off $\theta_{c}$, the social welfare is $E\left[\left(\theta e^{-\beta \theta t_{R}(\theta)}-c\right) \mathbb{1}\left\{\theta \geq \theta_{c}\right\}\right]$ where the first term, $\theta e^{-\beta \theta t_{R}(\theta)}$, is the valuation of a customer with type $\theta$ who purchases the item at time $\mathfrak{t}_{R}(\theta)$, and the second term is the production cost. Here, $\mathbb{1}\{A\}$ is an indicator function; that is $\mathbb{1}\{A\}=1$ when event A happens and is zero otherwise. Figure 3.6 shows that the relative welfare gain of DP to FP is not monotone in the production cost. When the production cost is small enough, social welfare under the DP policy is greater than that under the FP policy. Under the DP policy, the firm sells to more customers, as compared to the FP policy. However, under the DP policy, customers purchase the item later, as compared to the FP policy. While the former has a positive impact on the social welfare of the DP policy, the latter impacts its social welfare negatively. We observe that when the production cost is small enough, the former factor dominates the latter, and as a result, DP outperforms FP in terms of social welfare. But when the production cost is large, the latter factor dominates, and thus, the FP policy yields more social welfare than the DP policy.

Figure 3.6 shows that as the production cost gets smaller, DP earns more revenue than FP, because with a small production cost, the firm can afford to sell at a lower price and delay the allocation time of customers in order to extract more revenue from them.

### 3.4.2 Positive Holding Cost

In this section, we characterize an optimal mechanism for when the holding cost is positive. We will show that by having a positive holding cost, the firm motivates the customers to purchase earlier. Furthermore, by having a positive holding cost, the firm induces more customers to purchase the item immediately. This is in contrast with a positive production cost. Recall that with a positive production cost, the purchase time for all customers who make a purchase remain the same.

Before presenting the optimal mechanism with a positive holding cost, to get intuition on the impact of the holding cost, we revisit Example 3.3.3 when $h>0$.

Example 3.4.3 (Revisiting Example 3.3.3: Positive Holding Costs). Here, we present the optimal mechanism for the setting in Example 3.3.3 when the holding cost $h>0$ and the production cost $c=0$. Figure 3.7 shows how the optimal mechanism divides customers in different regions. A precise definition of the boundaries of these regions will be given later in Equations (3.7) and (3.8).

Observe that when the holding cost is small $\left(h \leq H_{l}:=0.004\right)$, there are four regions: high-type, medium-type, low-type, and no-allocation. We later define $H_{l}$ in Eq. (3.7) for any type distribution $F$.


Figure 3.5: The cut-off $\theta_{c}$ versus the production cost $c$ for the mechanism described in Theorem 3.4.1 and the FP policy. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and the holding cost $h=0$.

While customers in the high-type region get the item immediately, customers in the low-type and mediumtype regions delay their purchase time. Moreover, customers in the no-allocation region do not purchase the item at all. These customers and customers in the low-type region get zero utility. We note that as the holding cost increases, the low-type region shrinks, whereas other regions grow. This pattern continues until the holding cost hits $H_{l}$. At $h=H_{l}$, the low-type region vanishes and there will be only three regions: high-type, medium-type, and no-allocation. By increasing the holding cost from $H_{l}$ to $H_{h}:=0.025$, the high-type region gets larger while the medium-type region gets smaller; see the definition of $H_{h}$ in Eq. (3.7) for any type distribution $F$. In fact, at $h=H_{h}$, the medium-type region disappears. Finally, for $h \geq H_{h}$, there will be only two regions: high-type and no-allocation. That is, when the holding cost is high enough, the firm posts a fixed price, which only incentivizes the high-type customers to purchase the item at time zero.

Figure 3.8 shows the social welfare and revenue gain of DP as a percentage (relative to the FP policy) versus the holding cost. We note that the FP policy does not change as the holding cost varies. For any value of the holding cost, the FP policy posts a price of $\theta_{L}=\frac{1}{2}$. Thus, the social welfare of the FP policy is $E\left[\theta \times \mathbb{1}\left\{\theta \geq \theta_{L}\right\}\right]$. The social welfare of the DP policy is the expected value of customers at the time of purchase $\mathfrak{t}_{h}(\cdot)$ minus the holding cost, that is, $E\left[\left(\theta e^{-\beta \mathfrak{t}_{h}(\theta) \theta}-h \mathfrak{t}_{h}(\theta)\right) \times \mathbb{1}\left\{\theta \geq \underline{\theta}_{h}\right\}\right]$ where $\underline{\theta}_{h}$ is the lowest type that purchases the item. We formally define $\mathfrak{t}_{h}(\theta)$ and the cut-off $\underline{\theta}_{h}$ in Equations (3.8) and (3.7), respectively.

Interestingly, social welfare is not monotone in the holding cost. At first glance, we expect social welfare decreases when the holding cost gets larger, but this is only case when the holding cost is not too large. For


Figure 3.6: The social welfare and revenue gain of DP (relative to FP ) as the production cost varies. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and the holding cost $h=0$.
larger holding cost values, social welfare increases in $h$. To understand why, note that by increasing the holding cost, the firm incentivizes customers to purchase earlier as holding the items is costly; see Figure 3.9. This, in turn, enhances the social welfare as value increases for customers at their time of purchase. Furthermore, we observe that when the holding cost is not too large, the social welfare of the optimal DP mechanism is greater than that of FP. Thus, for small holding cost values, DP not only increases firm revenue but also social welfare.

Figure 3.8 shows that DP outperforms FP by a higher percentage when the holding cost is small, because a smaller holding cost allows the firm to lower prices and further delay the time of allocation to customers in order to extract more surplus from them.

Figures 3.9, 3.10, and 3.11, respectively, show the time of purchase of customers and their utility and payment as a function of the customer type $\theta$ for different holding cost values, in particular, $h=0$ (no holding cost), $H_{l} / 5$ (low holding cost), and $\left(H_{l}+H_{h}\right) / 2$ (medium holding cost). We further show the curves for $H_{l}$ and $H_{h}$ because at these holding costs, the structure of the optimal mechanism changes. Recall that $H_{l}$ is the lowest holding cost under which the firm only sells to high- and medium-type customers, and $H_{h}$ is the lowest holding cost under which the firm only sells to high-type customers.

We begin by comparing the curves for $h=0$ and $h=H_{l} / 5$. We observe that by increasing $h$ from 0 to $H_{l} / 5$, the purchase time of high- and low-type customers who purchase the item does not change. However, medium-type customers purchase the item earlier. Because of this, the utility of medium- and high-type customers increases slightly. We observe a similar phenomena when the holding cost increases from $H_{l} / 5$ to $H_{l}$. This implies that increasing the holding cost can benefit higher-type customers.


Figure 3.7: The structure of the optimal mechanism as a function of the holding cost $h$. Here, the customer type $\theta \sim U(0,1), \beta=0.1, V(\theta, t)=\theta e^{-\beta \theta t}$, and the production cost $c=0$.

Next, we compare the case of $h=H_{l}$ and $h=\left(H_{l}+H_{h}\right) / 2$. As expected, customers purchase earlier when $h=\left(H_{l}+H_{h}\right) / 2$. Moreover, relative to $h=H_{l}$, when $h=\left(H_{l}+H_{h}\right) / 2$, customers with higher-type customers pay a lower price, while lower-type customers end up paying more. As a result, by increasing $h$ from $H_{l}$ to $\left(H_{l}+H_{h}\right) / 2$, while the utility of higher-type customers increases, the utility of lower-type customers decreases.

Example 3.4.3 illustrates how the holding cost influences the structure of the optimal mechanism. Next, we formalize these observations by presenting the optimal mechanism with a positive holding cost. We will show that the optimal mechanism only sells to customers with initial valuation $\theta \geq \underline{\theta}_{h}$, where the cut-off $\underline{\theta}_{h}$ depends on the holding cost, $h$, and is given by

$$
\underline{\theta}_{h}:=\left\{\begin{array}{lll}
\max \left\{\underline{\theta}_{L}, \underline{\Theta}\right\} & \text { if } h<H_{l} & \text { Low Holding Cost; }  \tag{3.7}\\
\underline{\theta}_{M} & \text { if } h \in\left[H_{l}, H_{h}\right] & \text { Medium Holding Cost; } \\
\theta_{H} & \text { if } h>H_{h} & \text { High Holding Cost. }
\end{array}\right.
$$

Here, $H_{l}=\beta(\tilde{\theta})^{2} e^{-1}$ and $H_{h}=\beta\left(\theta_{L}\right)^{2}$ where $\tilde{\theta}$ solves $2 \tilde{\theta}+\alpha(\tilde{\theta})=0$ and $\theta_{L}$ is defined in Eq. (3.6). Considering the fact that $\tilde{\theta} \leq \theta_{L}$, it is easy to observe that $H_{l}<H_{h}$. We say the holding cost is low and medium when $h<H_{l}$ and $h \in\left[H_{l}, H_{h}\right]$, respectively. In addition, when $h>H_{h}$, we say the holding cost is high. For low holding costs $\left(h \leq H_{l}\right)$, the cut-off $\underline{\theta}_{h}=\underline{\theta}_{L}$, where $\underline{\theta}_{L}$ solves $\beta e^{-1}\left(\underline{\theta}_{L}\right)^{2}=h$. Observe that at $h=H_{l}$, where the cut-off $\underline{\theta}_{L}=\tilde{\theta}$, and $h=0$, we have $\underline{\theta}_{L}=0$.


Figure 3.8: The social welfare and the revenue gain of the optimal mechanism (relative to the FP policy) as a percentage versus the holding cost. Here, the customer type $\theta \sim U(0,1), \beta=0.1, V(\theta, t)=\theta e^{-\beta \theta t}$, and the production $\operatorname{cost} c=0$.

When the holding cost is medium, only customers with initial valuation $\theta \geq \underline{\theta}_{M}$ purchase the item where $\underline{\theta}_{M}$ solves $R\left(\underline{\theta}_{M}, \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right)\right)-h \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right)=0$ and $R(\theta, t)$ is defined in Eq. (B.14). Here, $\mathfrak{t}_{f}(\theta)$, called the FOC solution, is a solution that satisfies the first order condition (FOC). Precisely, $\mathfrak{t}_{f}(\theta)$ solves $\left.\frac{\partial R(\theta, t)-h t}{\partial t}\right|_{\mathfrak{t}_{f}(\theta)}=0$. Thus, one can show that when $h=H_{l}, \tilde{\theta}$ solves $R\left(\tilde{\theta}, \mathfrak{t}_{f}(\tilde{\theta})\right)-h \mathfrak{t}_{f}(\tilde{\theta})=0$.

We show that in the optimal mechanism, the time of sale of a customer with type $\theta \geq \underline{\theta}_{h}$ is

$$
\mathfrak{t}_{h}(\theta):=\left\{\begin{array}{lll}
0 & \text { if } \theta \geq \theta_{H}^{h} & \text { High-Type }  \tag{3.8}\\
\mathfrak{t}_{f}(\theta) & \text { if } \theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right] & \text { Medium-Type } \\
\frac{1}{\beta \theta} & \text { if } \theta \in\left[\underline{\Theta}, \theta_{L}^{h}\right] & \text { Low-Type }
\end{array}\right.
$$

Here, for $h>H_{h}$, we have $\theta_{H}^{h}=\theta_{L}$, and for $h \leq H_{h}, \theta_{H}^{h} \in\left[\theta_{L}, \theta_{H}\right]$ solves

$$
\left.\frac{\partial\left(R\left(\theta_{H}^{h}, t\right)-h t\right)}{\partial t}\right|_{t=0}=-\beta \theta_{H}^{h}\left(\theta_{H}^{h}+2 \alpha\left(\theta_{H}^{h}\right)\right)-h=0
$$

That is, for $h \leq H_{h}$, the FOC solution at $\theta_{H}^{h}$, that is, $\mathfrak{t}_{f}\left(\theta_{H}^{h}\right)$, is zero. We note that for any $\theta<\theta_{H}^{h}$, the FOC solution is negative, and $R(\theta, t)-h t$ is maximized at $t=0$. We also observe that at $h=0$ and $h=H_{h}, \theta_{H}^{h}$ is respectively $\theta_{H}$ and $\theta_{L}$. Furthermore, $\theta_{H}^{h}$ is decreasing in $h$, indicating that as the holding cost increases, more customers purchase the item at time zero.


Figure 3.9: Time of purchase versus type. Here, $H_{l}=0.004, H_{h}=0.025$, the customer type $\theta \sim U(0,1)$, $V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and production $\operatorname{cost} c=0$.

We now define $\theta_{L}^{h}$ in Eq. (3.8), where $h \leq H_{l}, \theta_{L}^{h} \in\left[\tilde{\theta}, \theta_{L}\right]$ solves

$$
\left.\frac{\partial\left(R\left(\theta_{L}^{h}, t\right)-h t\right)}{\partial t}\right|_{t=\frac{1}{\beta \theta_{L}^{h}}}=-\beta \theta_{L}^{h} e^{-1}\left(\theta_{L}^{h}+\alpha\left(\theta_{L}^{h}\right)\right)-h=0,
$$

and for $h>H_{l}$, we have $\theta_{L}^{h}=\tilde{\theta}$. That is, for $h \leq H_{l}$, we have $\mathfrak{t}_{f}\left(\theta_{L}^{h}\right)=\frac{1}{\beta \theta_{L}^{h}}$. We note that at $h=0$ and $h=H_{l}, \theta_{L}^{h}$ is respectively $\theta_{L}$ and $\tilde{\theta}$. Furthermore, $\theta_{L}^{h}$ is decreasing in $h$. This suggests that as the holding cost increases, the highest-type customer who make a purchase and get zero utility decreases. That is, the low-type group gets smaller. Figure 3.7 in Example 3.4 . 3 shows how $\theta_{L}^{h}$ and $\theta_{L}^{h}$ vary as the holding cost $h$ increases.

We now can describe the optimal mechanism by consolidating the time of purchase $\mathfrak{t}_{h}(\cdot)$ and the cut-off $\underline{\theta}_{h}$. In the optimal mechanism, when the holding cost is low, that is, $h \leq H_{l}$, the firm sells to high- and medium-type customers and some low-type customers with $\theta \in\left[\max \left\{\underline{\Theta}, \underline{\theta}_{L}\right\}, \theta_{L}^{h}\right]$. Thus, with low holding costs, the firm can extract the full surplus of some of low-type customers. However, under medium and high holding costs, the firm has no such opportunity, as it must end the sale early due to the high cost of carrying the items. In particular, when the holding cost is medium, $h \in\left[H_{l}, H_{h}\right]$, the firm only sells to high-type and some of medium-type customers, that is, those with $\theta \in\left[\underline{\theta}_{M}, \theta_{H}^{h}\right]$. Finally, when the holding cost is high, the firm does not benefit from the heterogeneity of valuation decay rates, and it simply posts a fixed price of $\theta_{L}$. Then, only customers with a type greater than $\underline{\theta}_{h}=\theta_{L}$ purchase the item at time zero.

The following theorem formally characterizes the optimal mechanism.


Figure 3.10: Payment versus type. Here, $H_{l}=0.004, H_{h}=0.025$, the customer type $\theta \sim U(0,1)$, $V(\theta, t)=\theta e^{-\beta \theta t}, \beta=0.1$, and production cost $c=0$.

Theorem 3.4.4 (Positive Holding Cost). If Assumption 3 holds and $\tilde{\theta}$ is the unique solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)$ $H_{l} \mathfrak{t}_{f}(\theta)=0$, then the optimal mechanism sells to customers of type $\theta \geq \underline{\theta}_{h}$ at time $\mathfrak{t}_{h}(\theta)$ and at price $\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{h}(\theta)\right)-\int_{\underline{\theta}_{h}}^{\theta} e^{-\beta \mathfrak{t}_{h}(z) z}\left(1-\beta \mathfrak{t}_{h}(z) z\right) d z$ where $\mathfrak{t}_{h}(\theta)$ and $\underline{\theta}_{h}$ are defined in Equations (3.7) and (3.8), respectively. Furthermore, for $\theta<\underline{\theta}_{h}, \mathfrak{p}(\theta)=\infty$.

In Theorem 3.4.4, we assume that at $h=H_{l}$, the solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$ is unique. In Lemma C.1.13 in Appendix C.1.8, we show that if $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)=0$ has a unique solution, then the solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$ is also unique for any $h \in\left[H_{l}, H_{h}\right]$. We use this assumption to characterize the optimal mechanism when the holding cost is medium and large ( $h \geq H_{l}$ ). This assumption ensures that $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta) \leq 0$ for any $\theta<\underline{\theta}_{h}=\underline{\theta}_{M}$. We note that this holds when the virtual value of customers, that is, $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)$ is increasing in the customer type $\theta$. In this sense, this assumption resembles the standard assumption in the standard mechanism design literature where it is assumed that the virtual value of customers is monotone in their types.

In Appendix C.1.8, we provide sufficient conditions to satisfy this assumption. We show that if for any $\theta \leq \tilde{\theta}, \alpha^{\prime}(\theta)$ is small enough, then this assumption holds. The aforementioned condition is satisfied for the Uniform, Exponential, and truncated Normal distributions.

Theorem 3.4.4 shows that a positive holding cost, similar to a positive production cost, introduces a cut-off $\underline{\theta}_{h}$. That is, the mechanism only sells the item to customers with type greater than or equal to $\underline{\theta}_{h}$. However, unlike the production cost, the holding cost changes both the time of sales and the price. Moreover, the thresholds that divide customers into different groups, that is, $\theta_{H}^{h}$ and $\theta_{L}^{h}$, also change with $h$.


Figure 3.11: Utility versus type. Here, $H_{l}=0.004, H_{h}=0.025$, the customer type $\theta \sim U(0,1), V(\theta, t)=$ $\theta e^{-\beta \theta t}, \beta=0.1$, and production cost $c=0$.

Observe that when $h=0$, we can recover the optimal mechanism with no holding cost, as presented in Theorem 3.3.2. To understand why, note that $\theta_{H}^{0}=\theta_{H}, \theta_{L}^{0}=\theta_{L}$, and $\underline{\theta}_{L}^{0}=0$. Furthermore, the FOC solution, $\mathfrak{t}_{f}(\theta)=\mathfrak{t}_{R}(\theta)=\frac{\theta+2 \alpha(\theta)}{\beta \theta \alpha(\theta)}$. When we increase the holding cost from 0 to $\epsilon \approx 0$, the time of sale remains the same for low-type and high-type customers, but medium-type customers purchase the item sooner, see Example 3.4.3.

In addition, the time of sale for low- and high-type customers does not depend on the holding cost. The holding cost only changes the time of sale for medium-type customers and the thresholds that separate customers into different groups; see Figure 3.9.

Next, we discuss the proof of Theorem 3.4.4 and the insights therein. In the proof of Theorem 3.4.4, we show that for any value of $h, e^{-\beta z \mathfrak{t}_{h}(z)}\left(1-\beta z \mathfrak{t}_{h}(z)\right) \geq 0$. Then, considering the fact that $u(\theta, \theta)=$ $\int_{\underline{\theta}_{h}}^{\theta} e^{-\beta z \mathfrak{t}_{h}(z)}\left(1-\beta z \mathfrak{t}_{h}(z)\right) d z$, we can conclude that the utility of customers, $u(\theta, \theta)$, is increasing in $\theta$. Thus, despite the fact that lower-type customers delay their time of purchase to maximize their utility, their utility is nevertheless not as high as higher-type customers.

In the proof, we further show that the purchase time $\mathfrak{t}_{h}(\cdot)$ is decreasing. Thus, lower-type customers purchase the item later. Note that in the optimal mechanism, high-type customers purchase the item immediately while others delay their purchase time.

In Appendix C.1.3, we divide the proof of Theorem 3.4.4 into three lemmas: Lemma C.1.2, C.1.3, and C.1.4. In Lemma C.1.2, C.1.3, and C.1.4, we characterize the optimal mechanism for when the holding cost is low, medium, and high, respectively. Note that the proof of Lemma C.1.2 encompasses the proof of Theorem 3.3.2. In the following, we briefly describe the main idea behind the proof. First we explain the proof for when the holding cost is low.

To characterize the optimal mechanism, by Lemma 3.2.2, we should solve the following optimization problem.

$$
\begin{array}{rr}
\max _{\{\mathfrak{t}(\theta), \mathfrak{p}(\theta): \theta \in[\underline{\theta}, \underline{\Theta}], u(\underline{\Theta}, \underline{\theta}) \geq 0\}} & \mathrm{E}[\max \{R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta), 0\}]-u(\underline{\Theta}, \underline{\Theta}) \\
\text { s.t. } & u(\theta, \theta) \geq u(\theta, \hat{\theta}) \quad \theta, \hat{\theta} \in[\underline{\Theta}, \bar{\Theta}] \\
& u(\theta, \theta) \geq 0 \tag{OPT-H}
\end{array} \quad \theta \in[\underline{\Theta}, \bar{\Theta}]
$$

Here the objective function is the virtual revenue and $R(\theta, \mathfrak{t}(\theta))$ is defined in Eq. (B.14). The first and second sets of constraints ensure that the mechanism is IC and IR, respectively. By Lemma 3.2.1, a mechanism is IC if and only if the interval and envelope conditions hold. This implies that one can replace the IC constraints with these two conditions. However, as we noted in Section 3.3, characterizing the optimal mechanism that satisfies these two conditions is rather complicated. Thus, we relax Problem OPT-H and we only consider the IR constraints and envelope conditions. That is, in the relaxed problem, we need to find an optimal mechanism that satisfies the IR constraints, that is, $u(\theta, \theta)=u(\underline{\Theta}, \underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta z \mathfrak{t}(z)}(1-\beta z \mathfrak{t}(z)) \geq 0$ for any $\theta \in[\underline{\Theta}, \bar{\Theta}]$. Considering that the term inside the integral, that is, $e^{-\beta z t(z)}(1-\beta z \mathfrak{t}(z))$, is not necessarily positive, to fulfill the IR constraints, it is not enough to set $u(\underline{\Theta}, \underline{\Theta}) \geq 0$. Note that in the traditional mechanism design, the utility of a customer is an integral of the probability of allocation, which is always positive. Therefore, the IR constraints are satisfied if we simply set $u(\underline{\Theta}, \underline{\Theta})$ to zero.

To characterize the optimal solution for the relaxed problem, we dualize the IR constraints to construct an upper bound on the relaxed problem. We show that the mechanism described in Theorem 3.4.4 obtains the upper bound, and thus, it is optimal. We then show that the solution of the relaxed problem also satisfies the interval condition. We verify the interval condition is satisfied by showing that the time of allocation $\mathfrak{t}_{h}(\cdot)$ is decreasing and $\left(1-\beta \theta \mathfrak{t}_{h}(\theta)\right) \geq 0$ for any $\theta \geq \underline{\theta}_{h}$. Recall that $u(\theta, \theta)=u(\underline{\Theta}, \underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta z \mathfrak{t}_{h}(z)}(1-$ $\left.\beta z \mathfrak{t}_{h}(z)\right)$. That the interval condition holds suggests that the mechanism given in Theorem 3.4.4 is revenueoptimal.

The proof of Theorem 3.4.4 is similar for the medium holding costs. One of the main challenges there is to show that the mechanism described in Theorem 3.4.4 is IR; that is, we need to verify that $\left(1-\beta \theta t_{h}(\theta)\right) \geq$ 0 for any $\theta \geq \underline{\theta}_{h}$. To this end, we first show that when $h \geq H_{l}$ and $\theta \geq \tilde{\theta}$, we have $1-\beta \mathfrak{t}_{h}(\theta) \theta \geq 0$. Then, through showing $\underline{\theta}_{M} \geq \tilde{\theta}$, we obtain $1-\beta \mathfrak{t}_{h}(\theta) \theta \geq 0$ for any $\theta \geq \underline{\theta}_{M}$.

### 3.4.3 Positive Production and Holding Costs

Up to here, we presented optimal mechanisms when either the holding cost $h$ is zero, or the production cost $c$ is zero. In this section, we characterize an optimal mechanism when both production and holding costs are positive.

The following is the main result of this section. In this theorem, with a slight abuse of notations, we define $\theta_{c} \geq \underline{\theta}_{h}$ as the smallest value that solves $R\left(\theta_{c}, \mathfrak{t}_{h}\left(\theta_{c}\right)\right)-h \mathfrak{t}_{h}\left(\theta_{c}\right)=c$ where $R(\theta, t), \mathfrak{t}_{h}(\theta)$, and $\underline{\theta}_{h}$
are defined in Equations (B.14), (3.8), and (3.7), respectively. We show that a positive production cost does not change the allocation time of customers with type $\theta \geq \theta_{c}$. That is, the firm can price such that customers with type $\theta \geq \theta_{c}$ purchase the item at time $\mathfrak{t}_{h}(\theta)$, and other customers do not purchase the item at all.

Theorem 3.4.5 (Positive Production and Holding Costs). If Assumption 3 holds, $\tilde{\theta}$ is the unique solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)=0$, the holding cost $h \geq 0$, and production cost $c \in[0, \bar{\Theta}]$, then the optimal mechanism sells to customers of type $\theta \geq \theta_{c}$ at time $\mathfrak{t}_{h}(\theta)$, given in Eq. (3.8), and at price

$$
\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{h}(\theta)\right)-\int_{\theta_{c}}^{\theta} e^{-\beta \mathfrak{t}_{h}(z) z}\left(1-\beta \mathfrak{t}_{h}(z) z\right) d z
$$

where $\theta_{c}$ solves $R\left(\theta_{c}, \mathfrak{t}_{h}\left(\theta_{c}\right)\right)-h \mathfrak{t}_{h}\left(\theta_{c}\right)=c$. Furthermore, for $\theta<\theta_{c}, \mathfrak{p}(\theta)=\infty$.

### 3.5 General Valuation Functions

In this section, we present an optimal selling mechanism under the general valuation function $V(\theta, t)=$ $\theta e^{-g(\theta) t}$, where $g(\theta) \geq 0$ is log-concave and increasing. Log-concavity implies that $\frac{g^{\prime}(x)}{g(x)}$ is decreasing. Every positive concave function is log-concave. However, the reverse does not necessarily hold ?. We show that the main insights of Section in 3.3 still hold under general valuation functions.

For convenience, we assume that the production and holding costs are zero. We show that in the optimal mechanism, customers with type $\theta$ purchase the item at time $\mathfrak{t}_{g}(\theta)$, defined below.

$$
\mathfrak{t}_{g}(\theta):=\left\{\begin{array}{lll}
0 & \text { if } \theta \geq \theta_{H}^{g} & \text { High-Type }  \tag{3.9}\\
\frac{g(\theta) \theta+\alpha(\theta)\left(g^{\prime}(\theta) \theta+g(\theta)\right)}{\alpha(\theta) g(\theta) g^{\prime}(\theta) \theta} & \text { if } \theta \in\left[\theta_{L}^{g}, \theta_{H}^{g}\right] & \text { Medium-Type } ; \\
\frac{1}{g^{\prime}(\theta) \theta} & \text { if } \theta \in\left[\underline{\Theta}, \theta_{L}^{g}\right] & \text { Low-Type }
\end{array}\right.
$$

where $\theta_{H}^{g}$ solves $\theta_{H}^{g}+\alpha\left(\theta_{H}^{g}\right)+\frac{\alpha\left(\theta_{H}^{g}\right) g^{\prime}\left(\theta_{H}^{g}\right) \theta_{H}^{g}}{g\left(\theta_{H}^{g}\right)}=0$, and $\theta_{L}^{g}$ solves $g\left(\theta_{L}^{g}\right)+\alpha\left(\theta_{L}^{g}\right) g^{\prime}\left(\theta_{L}^{g}\right)=0$.

The following theorem presents the optimal mechanism under a general valuation function.
Theorem 3.5.1 (General Valuation Functions). Suppose that Assumption 3 holds and the customer valuation function is given by $V(\theta, t)=\theta e^{-g(\theta) t}$, where $g(\cdot) \geq 0$ is non-decreasing and log-concave. Then, given that
time of allocation $\mathfrak{t}_{g}(\theta)$, defined in Eq. (3.9), is non-decreasing, the optimal mechanism sells to customers of type $\theta$ at time $\mathfrak{t}_{g}(\theta)$, and at price $\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{g}(\theta)\right)-\int_{\underline{\Theta}}^{\theta} e^{-g(z) \mathfrak{t}_{g}(z)}\left(1-\mathfrak{t}_{g}(z) g^{\prime}(z) z\right) d z$.

The assumption that $g(\theta)$ is non-decreasing indicates that the customers with higher initial valuations lose interest in the item much faster than customers with lower initial valuations. The log-concavity of $g(\cdot)$ implies that $1-\mathfrak{t}_{g}(\theta) g^{\prime}(\theta) \theta \geq 0$ for any $\theta$; see the proof of Theorem 3.5.1. Given that $1-\mathfrak{t}_{g}(\theta) g^{\prime}(\theta) \theta \geq 0$, we can conclude that the described mechanism is IR; that is $u(\theta, \theta) \geq 0$. This leads to monotonicity of $u(\theta, \theta)$ in $\theta$.

The monotonicity of time of allocation, $\mathfrak{t}_{g}(\cdot)$, is a standard assumption in the mechanism design literature Myerson (1981). We note that $\mathfrak{t}_{g}(\theta)$ is decreasing if $\theta g^{\prime}(\theta)$ and $\alpha(\theta) g^{\prime}(\theta)$ are increasing. These conditions hold when $g(\theta)=\theta^{a}, a \in[0,1]$. Later, in Example 3.5.2, we will show that when $g(\theta)=\theta^{a}$ with $a>1$ and $\theta \sim U(0,1)$, the allocation time $\mathfrak{t}_{g}(\theta)$ is monotone.

The proof of Theorem 3.5.1 is similar to that of Theorem 3.3.2. We start by relaxing the problem by ignoring the interval condition. Then, we show that the solution of the relaxed problem satisfies the interval condition, and is thus optimal.

Again, in the optimal mechanism there exist three groups of customers: high-type, medium-type, and low-type. High-type customers are customers with type $\theta \geq \theta_{H}^{g}$. These customers purchase the item at time zero. Medium-type customers are customers with type $\theta \in\left(\theta_{L}^{g}, \theta_{H}^{g}\right)$, and these customers delay their purchase time and enjoy a positive utility. Low-type customers are customers with type $\theta \leq \theta_{L}^{g}$ who also delay their purchase time, but obtain zero utility.

To get insight into Theorem 3.5.1, we revisit Example 3.3.3.

Example 3.5.2 (Revisiting Example 3.3.3: General Valuation Functions). Here we present the optimal mechanism for the setting described in Example 3.3 .3 when $V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}, a \geq 0$, and $h=c=0$. Note that as " $a$ " increases, the customers' decay rates become more heterogeneous. For this valuation function, thresholds $\theta_{H}^{g}$ and $\theta_{L}^{g}$ are respectively $\frac{a+1}{a+2}$ and $\frac{a}{a+1}$. Note that both $\theta_{H}^{g}$ and $\theta_{L}^{g}$ are increasing in a and converges to 1 when a goes to infinity. Furthermore, at $a=1, \theta_{H}^{g}=\theta_{H}=\frac{2}{3}$ and $\theta_{L}^{g}=\theta_{L}=\frac{1}{2}$. Thresholds $\theta_{H}^{g}$ and $\theta_{L}^{g}$ are depicted in Figure 3.12. We observe that as a increases, the high and medium-type regions shrink while the low-type region expands. Recall that the low-type customers get zero utility. Thus, when heterogeneity among customers increases, i.e., a increases, then the firm can
extract the entire surplus of more customers. As a result, as depicted in Figure 3.13, the revenue of the firm increases when decay rate of customers gets more heterogeneous.

Figure 3.13 shows that the revenue gain of the optimal mechanism relative to the FP policy. We observe that as a increases, the firm increases his revenue to more than $90 \%(a t a=10)$ by employing $D P$. The reason is that by increasing a, the valuation decay rates gets more heterogeneous. This, in turn, increases the value of differentiating customers via using DP.

Figure 3.13 also shows that the social welfare of the customers and firm increases when a increases. Furthermore, for any value of $a, D P$ outperforms FP in terms of obtained social welfare.

Next, we discuss the time of purchase, payment and utility of customers. Given that $V(\theta, t)=\theta e^{-g(\theta) t}$ with $g(\theta)=\theta^{a}$, the time of purchase in the optimal mechanism is given by

$$
\mathfrak{t}_{g}(\theta)= \begin{cases}0 & \text { if } \theta \geq \frac{a+1}{a+2} ;  \tag{3.10}\\ \frac{(a+2) \theta-(a+1)}{a(\theta-1) \theta^{a}} & \text { if } \theta \in\left[\frac{a}{a+1}, \frac{a+1}{a+2}\right] ; \\ \frac{1}{a \theta^{a}} & \text { if } \theta \in\left[0, \frac{a}{a+1}\right],\end{cases}
$$

The time of purchase is shown in Figure 3.14 for $a=0.2,0.6,1$, and 1.4. The figure shows that when a increases, customers with lower type and higher type purchase the item later. Specifically, the time of purchase of low-type customers increases significantly when we increase a from 0.2 to 1.4 while the time of purchase of other customers does not vary remarkably.

Figures 3.15 and 3.16, respectively, show the payment and utility of customers as a function of their types in the optimal mechanism. The utility of customers decreases when a increases. This is so because when a is large, the firm can better differentiate customers. Furthermore, Figures 3.15 shows that by increasing a, only payment of customers with lower and higher types increases.

### 3.6 Conclusion

Dynamic pricing is a common practice in many industries. Such practice has shown to be an effective tool to mitigate negative impact of demand uncertainty. This work contributes to the literature by showing dynamic pricing can be significantly beneficial even in the absence of demand uncertainty. Specifically, we show that


Figure 3.12: The thresholds versus the exponent $a$ for the mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.


Figure 3.13: The social welfare and revenue gain of DP (relative to the FP policy) in percentage as a function of the exponent $a$. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.
when customers' valuations are time-sensitive and decay at different rates, even with deterministic demand, the firms can increase their revenue by deploying dynamic pricing.

The heterogeneity in valuation decay rates enables the firms to differentiate customers effectively, and extract more revenue from them, relative to the case where customers' valuation does not even decay over time. With heterogeneous valuation decay rates, the relative ranking of customers in terms of their valuations vary over time. The change in customers' ranking motivates the firms to revisit their prices as time progresses. This way, they can sell the products to customers that have lower valuation/rank initially. Selling the products to these customers not only can increase the revenue of the firms but also improve the social welfare.

We show that dynamic pricing can be still beneficial when when producing and carrying the products are costly. However, for lower values of production and holding costs, the firms have more flexibility to delay the sales, and as a result, they can gain more from dynamic pricing.


Figure 3.14: Time of purchase versus type in the optimal mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.


Figure 3.15: Payment versus type in the optimal mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.


Figure 3.16: Utility versus type in the optimal mechanism described in Theorem 3.5.1. Here, the customer type $\theta \sim U(0,1), V(\theta, t)=\theta e^{-g(\theta) t}, g(\theta)=\theta^{a}$, and $h=c=0$.

## References

Abraham, I., Athey, S., Babaioff, M. \& Grubb, M. (2011), Peaches, lemons, and cookies: Designing auction markets with dispersed information, Technical report, Tech. Rep. MSRTR-2011-68, Microsoft Research. January.

Acimovic, J. \& Graves, S. (2011), Making better fulfillment decisions on the fly in an online retail environment. Working Paper, MIT.

Agrawal, S., Wang, Z. \& Ye, Y. (2009), A dynamic near-optimal algorithm for online linear programming. Working paper, Stanford University.

Akan, M., Ata, B. \& Dana, J. (2009), 'Revenue management by sequential screening', Unpublished manuscript, Carnegie Mellon University pp. 1792-1811.

Amazon's Recommendation Systems (2012), 'How recommendations work?', http://www.amazon. com/gp/help/customer/display.html/ref=pd_ys_help_iyr?ie=UTF8\&nodeId= 13316081.

Araman, V. F. \& Caldentey, R. (2009), 'Dynamic pricing for nonperishable products with demand learning', Operations research 57(5), 1169-1188.

Athey, S. (2001), 'Single crossing properties and the existence of pure strategy equilibria in games of incomplete information', Econometrica 69(4), 861-889.

Athey, S. \& Levin, J. (1999), Information and competition in us forest service timber auctions, Technical report, National Bureau of Economic Research.

Aviv, Y. \& Pazgal, A. (2008), 'Optimal pricing of seasonal products in the presence of forward-looking consumers', Manufacturing \& Service Operations Management 10(3), 339-359.

Aviv, Y., Wei, M. M. \& Zhang, F. (2015), Responsive pricing of fashion products: The effects of demand learning and strategic consumer behavior, Technical report, Working Paper, Washington University.

Azar, Y., Birnbaum, B. E., Karlin, A. R. \& Nguyen, C. T. (2009), On revenue maximization in secondprice ad auctions, in A. Fiat \& P. Sanders, eds, 'ESA', Vol. 5757 of Lecture Notes in Computer Science, Springer, pp. 155-166.

Ball, M. O. \& Queyranne, M. (2009), ‘Toward robust revenue management: Competitive analysis of online booking', Operations Research 57(4), 950-963.

Bansal, M. \& Maglaras, C. (2009), 'Dynamic pricing when customers strategically time their purchase: Asymptotic optimality of a two-price policy', Journal of Revenue and Pricing management 8(1), 42-66.

Battaglini, M. \& Lamba, R. (2012), Optimal dynamic contracting. Economic Theory Center Working Paper.

Beil, D. R., Chen, Q., Duenyas, I. \& See, B. D. (2015), 'When to deploy test auctions in sourcing', Working Paper.

Bergemann, D. \& Bonatti, A. (2011), 'Targeting in advertising markets: implications for offline versus online media', The RAND Journal of Economics 42(3), 417-443.

Bergemann, D. \& Bonatti, A. (2013), Selling cookies. Working Paper.
Bergemann, D. \& Said, M. (2011a), 'Dynamic auctions', Wiley Encyclopedia of Operations Research and Management Science .

Bergemann, D. \& Said, M. (2011b), 'Dynamic auctions: A survey', Wiley Encyclopedia of Operations Research and Management Science .

Bergemann, D. \& Välimäki, J. (2002), 'Information acquisition and efficient mechanism design', Econometrica 70(3), 1007-1033.

Bernstein, F., Kök, A. G. \& Xie, L. (2011), Dynamic assortment customization with limited inventories. Working Paper, Duke University.

Besanko, D. \& Winston, W. L. (1990), 'Optimal price skimming by a monopolist facing rational consumers', Management Science 36(5), 555-567.

Besbes, O. \& Lobel, I. (2015), 'Intertemporal price discrimination: Structure and computation of optimal policies', Management Science 61(1), 92-110.

Besbes, O. \& Sauré, D. (2012), Dynamic pricing strategies in the presence of demand shocks. Working Paper.

Besbes, O. \& Zeevi, A. (2011), 'On the minimax complexity of pricing in a changing environment', Operations Research 59(1), 66-79.

Bhat, C. R. (2002), 'Recent methodological advances relevant to activity and travel behavior analysis', In Perpetual Motion: Travel Behavior Research Opportunities and Application Challenges, edited by H.S. Mahmassani, Pergamon pp. 381-414.

Bhawalkar, K., Hummel, P. \& Vassilvitskii, S. (2014), Value of targeting, in 'Algorithmic Game Theory', Springer, pp. 194-205.

Bitran, G. \& Caldentey, R. (2003), 'An overview of pricing models for revenue management', Manufacturing \& Service Operations Management 5(3), 203-229.

Bitran, G. R. \& Mondschein, S. V. (1997), 'Periodic pricing of seasonal products in retailing', Management Science 43(1), 64-79.

Board, S. \& Skrzypacz, A. (2016), 'Revenue management with forward-looking buyers', Journal of Political Economy 124(4), 1046-1087.

Boleslavsky, R. \& Said, M. (2013), 'Progressive screening: Long-term contracting with a privately known stochastic process', Review of Economic Studies 80(1), 1-34.

Borgs, C., Candogan, O., Chayes, J., Lobel, I. \& Nazerzadeh, H. (2014), 'Optimal multiperiod pricing with service guarantees', Management Science 60(7), 1792-1811.

Buchbinder, N., Jain, K. \& Naor, J. (2007), Online primal-dual algorithms for maximizing ad-auctions revenue, in L. Arge, M. Hoffmann \& E. Welzl, eds, 'ESA', Vol. 4698 of Lecture Notes in Computer Science, Springer, pp. 253-264.

Buchbinder, N. \& Naor, J. (2007), 'The design of competitive online algorithms via a primal-dual approach', Foundations and Trends in Theoretical Computer Science 3(2-3), 93-263.

Cachon, G. P. \& Swinney, R. (2011), ‘The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior', Management Science 57(4), 778-795.

Chan, C. W. \& Farias, V. F. (2009), 'Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems', Mathematics of Operations Research 34(2), 333-350.

Chan, L. M., Shen, Z. M., Simchi-Levi, D. \& Swann, J. L. (2004), Coordination of pricing and inventory decisions: A survey and classification, in 'Handbook of quantitative supply chain analysis', Springer, pp. 335-392.

Chen, Y. \& Farias, V. F. (2013), ‘Simple policies for dynamic pricing with imperfect forecasts’, Operations Research 61(3), 612-624.

Chen, Y. \& Farias, V. F. (2015), Robust dynamic pricing with strategic customers., in 'EC', p. 777.

Ciocan, D. F. \& Farias, V. F. (2013), Dynamic allocation problems with volatile demand. To appear in Mathematics of Operations Research.

Clifford, S. (2012), 'Shopper alert: Price may drop for you alone', The New York Times, August 9 .

Coase, R. H. (1972), 'Durability and monopoly', The Journal of Law and Economics 15(1), 143-149.

Conlisk, J., Gerstner, E. \& Sobel, J. (1984), 'Cyclic pricing by a durable goods monopolist', The Quarterly Journal of Economics 99(3), 489-505.

Cremer, J. \& Khalil, F. (1992), 'Gathering information before signing a contract', The American Economic Review pp. 566-578.

Crémer, J., Spiegel, Y. \& Zheng, C. Z. (2009), 'Auctions with costly information acquisition', Economic Theory 38(1), 41-72.

Davis, J., Gallego, G. \& Topaloglu, H. (2011), Assortment optimization under variants of the nested logit model. Working Paper, Cornell University.

Devenur, N. R. \& Hayes, T. P. (2009), The adwords problem: online keyword matching with budgeted bidders under random permutations, in 'Proceedings of the 10th ACM conference on Electronic commerce', EC '09, ACM, New York, NY, USA, pp. 71-78.

Elmaghraby, W., Gülcü, A. \& Keskinocak, P. (2008), 'Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands', Manufacturing \& Service Operations Management 10(1), 126-148.

Emek, Y., Feldman, M., Gamzu, I., Leme, R. P. \& Tennenholtz, M. (2012), Signaling schemes for revenue maximization, in B. Faltings, K. Leyton-Brown \& P. Ipeirotis, eds, 'ACM Conference on Electronic Commerce', ACM, pp. 514-531.

Ëso, P. \& Szentes, B. (2007), 'Optimal information disclosure in auctions and the handicap auction', Review of Economic Studies 74(3), 705-731.

Farias, V. F., Jagabathula, S. \& Shah, D. (2011), Assortment optimization under a general choice model. Working Paper, MIT.

Farias, V. F., Jagabathula, S. \& Shah, D. (2013), 'A nonparametric approach to modeling choice with limited data', Management Science 59(2), 305-322.

Federgruen, A. \& Heching, A. (1999), 'Combined pricing and inventory control under uncertainty', Operations research 47(3), 454-475.

Feldman, J., Henzinger, M., Korula, N., Mirrokni, V. \& Stein, C. (2010), 'Online stochastic packing applied to display ad allocation', Algorithms-ESA 2010, Springer pp. 182-194.

Feng, Y. \& Gallego, G. (1995), 'Optimal starting times for end-of-season sales and optimal stopping times for promotional fares', Management Science 41(8), 1371-1391.

Gallego, G., Iyengar, G., Phillips, R. \& Dubey, A. (2004), Managing flexible products on a network. Working Paper, Columbia University.

Gallego, G. \& Topaloglu, H. (2012), Constrained assortment optimization for the nested logit model. Working Paper, Cornell University.

Gallego, G. \& van Ryzin, G. (1994a), 'On the relationship between inventory costs and variety benefits in retail assortments', Management Science 40(8), 999-1020.

Gallego, G. \& Van Ryzin, G. (1994b), 'Optimal dynamic pricing of inventories with stochastic demand over finite horizons', Management science 40(8), 999-1020.

Garrett, D. (2011), 'Durable goods sales with dynamic arrivals and changing values', Unpublished manuscript, Northwestern University .

Gaur, V. \& Honhon, D. (2006), 'Assortment planning and inventory decisions under a locational choice model', Management Science 52(10), 1528-1543.

Ghosh, A., Nazerzadeh, H. \& Sundararajan, M. (2007), Computing optimal bundles for sponsored search, in 'Internet and Network Economics, Third International Workshop (WINE)', pp. 576-583.

Goel, A., Guha, S. \& Munagala, K. (2010), 'How to probe for an extreme value', ACM Transactions on Algorithms 7(1), 12.

Goel, A., Mahdian, M., Nazerzadeh, H. \& Saberi, A. (2010), 'Advertisement allocation for generalized second-pricing schemes', Operations Research Letters 38(6), 571-576.

Golrezaei, N. \& Nazerzadeh, H. (2016), ‘Auctions with dynamic costly information acquisition', Operations Research .

Golrezaei, N., Nazerzadeh, H. \& Rusmevichientong, P. (2014), 'Real-time optimization of personalized assortments', Management Science 60(6), 1532-1551.

Google AdX Documentation (2015), 'Private auctions overview', https://support.google.com/ adxbuyer/answer/2839853?hl=en. Accessed: 2015-04-1.

Google DoubleClick Documentation (2016), 'Distinguish between the open auction, private auctions, and preferred deals', https://support.google.com/dfp_premium/answer/2710459? hl=en. Accessed: 2016-04-1.

Goyal, V., Levi, R. \& Segev, D. (2011), Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand. Working Paper, Columbia University.

Guha, S., Munagala, K. \& Sarkar, S. (2006), Optimizing transmission rate in wireless channels using adaptive probes, in 'Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS/Performance’, pp. 381-382.

Gul, F., Sonnenschein, H. \& Wilson, R. (1986), 'Foundations of dynamic monopoly and the coase conjecture', Journal of Economic Theory 39(1), 155-190.

Hatfield, J. W., Kojima, F. \& Kominers, S. D. (2015), 'Strategy-proofness, investment efficiency, and marginal returns: an equivalence', Working paper .

Helft, M. \& Vega, T. (2010), 'Retargeting ads follow surfers to other sites', The New York Times . URL: http://www.nytimes.com/2010/08/30/technology/30adstalk.html

Hendricks, K., Pinkse, J. \& Porter, R. H. (2003), 'Empirical implications of equilibrium bidding in firstprice, symmetric, common value auctions', The Review of Economic Studies 70(1), 115-145.

Honhon, D., Gaur, V. \& Seshadri, S. (2010), 'Assortment planning and inventory decisions under stock-out based substitution', Operations Research 58(5), 1364-1379.

Hummel, P. \& McAfee, P. (2012), ‘When does improved targeting increase revenue?’.

Jaillet, P. \& Lu, X. (2012), Near-optimal online algorithms for dynamic resource allocations. Working Paper, MIT.

Jasin, S. \& Kumar, S. (2012), 'A re-solving heuristic with bounded revenue loss for network revenue management with customer choice', Mathematics of Operations Research 37(2), 313-345.

Jing, B. (2011), 'Exogenous learning, seller-induced learning, and marketing of durable goods', Management Science 57(10), 1788-1801.

Kakade, S. M., Lobel, I. \& Nazerzadeh, H. (2013), 'Optimal dynamic mechanism design and the virtual pivot mechanism', Operations Research 61(4), 837-854.

Kalyanasundaram, B. \& Pruhs, K. (2000), 'An optimal deterministic algorithm for online b-matching', Theorictal Computer Science 233(1-2), 319-325.

Kök, A. G., Fisher, M. \& Vaidyanathan, R. (2008), Assortment planning: Review of literature and industry practice, in 'Retail Supply Chain Management', Springer.

Kristol, D. M. (2001), 'Http cookies: Standards, privacy, and politics’, ACM Transactions on Internet Technology (TOIT) 1(2), 151-198.

Larson, K. \& Sandholm, T. (2001), Costly valuation computation in auctions, in 'Proceedings of the 8th Conference on Theoretical Aspects of Rationality and Knowledge', TARK '01, pp. 169-182.

Lazear, E. P. (1984), 'Retail pricing and clearance sales’.

Lewis, G. (2011), ‘Asymmetric information, adverse selection and online disclosure: The case of ebay motors', The American Economic Review pp. 1535-1546.

Li, G., Rusmevichientong, P. \& Topaloglu, H. (2013), The d-level nested logit model: Assortment and price optimization problems. Working Paper, Marshall School of Business.

Liu, Q. \& van Ryzin, G. J. (2008a), 'On the choice-based linear programming model for network revenue management', Manufacturing and Service Operations Management 10(2), 288-310.

Liu, Q. \& Van Ryzin, G. J. (2008b), 'Strategic capacity rationing to induce early purchases', Management Science 54(6), 1115-1131.

Lobel, I. \& Xiao, W. (2013), Optimal long-term supply contracts with asymmetric demand information. Working Paper.

Lu, J. \& Ye, L. (2014), 'Optimal two-stage auctions with costly information acquisition', Working Paper .
Mahajan, S. \& van Ryzin, G. J. (2001), ‘Stocking retail assortments under dynamic consumer substitution', Operations Research 49(3), 334-351.

Mahdian, M., Nazerzadeh, H. \& Saberi, A. (2007), Allocating online advertisement space with unreliable estimates, in J. K. MacKie-Mason, D. C. Parkes \& P. Resnick, eds, 'ACM Conference on Electronic Commerce', ACM, pp. 288-294.

Mahdian, M., Nazerzadeh, H. \& Saberi, A. (2012), 'Online optimization with uncertain information', ACM Transactions on Algorithms 8(1), 2.

Mattioli, D. (2012), 'On orbitz, mac users steered to pricier hotels', The Wall Street Journal, June 26.

McAfee, R. P. \& McMillan, J. (1988), ‘Search mechanisms’, Journal of Economic Theory 44(1), 99-123.
Mehta, A., Saberi, A., Vazirani, U. V. \& Vazirani, V. V. (2007), 'Adwords and generalized online matching', Journal of the ACM 54(5).

Milgrom, P. \& Segal, I. (2002), 'Envelope theorems for arbitrary choice sets', Econometrica 70(2), 583-601.
Mirrokni, V. S., Gharan, S. O. \& Zadimoghaddam, M. (2012), Simultaneous approximations for adversarial and stochastic online budgeted allocation, in 'Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms', pp. 1690-1701.

Muthukrishnan, S. (2009), Ad exchanges: Research issues, in 'Internet and network economics', Springer, pp. 1-12.

Myerson, R. (1986), 'Multistage games with communications', Econometrica 54(2), 323-358.
Myerson, R. B. (1981), 'Optimal auction design', Mathematics of operations research 6(1), 58-73.
Pancs, R. (2013), 'Sequential negotiations with costly information acquisition', Games and Economic Behavior 82, 522-543.

Pavan, A., Segal, I. \& Toikka, J. (2014), 'Dynamic mechanism design: A myersonian approach', Econometrica .

Quint, D. \& Hendricks, K. (2012), 'Selecting bidders via non-binding bids when entry is costly', Working Paper.

Rayo, L. \& Segal, I. (2010), 'Optimal information disclosure’, Journal of Political Economy 118(5), 949987.

RFC6265 (2011), 'Http state management mechanism’, http://tools.ietf.org/html/rfc6265.
Roberts, J. W. \& Sweeting, A. (2010), Entry and selection in auctions, Technical report, National Bureau of Economic Research.

Rusmevichientong, P., Shen, Z.-J. M. \& Shmoys, D. B. (2009), 'A PTAS for capacitated sum-of-ratios optimization', Operations Research Letters 37(4), 230-238.

Rusmevichientong, P., Shen, Z.-J. M. \& Shmoys, D. B. (2010), ‘Dynamic assortment optimization with a multinomial logit choice model and capacity constraint', Operations Research 58(6), 1666-1680.

Shen, Z.-J. M. \& Su, X. (2007), 'Customer behavior modeling in revenue management and auctions: A review and new research opportunities', Production and operations management 16(6), 713-728.

Shi, X. (2012), ‘Optimal auctions with information acquisition’, Games and Economic Behavior 74(2), 666686.

Singer, N. (2012), 'Mapping, and sharing, the consumer genome', The New York Times .
URL: http://www.nytimes.com/2012/06/17/technology/acxiom-the-quiet-giant-of-consumer-databasemarketing.html

Smith, S. A. \& Agrawal, N. (2000), 'Management of multi-item retail inventories systems with demand substitution', Operations Research 48, 50-64.

Steel, E. \& Angwin, J. (2010), 'The web’s cutting edge, anonymity in name only', The Wall Street Journal, August 3 .

Stokey, N. L. (1979), 'Intertemporal price discrimination', The Quarterly Journal of Economics pp. 355371.

Su, X. (2007), 'Intertemporal pricing with strategic customer behavior', Management Science 53(5), 726741.

Syrgkanis, V., Kempe, D. \& Tardos, E. (2013), Information asymmetries in common-value auctions with discrete signals. Working Paper.

Szalay, D. (2009), 'Contracts with endogenous information', Games and Economic Behavior 65(2), 586625.

Talluri, K. \& Van Ryzin, G. (2004a), 'Revenue management under a general discrete choice model of consumer behavior', Management Science 50(1), 15-33.

Talluri, K. \& van Ryzin, G. J. (2004b), 'Revenue management under a general discrete choice model of consumer behavior', Management Science 50(1), 15-33.

Talluri, K. \& van Ryzin, G. J. (2004c), The Theory and Practice of Revenue Management, Springer, New York.

Thompson, D. (2012), 'The 11 ways that consumers are hopeless at math', the Atlantic .

Topaloglu, H. (2013), 'Joint stocking and product offer decisions under the multinomial logit model', Production and Operations Management 22(5).

Vallen, M. A. \& Bullinger, C. D. (1999), 'The due diligence process for acquiring and building power plants', The Electricity Journal 12(8), 28-37.
van Ryzin, G. J. \& Mahajan, S. (1999), 'On the relationship between inventory costs and variety benefits in retail assortments', Management Science 45, 1496-1509.

Wang, R. (1993), 'Auctions versus posted-price selling', The American Economic Review pp. 838-851.

Wortham, J. (2012), 'Rather than share your location, Foursquare wants to suggest one', The New York Times, June 7.

Yao, A. (1977), 'Probabilistic computations: Toward a unified measure of complexity', Proceedings of the 18th IEEE Symposium on Foundations of Computer Science (FOCS) pp. 222-227.

Ye, L. (2007), 'Indicative bidding and a theory of two-stage auctions', Games and Economic Behavior 58(1), 181-207.

Yu, M., Debo, L. \& Kapuscinski, R. (2015), ‘Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods', Management Science 62(2), 410-435.

## Appendix A

## Technical Appendix to Chapter 1

## A. 1

Here, we present the proof of Lemma 1.2.1, Theorem 1.3.3, and Proposition 1.4.2.

## A.1.1 Proof of Lemma 1.2.1:

Fix an arbitrary sequence $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ of customers and an algorithm $A$. Let the random variables $S_{1}, \ldots, S_{T}$ denote the sequence of assortments offered by the algorithm $A$, and let $\xi_{1}, \ldots, \xi_{T}$ be the random variables corresponding to the customer's choices. Note that $S_{t}$ may depend on $S_{1}, \ldots, S_{t-1}$ and $\xi_{1}, \ldots, \xi_{t-1}$. The expected revenue of the algorithm $A$ is given by

$$
\mathrm{E}\left[\sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \sum_{i \in S} r_{i} \mathbb{1}\left[S_{t}=S, \xi_{t}=i\right]\right]=\sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \sum_{i \in S} r_{i} \mathrm{E}\left[\phi_{i}^{z_{t}}\left(S_{t}\right) \mathbb{1}\left[S_{t}=S\right]\right]=\sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \sum_{i=1}^{n} r_{i} \phi_{i}^{z_{t}}(S) \operatorname{Pr}\left\{S_{t}=S\right\},
$$

where the first equality follows from the tower property of conditional expectation and the fact that $\mathrm{E}\left[\mathbb{1}\left[\xi_{t}=i\right] \mid S_{t}\right]=\phi_{i}^{z_{t}}\left(S_{t}\right)$, and the last equality follows from the fact that $\mathrm{E}\left[\mathbb{1}\left[S_{t}=S\right]\right]=\operatorname{Pr}\left\{S_{t}=S\right\}$.

For $t=1, \ldots, T$ and $S \in \mathcal{S}$, let $\bar{y}^{t}(S)=\operatorname{Pr}\left\{S_{t}=S\right\}$. To complete the proof, it suffices to show that $\bar{y}^{t}(S)$ satisfies the constraints of the $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$. Clearly, $\bar{y}^{t}(S) \geq 0$. Also, by definition, $\sum_{S \in \mathcal{S}} \bar{y}^{t}(S)=1$. Finally, since the algorithm cannot sell more than the initial inventory of each product, we have that for every product $i$, with probability one,

$$
\sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \mathbb{1}\left[S_{t}=S, \xi_{t}=i\right] \leq c_{i}
$$

and by taking expectation on both sides, we have $\sum_{t=1}^{T} \sum_{S \in \mathcal{S}} \phi_{i}^{z_{t}}(S) \bar{y}^{t}(S) \leq c_{i}$. This shows that $\left(\bar{y}^{t}(S): S \in \mathcal{S}\right)$ is a feasible solution of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$, which is the desired result.

## A.1.2 Proof of Theorem 1.3.3

In this section, we prove Theorem 1.3.3. We start with the following lemma.

Lemma A.1.1. For any increasing, concave, twice-differentiable penalty function $\Psi:[0,1] \rightarrow[0,1]$, the function $x \mapsto \frac{1-x}{2-x-\Psi(x)}$ is increasing on $[0,1]$, and for any $a \in[0,1]$, the function $C \mapsto \frac{1}{C}+\int_{a+\frac{1}{C}}^{1} \Psi(y) d y$ is decreasing on $[1 /(1-a), \infty)$.

Proof. Proof: The derivative of $x \mapsto \frac{1-x}{2-x-\Psi(x)}$ is equal to $\frac{-1+\Psi(x)+(1-x) \Psi^{\prime}(x)}{(2-x-\Psi(x))^{2}}$. The numerator $-1+\Psi(x)+$ $(1-x) \Psi^{\prime}(x)$ is equal to 0 at $x=1$, and its derivative is equal to $(1-x) \Psi^{\prime \prime}(x) \leq 0$ because $\Psi$ is concave. Thus, $-1+\Psi(x)+(1-x) \Psi^{\prime}(x) \geq 0$ for all $x \in[0,1]$. To complete the proof, note that the derivative of $C \mapsto \frac{1}{C}+\int_{a+\frac{1}{C}}^{1} \Psi(y) d y$ is equal to $-(1 / C)^{2}\left(1-\Psi\left(a+\frac{1}{C}\right)\right) \leq 0$.

Let $\left\{z_{t}\right\}_{t=1}^{T}$ be an arbitrary sequence of customers. Note that by Lemma 1.3.2, the Inventory-Balancing algorithm respects the capacity constraints of the problem. However, its solution may not correspond to a feasible solution of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$. In order to compare the expected revenue of our algorithm with the upper-bound given by $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$, we construct a sequence of feasible dual solutions. The dual of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ is given below:

$$
\begin{aligned}
& \text { minimize } \quad \sum_{t=1}^{T} \lambda^{t}+\sum_{i=1}^{n} \theta_{i} c_{i} \\
& \text { (Dual }\left(\left\{z_{t}\right\}_{t=1}^{T}\right) \text { ) } \\
& \text { SUbJECT TO: } \quad \lambda^{t} \geq \sum_{i=1}^{n} \phi_{i}^{z_{t}}(S)\left(r_{i}-\theta_{i}\right) \quad 1 \leq t \leq T, S \in \mathcal{S}, \\
& \theta_{i} \geq 0 \quad 1 \leq i \leq n .
\end{aligned}
$$

Based on the realization of customers' choices, we construct a feasible solution for the linear program Dual $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ as follows:

$$
\begin{aligned}
\theta_{i} & =r_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right) \\
\lambda^{t} & =\sum_{i \in S^{t}} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) \phi_{i}^{z_{t}}\left(S^{t}\right)
\end{aligned} \quad t=1,2, \ldots, n, ~ 子, 2, \ldots, T .
$$

Note that $\theta_{i}$ and $\lambda^{t}$ are random variables because they depend on the inventory levels, which are random. However, they form a feasible solution for the dual, with probability one, because

$$
\lambda^{t}=\sum_{i \in S^{t}} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) \phi_{i}^{z_{t}}\left(S^{t}\right) \geq \sum_{i \in S^{t}} r_{i} \Psi\left(I_{i}^{T} / c_{i}\right) \phi_{i}^{z_{t}}\left(S^{t}\right)=\sum_{i \in S^{t}}\left(r_{i}-\theta_{i}\right) \phi_{i}^{z_{t}}\left(S^{t}\right),
$$

where the inequality follows from the fact that $\Psi$ is increasing and $I_{i}^{t-1} \geq I_{i}^{T}$ an the equality follows from the definition of $\theta_{i}$, that is $r_{i} \Psi\left(I_{i}^{T} / c_{i}\right)=\left(r_{i}-\theta_{i}\right)$.

We now calculate the expected value of this dual solution, which will provide an upper bound on the value of $\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ by the Weak Duality Theorem. Since the sequence of the customers $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ is fixed, the expectation is with respect to the realization of each customer's choice. Recall that $Q_{i}^{t}$ is a binary random variable that is equal to 1 if the customer chooses product $i$ in period $t$, and 0 otherwise. Thus,

$$
\begin{aligned}
\mathrm{E}\left[\sum_{t=1}^{T} \lambda^{t}\right] & =\mathrm{E}\left[\sum_{t=1}^{T} \sum_{S^{\prime} \subseteq S^{t}} \sum_{i \in S^{\prime}} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) \phi^{z_{t}}\left(S^{\prime}, S^{t}\right)\right]=\mathrm{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{n} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) Q_{i}^{t}\right] \\
& =\mathrm{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{n} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right)\left(I_{i}^{t-1}-I_{i}^{t}\right)\right]=\mathrm{E}\left[\sum_{i=1}^{n} r_{i} \sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)\right],
\end{aligned}
$$

where the second equality follows from the tower property of conditional expectation and the fact that $\mathrm{E}\left[Q_{i}^{t} \mid I_{1}^{t-1}, \ldots, I_{n}^{t-1}\right]=\phi_{i}^{z_{t}}\left(S^{t}\right)$, since $S^{t}$ is a function of $I_{1}^{t-1}, \ldots, I_{n}^{t-1}$. The third equality follows from the observation that $I_{i}^{t-1}-I_{i}^{t}=Q_{i}^{t}$. The final equality follows because the $k^{t h}$ sold unit of product $i$ contributes an amount of $\Psi\left(\left(c_{i}-k+1\right) / c_{i}\right)$ to the summation.

Since $\theta_{i}$ and $\lambda^{t}$ are dual feasible, it follows from the Weak Duality Theorem that

$$
\mathrm{E}\left[\sum_{t=1}^{T} \lambda^{t}+\sum_{i=1}^{n} c_{i} \theta_{i}\right]=\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(\sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+c_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right)\right)\right] \geq \operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right) .
$$

On the other hand, the expected revenue of the Inventory-Balancing algorithm is equal to $\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right)\right]$. Therefore, the competitive ratio is at least

$$
\frac{\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right)\right]}{\operatorname{Primal}^{m}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)} \geq \frac{\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right)\right]}{\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(\sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+c_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right)\right)\right]}
$$

Note that if $I_{i}^{T}=c_{i}$, then the contribution of product $i$ to both the revenue of our algorithm and to the constructed dual solution is zero. Therefore, the competitive ratio of the algorithm is at least

$$
\min _{\left(c_{i}, I_{i}^{T}\right): I_{i}^{T} \leq c_{i}-1} \frac{c_{i}-I_{i}^{T}}{\sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+c_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right)}=\min _{\left(c_{i}, x\right): x \leq 1-\frac{1}{c_{i}}} \frac{1-x}{\frac{1}{c_{i}} \sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+(1-\Psi(x))},
$$

where the equality follows from the variable transformation $x=I_{i}^{T} / c_{i}$. Because $\Psi(1)=1$ and $\Psi$ is increasing, we have

$$
\frac{1}{c_{i}} \sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)=\frac{1}{c_{i}}\left(1+\sum_{t=I_{i}^{T}+1}^{c_{i}-1} \Psi\left(t / c_{i}\right)\right) \leq \frac{1}{c_{i}}+\int_{\frac{I_{i}^{T}+1}{c_{i}}}^{1} \Psi(y) d y
$$

Putting everything together, we have the following lower bound on the competitive ratio:

$$
\min _{\left(c_{i}, x\right) \in \mathbb{R}_{+} \times\left[0,1-\frac{1}{c_{i}}\right]}\left\{\frac{1-x}{\frac{1}{c_{i}}+1-\Psi(x)+\int_{x+\frac{1}{c_{i}}}^{1} \Psi(y) d y}\right\}
$$

To complete the proof, it suffices to show that the above ratio is lower bounded by $\alpha_{C_{\text {MIN }}}(\Psi):=$ $\min _{x \in\left[0,1-\frac{1}{c_{\text {MIN }}}\right]} \frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-\Psi(x)+\int_{x+\frac{1}{1}}^{c_{\text {MIN }}}} \Psi(y) d y$ defined in Theorem 1. Consider an arbitrary $\left(c_{i}, x\right) \in \mathbb{R}_{+} \times$ $\left[0,1-\frac{1}{c_{i}}\right]$. There are two cases to consider: $x \leq 1-\frac{1}{c_{\text {MIN }}}$ and $x>1-\frac{1}{c_{\text {MIN }}}$. In the first case, since the function $C \mapsto \frac{1}{C}+\int_{x+\frac{1}{C}}^{1} \Psi(y) d y$ is decreasing by Lemma 3 and $c_{i}>c_{\text {MIN }}$, we have

$$
\frac{1-x}{\frac{1}{c_{i}}+1-\Psi(x)+\int_{x+\frac{1}{c_{i}}}^{1} \Psi(y) d y} \geq \frac{1-x}{\frac{1}{c_{\mathrm{MIN}}}+1-\Psi(x)+\int_{x+\frac{1}{c_{\mathrm{MIN}}}}^{1} \Psi(y) d y} \geq \alpha_{c_{\mathrm{MIN}}}(\Psi) .
$$

In the second case, we have $x>1-\frac{1}{c_{\text {MIN }}}$. Recall that to compute the minimum competitive ratio, we need to consider $I_{i}^{T}<c_{i}$, thus, $\left(c_{i}, x\right) \in \mathbb{R}_{+} \times\left[0,1-\frac{1}{c_{i}}\right]$ and as a result we have $x \leq 1-\frac{1}{c_{i}}$, or equivalently, $c_{i} \geq 1 /(1-x)$. Applying Lemma 3 once again with $\frac{1}{1-x}$ as a lower bound for $c_{i}$, we get

$$
\begin{aligned}
\frac{1-x}{\frac{1}{c_{i}}+1-\Psi(x)+\int_{x+\frac{1}{c_{i}}}^{1} \Psi(y) d y} & \geq \frac{1-x}{1-x+1-\Psi(x)+\int_{x+(1-x)}^{1} \Psi(y) d y}=\frac{1-x}{2-x-\Psi(x)} \\
& \geq \frac{1-\left(1-\frac{1}{c_{\text {MIN }}}\right)}{2-\left(1-\frac{1}{c_{\text {MIN }}}\right)-\Psi\left(1-\frac{1}{c_{\text {MIN }}}\right)} \geq \alpha_{c_{\text {MIN }}}(\Psi),
\end{aligned}
$$

where the second inequality follows from the fact that $x>1-\frac{1}{c_{\text {MIN }}}$ and Lemma A.1.1, which shows that $\frac{1-x}{2-x-\Psi(x)}$ is increasing in $x$. This completes the proof.

## A.1.3 Proof of Proposition 1.4.2

In this section, we characterize the performance guarantee of the EIB algorithm in the I.I.D. arrival model. Precisely, we will look at a random arrival model which encompasses the I.I.D model. ${ }^{1}$ In the random arrival model the total number of the customers $T$ and number of customers of each type are chosen by an adversary, but the arrival order is chosen uniformly at random.

The performance guarantee of the EIB algorithm in the random arrival model is stated in the following proposition.

Proposition A.1.2 (Performance Guarantee in the Random Arrival Model). Suppose the penalty function is exponential (EIB), $\Psi(x)=\frac{e}{e-1}\left(1-e^{-x}\right), x \in[0,1]$. If $c_{\text {MIN }} \rightarrow \infty$ and $\frac{1}{\epsilon \times c_{\mathrm{MIN}}}=O(1)^{2}$, the ratio of the expected revenue of the EIB algorithm, $E_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{E I B}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]$, to the expected revenue of the optimal solution, Primal-S, is bounded below by the solution of the following linear program.

$$
\operatorname{MINIMIZE}_{\rho, \gamma, \chi} \quad \chi\left(\frac{1}{\epsilon}\right)-\frac{e}{e-1} \sum_{j=0}^{\frac{1}{\epsilon}-1} \rho^{j, \frac{1}{\epsilon}}\left(\frac{j \epsilon}{e}-e^{j \epsilon-1}+e^{-1}\right)
$$

Subject To:

$$
\begin{array}{rlrl}
\sum_{j=0}^{\frac{1}{\epsilon}-1} \gamma^{j, k} & =1 & & 0<k \leq \frac{1}{\epsilon}, \\
\gamma^{j, k} & \leq \rho^{j, k} & & 0 \leq j \leq \frac{1-\epsilon}{\epsilon}, 0<k \leq \frac{1}{\epsilon}, \\
\sum_{l=j}^{\frac{1}{\epsilon}-1} \rho^{l, k+1} & \geq \sum_{l=j}^{\frac{1}{\epsilon}-1} \rho^{l, k} & & 0 \leq j \leq \frac{1-\epsilon}{\epsilon}, 0<k \leq \frac{1-\epsilon}{\epsilon}, \\
\sum_{j=0}^{\frac{1}{\epsilon}-1} \rho^{j, k} \int_{y=1-j \epsilon}^{1} \Psi(y) d y & \leq \chi(k) & & 0<k \leq \frac{1}{\epsilon}, \\
\epsilon \sum_{j=0}^{\frac{1}{\epsilon}-1} \gamma^{j, k+1} \Psi(1-(j+1) \epsilon) & \leq \chi(k+1)-\chi(k) & 0<k \leq \frac{1}{\epsilon}-1 .
\end{array}
$$

[^27]${ }^{2}$ We use the standard Landau notation: $f(n)=O(g(n))$ denotes $|f(n)| \leq d|g(n)|$ where $d$ is some constant.
where $\epsilon$ is a discretization parameter.For the linear penalty function (LIB), $\Psi(x)=x$, the ratio of the expected revenue of the Inventory-Balancing algorithm, $E_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{L I B}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]$ to Primal-S is at least equal to the solution of the linear program FRLP with objective function $\chi\left(\frac{1}{\epsilon}\right)+\frac{1}{2} \sum_{j=0}^{\frac{1}{\epsilon}-1} \rho^{j, \frac{1}{\epsilon}}(j \epsilon)^{2}$.

Proposition A.1.2 shows that the expected revenue of the EIB and LIB algorithms is at least equal to Primal-S times the solution of the linear program FRLP. The proof is given in Section A.4.2. The main idea comes from the factor-revealing linear programs by Mirrokni et al. (2012), where we bound the competitive ratio of an algorithm by the solution of a linear program. In particular, we show that the objective function corresponds to the revenue of the IB algorithms when the value of Primal-S is normalized to one and any solution of the IB algorithms corresponds to a feasible solution of FRLP. The revenue of the IB algorithms is evaluated by potential function $\chi$, which is an integral of the penalty function; see the precise definition in Section A.1.2. That is, the IB algorithms perform well if they can maximize the final value of potential function, $\chi\left(\frac{1}{\epsilon}\right)$. Thus, we need to monitor the increase in the potential function during the run of the algorithms. ${ }^{3}$ Interestingly, the last set of constraints implies that the potential function has a significant increase over time, with high probability. Hence, we can lower bound the final value $\chi\left(\frac{1}{\epsilon}\right)$ or equivalently the overall performance of the IB algorithms, as a function of the optimal solution, measured by $\gamma_{j, k}$ s.

As stated, $\epsilon$ is a discretization parameter, that is, the horizon $T$ is divided into $\frac{1}{\epsilon}$ time slots, each with an integer length $\epsilon T$. Considering the fact that $\epsilon$ is a discretization parameter, we prefer to choose $\epsilon$ as small as possible. Particularly, as $\epsilon \rightarrow 0, T$ should approach infinity $(T \rightarrow \infty)$. Note that we have assumed $\frac{1}{\epsilon \times c_{\text {MIN }}}=O(1)$. Thus depending on how $\epsilon T$ scales, we can determine how fast $T$ increases toward infinity, e.g., given that $\epsilon T$ is a finite integer, $T=O\left(c_{\text {min }}\right)$.

The optimal value of the linear program FRLP for the LIB algorithm for $\epsilon=\frac{1}{20}, \frac{1}{30}$ and $\frac{1}{50}$ are respectively $0.69,0.71$, and 0.72 . The corresponding values for the EIB algorithm are respectively $0.72,0.74$, and 0.75. We observe that as $\epsilon$ gets smaller, the optimal value of FRLP becomes more accurate and converges to 0.72 and 0.75 respectively for the LIB and EIB algorithms. Considering the fact that random arrival model encompasses the I.I.D. model, we get the performance guarantee of the LIB and EIB algorithms in the I.I.D. model given in Proposition 1.4.2.

[^28]
## A.1.4 Learning the Customer Types under the Multinomial Logit Model

Suppose the choice model of each customer type is described by a multinomial logit. If we show all products to $m$ independent customers of type $z$ and compute the maximum likelihood estimates $\left(V_{0}^{z}, \ldots, V_{n}^{z}\right)$, then it is a standard result that $\operatorname{Pr}\left\{\max _{i, S}\left|\phi_{i}^{z}(S)-\bar{\phi}_{i}^{z}(S)\right|>\delta\right\} \leq d e^{-\delta^{2} m}$. Here, $\phi_{i}^{z}(S)=V_{i}^{z} /\left(V_{0}^{z}+\sum_{\ell \in S} V_{\ell}^{z}\right)$ and $d$ is a constant, see, for example, Rusmevichientong et al. 2010. Therefore, $\mathrm{E}\left[\max _{i, S}\left|\phi_{i}^{z}(S)-\bar{\phi}_{i}^{z}(S)\right|\right] \leq \delta+d e^{-\delta^{2} m}$. Now consider the following variation of the IB algorithm.

Upon the arrival of the customer in period $1 \leq t \leq T$, of type $z_{t}$, with probability $0<\gamma<1$, we do explore, i.e., we show all the products to the customer and with probability $1-\gamma$, we offer an assortment $S^{t}=\arg \max _{S \in \mathcal{S}} \sum_{i \in S} \Psi\left(I_{i}^{t-1} / c_{i}\right) r_{i} \phi_{i}^{z_{t}}(S)$, where $\phi_{i}^{z_{t}}(S)$ is the estimated selection probability, as described above, using previous sales data.

Note that the number of observations up to period $t$, with high probability, is approximately $\theta(\gamma t)$. Hence, by setting $\delta=\frac{1}{t^{\left(1-\delta_{1}\right) / 2}}$ and $m=\gamma t$ where $0<\delta_{1}<1$, we have $\mathrm{E}\left[\max _{i, S}\left|\phi_{i}^{z_{t}}(S)-\bar{\phi}_{i}^{z_{t}}(S)\right|\right]=$ $O\left(\frac{1}{t^{\left(1-\delta_{1}\right) / 2}}+e^{-\gamma t^{\delta_{1}}}\right)$, which implies that $\mathrm{E}\left[\sum_{t=1}^{T} \epsilon^{t}\right]=o(T) .{ }^{4}$ Observe that as $c_{\text {MIN }}$ and $T$ proportionally grow, $\mathrm{E}\left[\sum_{t=1}^{T} \epsilon^{t}\right] / c_{\text {Min }}$ approaches 0 . By Proposition 1.5.1 and the fact that algorithm loses at most $\gamma$ fraction of its revenue during explorations, ${ }^{5}$ the competitive ratio of the modified algorithm, as $c_{\text {MIN }}$ and $T$ proportionally tend to infinity, would be equal to $(1-\gamma) \alpha(\Psi)$.

Note that the modified algorithm, because of the constant rate of sampling, would still perform well if the choice models change slowly over time.

## A. 2 Computational Complexity of Inventory-Balancing Algorithms

The complexity of the Inventory-Balancing Algorithm is measured by the complexity the following combinatorial optimization problem $\max _{S \in \mathcal{S}} \sum_{i \in S} r_{i} \Psi\left(I^{t-1} / c_{i}\right) \phi_{i}^{z t}(S)$, where $r_{i} \Psi\left(I^{t-1} / c_{i}\right)$ is the discounted revenue of product $i$ in period $t$. As shown in the following examples, for a broad class of choice models this

$$
{ }^{4} f(x)=o(g(x)) \text { if } \lim _{x \rightarrow \infty} \frac{f(x)}{g(x)}=0 .
$$

${ }^{5} \mathrm{We}$ allocate $\gamma$ fraction of the inventory of each product for "exploration".
problem can be solved efficiently. Also, all of the choice models in these examples satisfy Assumption 1. We note that in Example A.2.1, no constraints on the assortments are included.

Example A.2.1 (Multinomial and Nested Logit). Under the multinomial logit (MNL) choice model,

$$
\phi_{i}^{z}(S)= \begin{cases}v_{i}^{z} /\left(v_{0}^{z}+\sum_{\ell \in S} v_{\ell}^{z}\right) & \text { if } \quad i \in S \cup\{0\}, \\ 0 & \text { otherwise }\end{cases}
$$

where for each $i \in\{0,1,2, \ldots, n\}, v_{i}^{z} \in \mathbb{R}_{+}$denote the preference weight parameter associated with product i for a customer type $z$. As shown in Talluri \& van Ryzin (2004b), the assortment optimization problem $\max _{S} \sum_{i \in S} w_{i} \phi_{i}^{z}(S)$ can be solved efficiently by simply sorting the products in a descending order of the marginal revenue $w_{i}$, and the optimal assortment can be found among the revenue-ordered assortments $\{1\},\{1,2\}, \ldots,\{1,2, \ldots, n\}$. Li et al. (2013) showed that the above assortment optimization problem can also be solved in $O(d n \log n)$ operations for a d-level nested logit choice model, which generalizes to $M N L$ to allow for a product taxonomy that is described by a tree with d levels (the MNL corresponds to the special case where $d=1$ ).

Example A.2.2 (Choice Models with Constraints). Our formulation automatically allows for constraints on the assortments, through the specification of the set $\mathcal{S}$ of feasible assortments. For example, if we have a budget or shelf-space constraint, we can define $\mathcal{S}=\left\{S: \sum_{i \in S} d_{i} \leq B\right\}$, where $d_{i}$ is the cost (or space) associated with showing product $i$ and $B$ is the budget. When the underlying choice model is an MNL or a nested logit model, the assortment optimization problem can still be solved efficiently; see, for example, Rusmevichientong et al. (2009, 2010), Gallego \& Topaloglu (2012).

Example A.2.3 (General Choice Models). Farias et al. (2013) has pioneered a novel algorithm that learns a general nonparametric choice model from transaction data. Under their framework, for each customer of type $z$, the choice model corresponds to a probability distribution $\eta^{z}$ over the set of all permutations of $\{1, \ldots, n\}$ where for each $i, \sigma(i) \in\{1,2, \ldots, n\}$ denote the rank of product $i$, with 1 be the highest rank (most preferred). They assume that customer chooses the product with the highest rank. In this case,

$$
\phi_{i}^{z}(S)=\sum_{\sigma} \eta^{z}(\sigma) \mathbb{1}\left[\sigma(i) \leq \min _{k \in S \cup\{0\}} \sigma(k)\right] .
$$

It is easy to verify that $\phi_{i}^{z}(S) \geq \phi_{i}^{z}(S \cup\{j\})$ for $j \neq i$. So, this choice model satisfies Assumption 1. Moreover, Farias et al. (2011) described an algorithm for solving the assortment optimization problem
$\max _{S} \sum_{i \in S} w_{i} \phi_{i}^{z}(S)$ under their general choice model framework. This algorithm which can be used as a subroutine in our Inventory-Balancing algorithm is efficient for MNL choice model.

## A. 3 Numerical Experiments: Appendix to Section 1.6

## A.3.1 Known Length of the Horizon

In this Section we compare the performance of EIB, LIB and Hybrid algorithms to the Myopic policy and the LP-based heuristics when the length of the horizon is known in advance. We set the initial inventory levels to 100 , i.e., $c_{i}=100, i=1,2, \ldots, 73$.

Performance Evaluation: In Table A.1, we present the average revenue of each algorithm as a percentage of the upper bound, which is averaged over all 250 problem instances, for loading factors 1.4, 1.6 and 1.8 and for coefficients of variation of $0.1,1$, and 2 . As the table shows, when the number of customers is known in advance, $\mathrm{LPR}^{500}$ algorithm can obtain more that $99 \%$ of the optimal solution for all the considered problem classes which implies that having more resolving periods is not necessary.

| Problem Class |  | $\begin{gathered} \text { Upper } \\ \text { Bound } \\ \text { (in } \$ 1000 \text { ) } \end{gathered}$ | Avg. Revenue Under Different Policies (as \% of the Upper Bound) |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Inventory-Balancing | Myopic Policy | One-shot LP |  | $\frac{\text { LP Resolving }}{\text { LPR }^{500}}$ | Hybrid |  |
| LF | CV |  |  | EIB | LIB |  | LPO | ALPO | $\mathcal{H}_{1.5}$ | $\mathcal{H}_{2}$ |
| 1.4 | 2.0 |  | 173 | 97.3 | 97.2 | 96.2 | 69.3 | 77.9 | 99.3 | 98.6 | 98.9 |
|  | 1.0 | 179 | 97.6 | 97.8 | 95.4 | 83.5 | 89.2 | 99.3 | 98.5 | 98.8 |
|  | 0.1 | 182 | 98.1 | 98.5 | 95.8 | 95.8 | 98.0 | 99.5 | 98.9 | 99.1 |
| 1.6 | 2.0 | 175 | 98.2 | 98.3 | 97.2 | 68.4 | 75.4 | 99.4 | 98.9 | 99.0 |
|  | 1.0 | 181 | 98.7 | 98.9 | 97.4 | 83.6 | 88.3 | 99.6 | 99.0 | 99.0 |
|  | 0.1 | 183 | 99.3 | 99.3 | 97.8 | 95.8 | 97.8 | 99.7 | 99.4 | 99.5 |
| 1.8 | 2.0 | 177 | 98.8 | 98.9 | 98.0 | 65.5 | 72.1 | 99.4 | 99.1 | 99.1 |
|  | 1.0 | 182 | 99.2 | 99.3 | 98.4 | 79.8 | 84.6 | 99.5 | 99.3 | 99.4 |
|  | 0.1 | 183 | 99.7 | 99.8 | 99.4 | 95.5 | 97.6 | 99.8 | 99.8 | 99.8 |

Table A.1: Revenue comparison when the length of horizon is known in advance. The standard errors of all numbers are less than $0.1 \%$.

We note that both the LIB and EIB algorithms outperform Myopic and LPO algorithms. Moreover, the revenue of the EIB and LIB algorithms is within $\pm 2 \%$ of that of the resolving heuristics. Comparing the performance of $\mathrm{LPR}^{500}$ in Tables A. 1 and 1.3 implies that that LPR heuristics are sensitive to uncertainty
in number of customers. Precisely, the performance of LPR heuristic decreases significantly when it does not know the exact length of the horizon. In all problem classes, the Hybrid algorithms yield more revenue than the IB polices since they incorporates additional information about arrival sequence by using the LP resolving heuristic.

Again, in all cases, the LPO algorithm has the lowest revenue and its performance decreases by increasing CV and loading factor. Observe that when $\mathrm{CV}=0.1$, One-shot LP heuristics obtain more than $95 \%$ of the optimal clairvoyant solution. For small value of CV, the number of customers of each type is very concentrated around its average. Therefore, these heuristics do not suffer from fixing their strategies at the beginning of the horizon. Note that even for small value of CV, our IB algorithms perform better than One-shot LP heuristics.

Transient Behavior: Figures A. 1 shows the cumulative revenue over time for Myopic, LIB, LPR $^{500}$, LPO and ALPO algorithms with LF $=1.8$ and $\mathrm{CV}=2$. We observe that the Myopic policy and the LIB algorithm are very aggressive during the initial periods, resulting in higher cumulative revenues than One-shot LP and resolving heuristics. Since resolving heuristics know exactly the number of customers in advance, they manage to earn revenue linearly over time. This implies that knowing the true estimate of the length of the horizon (number of customers) is essential for the resolving heuristics, that is, if the number of customers is less than its estimated value, these heuristics will suffer from significant revenue loss, see Section 1.6.2.

## A.3.2 Worst-Case Performance

In Section 1.6.2, we have compared different polices in term of their average performance. Here, we investigate the worst-case performance of different policies. To this aim, we consider 250 random arrival sequences. For each of them we compute the ratio of revenue collected by each policy and the corresponding optimal clairvoyant solution. Then, the worst-case performance of any policy is defined as the minimum of these ratios. Table A. 2 presents the worst-case performance of all policies for $\mathrm{LF}=1.4,1.6$, and 1.8 and $\mathrm{CV}=0.1,1$ and 2 when the length of the horizon is drawn from the uniform distribution with $\bar{T}-\underline{T}=\mathrm{E}[T]$. Our IB polices outperform other policies in term of worst-case performance, that is they can obtain at least $91 \%$ of the optimal clairvoyant solution, which is much higher than the theoretical bounds, i.e., $63 \%$ for the EIB policy and $50 \%$ for the LIB policy. We observe that the LP resolving heuristics perform poorly compare


Figure A.1: The cumulative revenue over time for $L F=1.8$ and $C V=2$ when the length of the horizon is known in advance.
to IB and Hybrid algorithms. Furthermore, One-shot LP heuristics are very sensitive to uncertainty in arrival sequence (large CV). For instance, when $\mathrm{LF}=1.8$ and $\mathrm{CV}=2$ there is an arrival sequence in which they only get $3.8 \%$ of the optimal clairvoyant solution.

| Problem Class |  | Worst Case Revenue Under Different Policies (as \% of the Upper Bound) |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Inventory-Balancing |  | Myopic <br> Policy | One-shot LP |  | LP Resolving |  | hybrid |  |
| LF | CV | EIB | LIB |  | LPO | ALPO | $\mathrm{LPR}^{500}$ | $\mathrm{LPR}^{50}$ | $\mathcal{H}_{1.5}$ | $\mathcal{H}_{2}$ |
| 1.4 | 2.0 | 91.8 | 91.4 | 91.0 | 13.2 | 14.4 | 75.8 | 76.7 | 88.4 | 83.2 |
|  | 1.0 | 92.2 | 92.0 | 91.9 | 16.8 | 17.8 | 76.7 | 77.4 | 88.8 | 84.7 |
|  | 0.1 | 92.2 | 91.8 | 91.4 | 72.6 | 73.3 | 78.9 | 79.1 | 89.5 | 85.4 |
| 1.6 | 2.0 | 92.5 | 92.0 | 92.0 | 8.6 | 8.8 | 70.9 | 73.3 | 90.0 | 86.9 |
|  | 1.0 | 93.2 | 91.7 | 91.2 | 41.8 | 41.9 | 72.7 | 73.1 | 89.8 | 87.3 |
|  | 0.1 | 92.7 | 92.8 | 91.3 | 66.5 | 67.4 | 73.4 | 74.5 | 91.1 | 87.3 |
| 1.8 | 2.0 | 92.4 | 92.3 | 91.2 | 3.8 | 3.8 | 66.2 | 67.4 | 90.5 | 87.1 |
|  | 1.0 | 92.8 | 92.5 | 92.0 | 20.9 | 20.3 | 68.4 | 69.0 | 91.8 | 88.9 |
|  | 0.1 | 93.1 | 93.2 | 91.6 | 60.3 | 60.5 | 67.8 | 68.2 | 92.6 | 90.3 |

Table A.2: Worst-Case Performance Comparison when the length of horizon is unknown.

## A.3.3 Learning the Customer Types

Here we investigate the performance of the IB algorithms when we do not know the exact value of the selection probability $\bar{\phi}_{i}^{z_{t}}(S)$. Rather, we only have an estimate $\phi_{i}^{z}(S)$ based on data collected in the previous
periods. Since we assume the multinomial logit choice model for each customer type, we maintain an estimate $\mathbf{V}^{\mathbf{z}}(t)=\left(V_{0}^{z}(t), V_{1}^{z}(t), \ldots, V_{n}^{z}(t)\right)$ of the preference weight parameters, where for each product $i$, we set $V_{i}^{z}(t)$ to be proportional to the number of times that a customer of type $z$ purchases DVD during the previous $t-1$ periods, and we normalize $\mathbf{V}^{\mathbf{z}}(t)$ so that $V_{0}^{z}(t)=1$. Similar to the previous section, we have 10 customer types and 73 products with initial inventory of $c_{i}=30$.

| Problem Class |  | Upper <br> Bound | Revenue <br> (as \% of the Upper Bound) |  |
| :---: | :---: | :---: | :---: | :---: |
| Loading <br> Factor (LF) | Coefficient of <br> Variation (CV) |  |  |  |
| 1.2 | 0.2 | 526 | 91.3 | 91.4 |
|  | 0.5 | 524 | 89.6 | 89.6 |
|  | 0.8 | 522 | 86.2 | 86.4 |
| 1.4 | 0.2 | 546 | 95.4 | 95.7 |
|  | 0.5 | 544 | 92.9 | 93.2 |
|  | 0.8 | 541 | 89.9 | 90.6 |
|  | 0.2 | 549 | 98.6 | 98.5 |
|  | 0.5 | 548 | 96.7 | 96.9 |
|  | 0.8 | 546 | 93.8 | 93.9 |

Table A.3: The average revenue for the LIB and EIB algorithms when the underlying parameters are unknown, and each algorithm uses the estimated parameters based on data collected in the previous periods.

Table A. 3 shows the revenue of the IB algorithms when these algorithms only have estimates of the preference weight parameters. In absolute terms, the IB algorithms perform well despite not knowing the true parameter values; they obtain $83 \%-98 \%$ of the upper bound, depending of the coefficient of variations and the loading factor. We observe better performance for loading factor of 1.6 in compare to smaller loading factors. The reason is that larger loading factor or longer the horizon allows the algorithms to obtain better estimates of the unknown parameters. Note that IB algorithms perform well even with few observations. One of the reasons is that in the setting above, we do not impose any constraint on the size of the assortment that policies can offer to each customer. This could compensate for the inaccuracy in estimation of choice model since the algorithm can offer large assortments. Furthermore, by Proposition 1.5.1, we expect the IB algorithms to be robust with respect to the preference weight parameters.

## A. 4 Relegated Proofs

Proof. Proof of Lemma 1.3.2: When the customer arrives in period $t$, if product $j$ has no remaining inventory, then $I_{j}^{t-1}=0$, which implies that $r_{j} \Psi\left(I_{j}^{t-1} / c_{j}\right)=0$. By Assumption 1, under our choice model, adding product $j$ to an assortment does not increase the probability that a customer will select other products. Recall that in the case that both sets $S$ and $S \cup\{j\}$ have the maximum discounted revenue, we choose the set with the smaller number of products. Therefore, product $j$ will never be included as a part of the optimal assortment.

Proof. Proof of Corollary 1.3.4: First, observe that $\Psi(x) \geq x$. This is because $\Psi$ is increasing and concave, and we have $\Psi(0)=0$ and $\Psi(1)=1$. By this observation, we have

$$
\begin{aligned}
\alpha_{\Psi}\left(c_{\text {MIN }}\right) & \geq \min _{x \in\left[0,1-\frac{1}{c_{\text {MIN }}}\right]}\left\{\frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-\Psi(x)+\int_{x+\frac{1}{c_{\text {MIN }}}}^{1} \Psi(y) d y}\right\} \\
& \geq \min _{x \in\left[0,1-\frac{1}{c_{\text {MIN }}}\right]}\left\{\frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-x+\int_{x+\frac{1}{c_{\text {MIN }}}}^{1} d y}\right\} \\
& \geq \min _{x \in\left[0,1-\frac{1}{c_{\text {MIN }}}\right]}\left\{\frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-x+1-x-\frac{1}{c_{\text {MIN }}}}\right\} \\
& \geq \frac{1}{2} .
\end{aligned}
$$

The second equality follows from the fact that for any $x \in\left[0,1-\frac{1}{c_{\text {MIN }}}\right]$, the lower limit of integral, $x+\frac{1}{c_{\text {MIN }}}$, is less than the upper limit of integral, 1 , and $\Psi(x) \leq 1$.

In the following we show that the competitive ratio of the IB algorithm with an increasing, strictly concave, and differentiable penalty function is strictly greater than $\frac{1}{2}$. First note that for $x=0$, the lower bound of the competitive ratio, $\left\{\frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-\Psi(x)+\int_{x+\frac{1}{1}}^{c_{\text {MIN }}} \Psi(y) d y}\right\}$ is greater than $\frac{1}{2}$. This is because for a differentiable penally function $\Psi, \int_{\frac{1}{c_{\text {MIN }}}}^{1} \Psi(y) d y$ is strictly less than $1-\frac{1}{c_{\text {MIN }}}$ and $\Psi(0)=0$. Thus, the result holds because
$\min _{x \in\left(0,1-\frac{1}{c_{\text {MIN }}}\right.}\left\{\frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-\Psi(x)+\int_{x+\frac{1}{c_{\text {MIN }}}}^{1} \Psi(y) d y}\right\}>\min _{x \in\left(0,1-\frac{1}{c_{\text {MIN }}}\right]}\left\{\frac{1-x}{\frac{1}{c_{\text {MIN }}}+1-x+\int_{x+\frac{1}{c_{\text {MIN }}}}^{1} d y}\right\}=\frac{1}{2}$.

The inequality holds because for any differentiable and strictly concave penalty function $\Psi(x)$, we have $\Psi(x)>x$ for all $x \in(0,1)$.

## A.4.1 Appendix to Section 1.3.1

Proof. Proof of Theorem 1.3.6: From Theorem 1 we have:

$$
\begin{aligned}
\alpha(\Psi) & =\min _{x \in[0,1]}\left\{\frac{1-x}{1-\frac{e}{e-1}\left(1-e^{-x}-\int_{x}^{1}\left(1-e^{-y}\right) d y\right)}\right\} \\
& =\min _{x \in[0,1]}\left\{\frac{1-x}{1-\frac{e}{e-1}\left(1-e^{-x}-1+x-e^{-1}+e^{-x}\right)}\right\} \\
& =\min _{x \in[0,1]}\left\{\frac{1-x}{1-\frac{e}{e-1}\left(x-e^{-1}\right)}\right\}=\frac{e-1}{e}=1-\frac{1}{e} .
\end{aligned}
$$

The second part of the theorem is followed from Lemma 1.3.5.

Proof. Proof of Lemma 1.3.5: ${ }^{6}$ Consider a setting with $n$ products, indexed by $1, \cdots, n$, all with revenue equal to 1 and initial inventory of $\frac{1}{n} T$. Think of $T$, the length of the horizon, as a very large number (that would tend to infinity) and a multiple of $n$. The number of types is equal to $2^{n}-1$. Each type corresponds to a set $\Theta \neq \emptyset$ of products that a customer of that type equally likes; the "no-purchase" probability for all types is equal to zero.

The arrival process is defined as follows: customer arrives in $n$ phases of equal length, that is, the number of customers in each phase is $\frac{T}{n}$. All the customers in each phase have the same type. We denote the type of the customer in phase $j$ by $\Theta_{j}$. We have $\Theta_{1}=\{1,2, \ldots n\}$; for $j, 2 \leq j \leq n, \Theta_{j}=\Theta_{j-1} \backslash\left\{\theta_{j-1}\right\}$ where $\theta_{j-1}$ is a randomly chosen element of $\Theta_{j-1}$. In other words, the set of products of interest to customer during phase $j$ is the set of products of interest to customers in phase $j-1$ minus one of those products and $\theta_{n}$ is the only product of interest to customers in phase $n$, i.e., customers in phase $j$ randomly lose interest in one of the products of interest in phase $j-1$. An example of sequences of customer types in $n$ phases is

[^29]$\{\{1,2, \ldots, n\},\{1,2, \ldots, n-1\}, \ldots,\{1,2\},\{1\}\}$. Therefore, there are $n!$ sequences of customer arrivals, each with equal probability.

In Lemma A.4.1 in Section A.4, we show that the following Inventory-Balancing policy is optimal among all deterministic policies: offer to each customer all the products with the highest (positive) remaining inventory that are of interest to her. ${ }^{7}$

Each customer purchases one of the products (if any) offered to her because the no-purchase probability is zero. Hence, the policy described above, in each phase, sells equal portion of the remaining inventory of each product that is of interest to the customers in that phase (which are all of the same type). For instance, in the first phase $\frac{1}{n}$ fraction of the inventory of every product is sold. Note that the rounding error is negligible since $T$ is large. Recall that $\theta_{i}$ denotes the product that will be of no more interest to the customers arriving after and including phase $i+1$. Let $q_{i, j}$ be the fraction of customers in phase $j$ that bought product $\theta_{i}$. We have

$$
q_{i, j}= \begin{cases}\frac{1}{n-j+1} & j \leq i  \tag{A.1}\\ 0 & j>i\end{cases}
$$

where $n-j+1$ is the number of products of interest to customers in phase $j$. Therefore, the revenue obtained from product $\theta_{i}$ is $\frac{1}{n} T\left(\min \left\{\sum_{j=1}^{i} \frac{1}{n-j+1}, 1\right\}\right)$ and consequently the total revenue of the policy above is equal to $\frac{1}{n} T\left(\sum_{i=1}^{n} \min \left\{\sum_{j=1}^{i} \frac{1}{n-j+1}, 1\right\}\right)$. On the other hand, the optimal clairvoyant solution that knows the customers types in advance sells all units of product $\theta_{i}$ to customers in phase $i$ and obtains in total a revenue of $T$. This completes the proof.

Lemma A.4.1. For the arrival process described in the proof of Lemma 1.3.5, the following inventory balancing algorithm is optimal among all deterministic policies: offer to the customer all the products with the highest (positive) remaining inventory that are of interest to her.

Proof. Proof: Since the revenue from all the products is the same and no-purchase probability is zero, by Eq. (A.7), to prove the lemma, it suffices to show the following Claim.

[^30]Claim 1: Consider any two products $i$ and $j$ and remaining inventory levels $\left(x_{1}, \ldots, x_{n}\right)$ such that $x_{i}>$ $x_{j}$. If with $t$ periods remaining, the type of the arriving customer is $z$ such that $i \in z$ and $j \in z$, then we have

$$
V\left(t, x_{1}, \ldots, x_{i}, \ldots, x_{j}, \ldots, x_{n} \mid z\right) \leq V\left(t, x_{1}, \ldots, x_{i}-1, \ldots, x_{j}+1, \ldots, x_{n} \mid z\right)
$$

This claim implies that it would be better to equalize the inventory levels. Namely, if the inventory for product $i$ is higher than product $j$, the value (i.e., expected revenue) of the DP policy would increase if instead we have one additional unit of product $j$ and one less unit of product $i$.

We prove the claim using induction on the inventory levels, fixing product (any) two products $i$ and $j$ :
The induction basis is when $x_{i}=x_{j}+1$ (and no restriction on the inventory of other products). In this case, because of the symmetry in the problem, the value function does not change if we replace one unit of product $i$ with one unit of product $j$. The reason is that the current customer in interested in both $i$ and $j$ and because of the symmetry in the arrival process, the probability that a future customer is only interested in product $i$ but not product $j$ is the same as the probability that a future customer is only interested in product $j$ but not product $i$.

Induction Step: Consider initial inventory levels $\left(y_{1}, \ldots, y_{n}\right)$ such that $y_{i}>y_{j}+1$. Assume Claim 1 holds for any other initial inventory levels $\left(x_{1}, \ldots, x_{n}\right)$ such $x_{k} \leq y_{k}, 1 \leq k \leq n$, and at least one of these inequalities is strict. To prove the induction step, suppose the optimal dynamic program starting with inventory levels $\left(y_{1}, y_{2}, \ldots, y_{n}\right)$ offers set $S$ to the arriving customer. Hence, by conditioning on the type of the customer in the next period, denoted by $z^{\prime}$, we have:

$$
\begin{aligned}
& V\left(t, y_{1}, \ldots, y_{n} \mid z\right) \\
= & \sum_{z^{\prime} \in \mathcal{Z}} \operatorname{Pr}\left[\text { next customer is of type } z^{\prime} \mid z\right] \times\left(1+\frac{1}{|S|} \sum_{k \in S} V\left(t-1, y_{k}-1, y_{-\{k\}} \mid z^{\prime}\right)\right)
\end{aligned}
$$

where $\left(y_{k}-1, y_{-\{k\}}\right)$ represents the same inventory levels as before only with one less unit of product $k$. Now by applying the induction hypothesis, we get

$$
\begin{aligned}
& V\left(t, y_{1}, \ldots, y_{n} \mid z\right) \\
\leq & \sum_{z^{\prime} \in \mathcal{Z}} \operatorname{Pr}\left[\text { next customer is of type } z^{\prime} \mid z\right] \times\left(1+\frac{1}{|S|} \sum_{k \in S} V\left(t-1, y_{k}-1, y_{i}-1, y_{j}+1, y_{-\{i, j, k\}} \mid z^{\prime}\right)\right) \\
\leq & V\left(t-1, y_{i}-1, y_{j}+1, y_{-\{i, j\}} \mid z\right)
\end{aligned}
$$

The last inequality follows from the properties of the optimal dynamic program - note that the optimal policy starting with inventory levels $\left(y_{i}-1, y_{j}+1, y_{-\{i, j\}}\right)$ may find a set more profitable than $S$ to offer to the customer. Finally, we point out if $i \in S$ and $k=i$ is chosen by the customer, then $V\left(t-1, y_{k}-1, y_{i}-\right.$ $\left.1, y_{j}+1, y_{-\{i, j, k\}} \mid z^{\prime}\right)$ would be defined equivalent to $V\left(t-1, y_{i}-2, y_{j}+1, y_{-\{i, j\}} \mid z^{\prime}\right)$; note that in the induction step we assume $y_{i} \geq y_{j}+2$.

## A.4.2 Proof of Proposition A.1.2

We need to show that any solution of the Inventory-Balancing algorithm corresponds to a feasible solution of linear program FRLP, and when the value of Primal-S is normalized to one, the objective function of linear program FRLP is a lower bound for the expected revenue of the Inventory-Balancing algorithms. In the following, we only show the result for the EIB algorithm; a similar argument can be applied to the LIB algorithm.

Note that we divide the horizon $T$ into $\frac{1}{\epsilon}$ time slots such that $T \epsilon$ is an integer. We only observe the remaining inventory of each product $i$ in periods $k T \epsilon, 0<k \leq \frac{1}{\epsilon}$. Consider the solution (allocation) of the EIB algorithm for a sequence of customers $\left\{z_{t}\right\}_{t=1}^{T}$. For this solution, let $\rho^{j, k}$ be the sum of revenue times capacity of any product $i$ whose fraction of the remaining inventory $\frac{I_{i}^{k T \epsilon}}{c_{i}}$ in the $k^{t h}$ time slot (period $\left.k T \epsilon\right)$ is between $1-(j+1) \epsilon$ and $1-j \epsilon$ (inclusive); that is, $\frac{I_{i}^{k T \epsilon}}{c_{i}} \in(1-(j+1) \epsilon, 1-j \epsilon]$ where $0 \leq j \leq \frac{1}{\epsilon}-1$ and $0<k \leq \frac{1}{\epsilon}$. Similarly, let $\gamma^{j, k}$ be the total revenue obtained from product $i$ in the optimal solution (of Primal-S) where $\frac{I_{i}^{k T \epsilon}}{c_{i}} \in(1-(j+1) \epsilon, 1-j \epsilon]$, i.e.,

$$
\begin{aligned}
\rho^{j, k} & =\sum_{i: \frac{{ }_{i}^{k T \epsilon}}{c_{i}} \in(1-(j+1) \epsilon, 1-j \epsilon]} r_{i} c_{i}, \\
\gamma^{j, k} & =\sum_{\substack{i: \frac{I_{i}^{k T \epsilon}}{c_{i}} \in(1-(j+1) \epsilon, 1-j \epsilon]}} o_{i},
\end{aligned}
$$

where $o_{i}$ is the revenue obtained from selling product $i$ in the optimal solution. Note that $o_{i} \leq r_{i} c_{i}$. For any time slot $k$, we define

$$
\chi(k)=\sum_{i=1}^{n} c_{i} r_{i} \int_{y=\frac{I_{i}^{k T \epsilon}}{c_{i}}}^{1} \Psi(y) d y=\sum_{i=1}^{n} r_{i} \int_{y=I_{i}^{k T \epsilon}}^{c_{i}} \Psi\left(y / c_{i}\right) d y .
$$

Notice that $\chi(k)$ is an increasing function of $k$. Using the fact that in any period $t$ the IB algorithm chooses an assortment $S^{t}$ that maximizes $\sum_{i \in S} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) \phi_{i}^{z_{t}}(S)$, we will bound the change in $\chi$ function at two consecutive time slots, see the fifth sets of constraints in linear program FRLP.

Next, we will show that the objective function of FRLP is less than the total revenue of the EIB algorithm. Since the penalty function is exponential, $\Psi(x)=\frac{e}{e-1}\left(1-e^{-x}\right), x \in[0,1]$, the first term in the objective function is given by

$$
\chi\left(\frac{1}{\epsilon}\right)=\sum_{i=1}^{n} c_{i} r_{i} \int_{y=\frac{I_{i}^{T}}{c_{i}}}^{1} \Psi(y) d y=\frac{e}{e-1} \sum_{i=1}^{n} c_{i} r_{i}\left(1-\frac{I_{i}^{T}}{c_{i}}+e^{-1}-e^{-\frac{I_{i}^{T}}{c_{i}}}\right) .
$$

By definition of $\rho^{j, \frac{1}{\epsilon}}$, the second term of the objective function is lower bounded as follows

$$
\frac{e}{e-1} \sum_{j=0}^{1-\frac{1}{\epsilon}} \rho^{j, \frac{1}{\epsilon}}\left(\frac{j \epsilon}{e}-e^{j \epsilon-1}+e^{-1}\right) \geq \frac{e}{e-1} \sum_{i=1}^{n} c_{i} r_{i}\left(\frac{1-\frac{I_{i}^{T}}{c_{i}}}{e}-e^{-\frac{I_{i}^{T}}{c_{i}}}+e^{-1}\right) .
$$

The inequality holds because $\frac{x}{e}-e^{x-1}$ is a decreasing function of $x$. Therefore,

$$
\chi\left(\frac{1}{\epsilon}\right)-\frac{e}{e-1} \sum_{j=0}^{\frac{1}{\epsilon}-1} \rho^{j, \frac{1}{\epsilon}}\left(\frac{j \epsilon}{e}-e^{j \epsilon-1}+e^{-1}\right) \leq \sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right),
$$

where the left hand side is the objective function of linear program FRLP and the right hand side is the revenue of the EIB algorithm.

The next step is to show that any solution of the EIB algorithm corresponds to a feasible solution of linear program FRLP. Without loss of generality, we can normalize the revenue of the optimal solution, Primal-S, to 1 which implies the first set of constraints. The second set of constraints holds because $o_{i} \leq r_{i} c_{i}$. The third set of constraints follows from the definition of $\rho^{j, k}$ and the fact that $I_{i}^{k T \epsilon}$ is a decreasing function of $k$. The forth set of constraints holds because of the definition of $\chi$ and the fact that $\int_{y=x}^{1} \Psi(y) d y$ is a decreasing function of $x$. In the following, we prove Lemma A.4.2 stated below, which leads us to the last set of constraints. This set of constraints gives a lower bound for difference of $\chi(k+1)$ and $\chi(k)$ as a function of the optimal solution $\gamma^{j, k}$. To prove the lower bound, we show that under uniform permutation, the revenue obtained from product $i$ in the optimal solution during periods in $[k T \epsilon,(k+1) T \epsilon)$ denoted by $o_{i, k}$ is concentrated around its average $\epsilon O_{i}$. Using this concentration and the fact that the IB algorithm chooses a set that maximizes discounted revenue, we will get the desired bound for the change in function $\chi$. Hence,
any solution of the EIB algorithm, corresponds to a feasible solution of the linear program FRLP where the objective is less than the revenue of the algorithm. Considering the fact that the optimal solution, Primal-S, is normalized to one, the solution of the linear program FRLP is a number in $[0,1]$ and by minimizing the objective function, we obtain a lower bound for the performance of the EIB algorithm, namely the ratio of $\mathrm{E}_{\left\{z_{t}\right\}_{t=1}^{T}}\left[\operatorname{Rev}_{E I B}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)\right]$ to Primal-S.

Lemma A.4.2. Suppose $\Psi$ is increasing, concave, and twice differentiable with a bounded derivative, in $[0,1]$ and $\frac{1}{\epsilon \times c_{\mathrm{MIN}}}=O(1)$, then with high probability, for any $0<k<\frac{1}{\epsilon}$,

$$
\epsilon \sum_{j=0}^{\frac{1}{\epsilon}-1} \gamma^{j, k+1} \Psi(1-(j+1) \epsilon) \leq \chi(k+1)-\chi(k)+O\left(\frac{1}{c_{\text {МІІ }}}\right) .
$$

## Proof. Proof of Lemma A.4.2:

Note that the contribution of product $i$ to $\chi(k+1)-\chi(k)$ is equal to $r_{i} c_{i} \int_{I_{i}^{t_{1}} / c_{i}}^{\int_{i}^{t_{0}} /} \Psi(y) d y$ where $t_{0}=k T \epsilon$ and $t_{1}=(k+1) T \epsilon$. By the assumption that the derivative of $\Psi$ is bounded, we can substitute the integral with the sum and get

$$
\begin{equation*}
r_{i} c_{i} \int_{I_{i}^{t_{1}} / c_{i}}^{I_{i}^{t_{0}} / c_{i}} \Psi(y) d y=r_{i} \int_{I_{i}^{t_{1}}}^{I_{i}^{t_{0}}} \Psi\left(y / c_{i}\right) d y \geq r_{i} \sum_{z=I_{i}^{t_{1}}}^{I_{i}^{t_{0}}} \Psi\left(z / c_{i}\right)-O\left(\frac{1}{c_{\text {MIN }}}\right) \tag{A.2}
\end{equation*}
$$

Now let us consider the optimal solution. Let $S_{o p t}^{t}$ denote the set that is shown to the customer in period $t$ by the optimal solution. Since the Inventory-Balancing algorithm shows to each customer an assortment that maximizes $\sum_{i \in S} r_{i} \Psi\left(\frac{I_{i}^{t}}{c_{i}}\right) \phi_{i}^{t}(S)$, we have

$$
\begin{equation*}
\sum_{i=1}^{n} r_{i} \sum_{z=I_{i}^{t_{1}}}^{I_{i}^{t_{0}}} \Psi\left(z / c_{i}\right) \geq \sum_{t=t_{0}}^{t_{1}-1} \sum_{i \in S_{\text {opt }}^{t}} r_{i} \Psi\left(I_{i}^{t} / c_{i}\right) \phi_{i}^{t}\left(S_{o p t}^{t}\right) \geq \sum_{t=t_{0}}^{t_{1}-1} \sum_{i \in S_{\text {opt }}^{t}} r_{i} \Psi\left(I_{i}^{t_{1}} / c_{i}\right) \phi_{i}^{t}\left(S_{o p t}^{t}\right) \tag{A.3}
\end{equation*}
$$

The last inequality holds because $\Psi$ is an increasing function. Recall that $o_{i, k}$ is the revenue obtained by the optimal solution from product $i$ during periods in $[k T \epsilon,(k+1) T \epsilon)$. By this definition and the inequality above, we have

$$
\begin{equation*}
\sum_{i=1}^{n} r_{i} \sum_{z=I_{i}^{t_{1}}}^{I_{i}^{t_{0}}} \Psi\left(z / c_{i}\right) \geq \sum_{i=1}^{n} \Psi\left(I_{i}^{t_{1}} / c_{i}\right) o_{i, k} \tag{A.4}
\end{equation*}
$$

In Lemma A.4.3 which is borrowed from Mirrokni et al. (2012) (with some modifications), we show that under uniform permutation, with high probability, $o_{i, k}$ is concentrated around its average $\epsilon o_{i}$ where $\frac{1}{\epsilon}$ is the number of time slots. Note that this concentration holds under the uniform permutation and as we will discuss below it is useful when $\frac{1}{\epsilon \times c_{\text {MIN }}}=O(1)$. By this lemma and the above equation,

$$
\begin{equation*}
\sum_{i=1}^{n} r_{i} \sum_{z=I_{i}^{t_{1}}}^{I_{i}^{t_{0}}} \Psi\left(z / c_{i}\right) \geq \sum_{i=1}^{n} \Psi\left(I_{i}^{t_{1}} / c_{i}\right) \epsilon O_{i} \tag{A.5}
\end{equation*}
$$

with high probability. Therefore, by Equation (A.5) and the definition of $\gamma$, we have

$$
\begin{equation*}
\sum_{i=1}^{n} r_{i} \sum_{z=I_{i}^{t_{1}}}^{I_{i}^{t_{0}}} \Psi\left(z / c_{i}\right)-O\left(\frac{1}{c_{\text {MIN }}}\right) \geq \epsilon \sum_{j=0}^{\frac{1}{\epsilon}-2} \Psi(1-(j+1) \epsilon) \gamma^{j, k+1}-O\left(\frac{1}{c_{\text {MIN }}}\right) \tag{A.6}
\end{equation*}
$$

Since the left hand side is less that $\chi(k+1)-\chi(k)$, the proof is completed.
Lemma A.4.3. If the customers arrive according to a random order (i.e., a permutation chosen uniformly at random), $\sum_{i=1}^{n} o_{i}=1$, then for any $\delta>0$ and $\frac{1}{T \epsilon} \leq k \leq \frac{1}{\epsilon}$,

$$
\operatorname{Pr}\left[\sum_{i=1}^{n}\left|o_{i, k}-\epsilon o_{i}\right|>\frac{5}{\epsilon \times c_{\mathrm{MIN}} \times \delta}\right]<1-\delta .
$$

The assumption that $\sum_{i=1}^{n} o_{i}=1$ implies that we have normalized Primal-S to 1 . In this lemma, we need $\frac{5}{\epsilon \times c_{\text {MIN }}}$ to be either constant or go to 0 which justifies the assumption $\frac{1}{\epsilon \times c_{\text {MIN }}}=O(1)$ in Lemma A.4.2.

## A.4.3 Appendix to Section 1.5.2

Proof. Proof of Proposition 1.5.1: We prove the claim by revisiting the steps of the proof of Theorem 1.3.3. Let $\left\{z_{t}\right\}_{t=1}^{T}$ be the sequence of the customers. By Lemma 1.3.2, we never offer any product that has no inventory.

We now construct a solution for Dual $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$, with the true selection probabilities, as follows:

$$
\begin{aligned}
\theta_{i} & =r_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right) \\
\lambda^{t} & =\sum_{i=1}^{n} r_{i}\left[\Psi\left(I_{i}^{t-1} / c_{i}\right) \bar{\phi}_{i}^{z_{t}}\left(S^{t}\right)+2 \epsilon^{t}\right]
\end{aligned}
$$

where $S^{t}=\arg \max _{S \in \mathcal{S}} \sum_{i \in S} r_{i} \Psi\left(I_{i}^{t-1} / c_{i}\right) \phi_{i}^{z_{t}}(S)$ is the assortment offered by the IB algorithm. Note that we add the error term $\epsilon^{t}$ to the value of $\lambda^{t}$ because the assortment $S^{t}$ is computed using the estimated selection probability $\phi^{z_{t}}$. This construction gives us a feasible dual solution because

$$
\lambda^{t} \geq \sum_{i=1}^{n} r_{i} \Psi\left(I_{i}^{T} / c_{i}\right)\left(\phi_{i}^{z_{t}}\left(S^{t}\right)-\epsilon^{t}\right)+2 \sum_{i=1}^{n} r_{i} \epsilon^{t} \geq \sum_{i=1}^{n} r_{i} \Psi\left(I_{i}^{T} / c_{i}\right) \phi_{i}^{z_{t}}\left(S^{t}\right)+\sum_{i=1}^{n} r_{i} \epsilon^{t}
$$

where the first inequality follows from the fact that for all $i=1,2, \ldots, n$ and $S \in \mathcal{S}$, $\left|\bar{\phi}_{i}^{z_{t}}(S)-\phi_{i}^{z_{t}}(S)\right| \leq \epsilon^{t}$. The second inequality holds because $\Psi$ is increasing and $I_{i}^{t-1} \geq I_{i}^{T}$. By definition of $S^{t}$,

$$
\begin{aligned}
\lambda^{t} & \geq \max _{S}\left\{\sum_{i \in S} r_{i} \Psi\left(I_{i}^{T} / c_{i}\right) \phi_{i}^{z_{t}}(S)\right\}+\sum_{i=1}^{n} r_{i} \epsilon^{t} \\
& \geq \max _{S}\left\{\sum_{i \in S} r_{i} \Psi\left(I_{i}^{T} / c_{i}\right)\left(\bar{\phi}_{i}^{z t}(S)-\epsilon^{t}\right)\right\}+\sum_{i=1}^{n} r_{i} \epsilon^{t} \\
& \geq \max _{S} \sum_{i \in S}\left(r_{i}-\theta_{i}\right) \bar{\phi}_{i}^{z_{t}}(S),
\end{aligned}
$$

where the second inequality follows from the fact that for all $i$ and $S \in \mathcal{S},\left|\bar{\phi}_{i}^{z_{t}}(S)-\phi_{i}^{z_{t}}(S)\right| \leq \epsilon^{t}$. The third inequality follows from the definition of $\theta_{i}$ and the fact that $\Psi\left(I_{i}^{T} / c_{i}\right)$ is less than or equal to 1 .

It follows from the Weak Duality Theorem that
$\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right) \leq \mathrm{E}\left[\sum_{t=1}^{T} \lambda^{t}+\sum_{i=1}^{n} c_{i} \theta_{i}\right] \leq \mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right)+\sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+2 \sum_{t=1}^{T} \epsilon^{t}\right)\right]$
Hence,

$$
\frac{\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right)\right]}{\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)} \geq \frac{\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right)\right]}{\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}\left(1-\Psi\left(I_{i}^{T} / c_{i}\right)\right)+\sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+2 \sum_{t=1}^{T} \epsilon^{t}\right)\right]} .
$$

where $\mathrm{E}\left[\sum_{i=1}^{n} r_{i}\left(c_{i}-I_{i}^{T}\right)\right]$ is the revenue of the Inventory-Balancing algorithm. Note that the contribution of any product $i$ that is not sold by the optimal solution to Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ is zero. Thus, to find the competitive ratio, we only consider products that are sold by the optimal solution. Since it is assumed that
the IB algorithm sells at least one unit of any product that is sold by the optimal solution, the competitive ratio of the algorithm is at least

$$
\min _{\left(c_{i}, x\right): x \leq 1-\frac{1}{c_{i}}} \frac{1-x}{\frac{1}{c_{i}} \sum_{t=I_{i}^{T}+1}^{c_{i}} \Psi\left(t / c_{i}\right)+(1-\Psi(x))+\frac{2}{c_{i}} \mathrm{E}\left[\sum_{t=1}^{T} \epsilon^{t}\right]}
$$

where $i$ is a product sold by the optimal solution. Therefore, by the same argument as in the proof of Theorem 1.3.3, the competitive ratio is at least

$$
\min _{x \in\left[0,1-\frac{1}{c_{\mathrm{MIN}}}\right.}\left\{\frac{1-x}{\frac{1}{c_{\mathrm{MIN}}}+1-\Psi(x)+\int_{x+\frac{1}{c_{\mathrm{MIN}}}}^{1} \Psi(y) d y+\frac{2}{c_{\mathrm{MIN}}} \mathrm{E}\left[\sum_{t=1}^{T} \epsilon^{t}\right]}\right\}
$$

## A. 5 Asymptotic Optimality of the Dynamic Programming Policy

In this section, we show asymptotic optimality of the dynamic programming ( DP ) policy when the type of customers is drawn independently from a known distribution. Namely, we show that the value obtained by the DP policy approaches Primal-S asymptotically when both the capacities and the horizon scale proportionally.

Let $\eta^{z}>0, z \in \mathcal{Z}$, be the probability that in each period a customer of type $z$ arrives. ${ }^{8}$ Let $V\left(t, x_{1}, \ldots, x_{n} \mid z\right)$ denote the maximum expected revenue with $t$ periods remaining, given that a customer of type $z \in \mathcal{Z}$ arrives, and the remaining inventories are $\left(x_{1}, \ldots, x_{n}\right)$. Then, the dynamic programming formulation of this problem is given by

$$
\begin{aligned}
& V\left(t, x_{1}, \ldots, x_{n} \mid z\right) \\
= & \max _{S \in \mathcal{S}: x_{i} \geq 1 \forall i \in S}\left\{\sum_{i \in S} \phi_{i}^{z}(S)\left[r_{i}+V\left(t-1, x_{1}, \ldots, x_{i}-1, \ldots, x_{n}\right)\right]+\phi_{0}^{z}(S) V\left(t-1, x_{1}, \ldots, x_{n}\right)\right\}
\end{aligned}
$$

where $V\left(t, x_{1}, \ldots, x_{n}\right)=\sum_{z \in \mathcal{Z}} \eta^{z} V\left(t, x_{1}, \ldots, x_{n} \mid z\right)$. Also, the terminal condition is given by $V(0, \cdot)=$ 0 . We denote the optimal revenue under the dynamic programming formulation by $V(T, \boldsymbol{c})$ where $\boldsymbol{c}$ is

[^31]the vector of initial inventories. We note that in computing $V(T, \boldsymbol{c})$, we take expectation with respect to sequence of customers and the customers' choices. For simplicity, we assume that the policy can always offer an "empty" assortment with $S=\varnothing$. Thus, the maximum in the dynamic programming equation is always well-defined.

The asymptotic optimality result is stated in the following Proposition.
Proposition A.5.1 (Asymptotic Optimality of DP). Given that the type of customers is drawn independently from a known distribution such that the probability of arriving a customer of type $z \in \mathcal{Z}$ in any period $t$ is $\eta^{z}$, then

$$
\lim _{\beta \rightarrow \infty} \frac{V(\beta T, \beta \boldsymbol{c})}{\operatorname{Primal}-\mathrm{S}(\beta T, \beta \boldsymbol{c})}=1
$$

where Primal- $\mathbf{S}(\beta T, \beta \boldsymbol{c})$ is the linear programming Primal-S with initial inventories $\beta \boldsymbol{c}$ and the length of the horizon $\beta T$.

In the above proposition, we scale both the horizon and initial inventory with a scalar $\beta$. The corresponding problem is called $\beta$-scaled stochastic problem. Then, to see the asymptotic behavior of dynamic programming, we let $\beta$ go to infinity. We note that Proposition A.5.1 does not imply that dynamic programming policy is asymptotically optimal for every sequence of customer types. Instead it shows that it is asymptotically optimal only when take average over all sequences.

Proof. Proof of Proposition A.5.1:
By Lemma 1.4.1, $V(\beta T, \beta \boldsymbol{c}) \leq \operatorname{Primal-S}(\beta T, \beta \boldsymbol{c})$ for all $T, \boldsymbol{c}$ and $\beta>0$. Now, let $\left\{\bar{y}^{z}(S): S \in \mathcal{S}, z \in \mathcal{Z}\right\}$ denote an optimal solution for the (unscaled) Primal-S $(T, \boldsymbol{c})$ for all $T, \boldsymbol{c}$ and $\beta>0$. Then, it is easy to verify that $\left\{\beta \bar{y}^{z}(S): S \in \mathcal{S}, z \in \mathcal{Z}\right\}$ is an optimal solution to Primal-S $(\beta T, \beta \boldsymbol{c})$.

To show that $\lim _{\beta \rightarrow \infty} \frac{V(\beta T, \beta c)}{\operatorname{Primal}-\mathrm{S}(\beta T, \beta c)}=1$, we construct a deterministic policy $\mu$ for the $\beta$-scaled stochastic problem whose expected revenue approaches Primal-S $(\beta T, \beta \boldsymbol{c})$ as $\beta$ increases toward infinity. We show that this policy is admissible, that is the total sales of product $i$ is less than its initial inventory. Therefore, $V(\beta T, \beta \boldsymbol{c})$ also approaches Primal-S $(\beta T, \beta \boldsymbol{c})$ as $\beta \rightarrow \infty$.

The policy $\mu$ operates as follows: Offer a set $S \in \mathcal{S}$ to customers of type $z$ for up to $\beta \eta^{z} \bar{y}^{z}(S)$ times. The order in which the sets are offered is arbitrary. Under this policy, we will NOT accept all of the demands generated by offering $S$. Rather, we will limit the sales of product ifrom offering $S$ to customers of type $z$ to at most $\beta \eta^{z} \phi_{i}^{z}(S) \bar{y}^{z}(S)$.

Let $N(\beta T)=\left(N^{z}(\beta T): z \in \mathcal{Z}\right)$ be a multinomial random vector, where $N^{z}(\beta T)$ denotes the total number of customers of type $z$ over $\beta T$ periods. Note that $N^{z}(\beta T)$ has a binomial distribution with parameter $\beta T$ and $\eta^{z}$. We define the random variable $D_{i}^{z}(S, q)$ as the total number of customers of type $z$ who select product $i$ when $S$ is offered under the policy $\mu$, given that there are $q$ customers of type $z$. Since under the policy $\mu$ we do not accept all the demands, the total sales of product $i$ from customers of type $z$ generated from offering $S$ under the policy $\mu$ is given by

$$
\text { Sale }_{i}^{\mu, z}(S)=\min \left\{D_{i}^{z}\left(S, N^{z}(\beta T)\right), \beta \eta^{z} \phi_{i}^{z}(S) \bar{y}^{z}(S)\right\}
$$

We point out that Sale $_{i}^{\mu, z}(S)$ is a random variable because $D_{i}^{z}\left(S, N^{z}(\beta T)\right)$ is a random variable. Since $\beta \bar{y}^{z}(S)$ is a feasible solution of linear program $\operatorname{Primal-S}(\beta T, \beta \boldsymbol{c})$, we have

$$
\beta \sum_{z \in \mathcal{Z}} \sum_{S \in \mathcal{S}} \eta^{z} \phi_{i}^{z}(S) \bar{y}^{z}(S) \leq \beta c_{i}, \quad i=1, \ldots, n
$$

which implies that, with probability one,

$$
\sum_{z \in \mathcal{Z}} \sum_{S \in \mathcal{S}} \operatorname{Sale}_{i}^{\mu, z}(S) \leq \beta \sum_{z \in \mathcal{Z}} \sum_{S \in \mathcal{S}} \eta^{z} \phi_{i}^{z}(S) \bar{y}^{z}(S) \leq \beta c_{i}, \quad i=1, \ldots, n
$$

Therefore, the policy $\mu$ is admissible because the total sales of product $i$ does not exceed its initial inventory. The total revenue over $\beta T$ periods under policy $\mu$ is given by a random variable $\sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{z \in \mathcal{Z}} \operatorname{Sale}_{i}^{\mu, z}(S)$. Then,

$$
\begin{aligned}
\lim _{\beta \rightarrow \infty} \frac{1}{\beta} \sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{z \in \mathcal{Z}} \operatorname{Sale}_{i}^{\mu, z}(S) & =\lim _{\beta \rightarrow \infty} \frac{1}{\beta} \sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{z \in \mathcal{Z}} \min \left\{D_{i}^{z}\left(S, N^{z}(\beta T)\right), \beta \eta^{z} \bar{y}^{z}(S) \phi_{i}^{z}(S)\right\} \\
& =\sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{z \in \mathcal{Z}} \min \left\{\lim _{\beta \rightarrow \infty} \frac{1}{\beta} D_{i}^{z}\left(S, N^{z}(\beta T)\right), \eta^{z} \bar{y}^{z}(S) \phi_{i}^{z}(S)\right\} \\
& =\sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{z \in \mathcal{Z}} \eta^{z} \bar{y}^{z}(S) \phi_{i}^{z}(S)=\text { Primal-S }(T, \boldsymbol{c})
\end{aligned}
$$

To establish the third equality above, note that

$$
\frac{1}{\beta} D_{i}^{z}\left(S, N^{z}(\beta T)\right)=\frac{1}{\beta} \sum_{t=1}^{M^{z}} \mathbb{1}_{\left\{B_{t, z}^{i}(S)=1\right\}}=\frac{M^{z}}{\beta} \times \frac{1}{M^{z}} \sum_{t=1}^{M^{z}} \mathbb{1}_{\left\{B_{t, z}^{i}(S)=1\right\}}
$$

where $M^{z}:=\min \left\{N^{z}(\beta T), \beta \eta^{z} \bar{y}^{z}(S)\right\}$ and $B_{t, z}^{i}(S)=1$ denotes the event that the $t^{\text {th }}$ customer of type $z$ selects product $i$ when $S$ is offered, with $\mathrm{E}\left[\mathbb{1}_{\left\{B_{t, z}^{i}(S)=1\right\}}\right]=\phi_{i}^{z}(S)$. By SLLN, we know that $\lim _{\beta \rightarrow \infty} N^{z}(\beta T) / \beta=\eta^{z} T$ almost surely (a.s.). Since under the policy $\mu$, we only offer $S$ up to $\beta \eta^{z} \bar{y}^{z}(S)$ customers of type $z$,

$$
\lim _{\beta \rightarrow \infty} \frac{M^{z}}{\beta}=\lim _{\beta \rightarrow \infty} \frac{\min \left\{N^{z}(\beta T), \beta \eta^{z} \bar{y}^{z}(S)\right\}}{\beta}=\eta^{z} \bar{y}^{z}(S) \text { a.s. }
$$

By a similar argument, $\frac{1}{M^{z}} \sum_{t=1}^{M^{z}} \mathbb{1}_{\left\{B_{t, z}^{i}(S)=1\right\}}=\phi_{i}^{z}(S)$. Thus, with probability one,

$$
\lim _{\beta \rightarrow \infty} \frac{1}{\beta} D_{i}^{z}\left(S, N^{z}(\beta T)\right)=\eta^{z} \bar{y}^{z}(S) \phi_{i}^{z}(S)
$$

which gives us the desired result. Then, by the Dominated Convergence Theorem, it follows that $\lim _{\beta \rightarrow \infty} \frac{1}{\beta} \mathrm{E}\left[\sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{z \in \mathcal{Z}} \mathrm{Sale}_{i}^{\mu, z}(S)\right]=\operatorname{Primal}-\mathrm{S}(T, \boldsymbol{c})$. Since the policy $\mu$ is admissible,

$$
1 \geq \lim _{\beta \rightarrow \infty} \frac{V(\beta T, \beta \boldsymbol{c})}{\operatorname{Primal-S}(\beta T, \beta \boldsymbol{c})} \geq \lim _{\beta \rightarrow \infty} \frac{\frac{1}{\beta} \mathrm{E}\left[\sum_{i=1}^{n} r_{i} \sum_{S \in \mathcal{S}} \sum_{\boldsymbol{z} \in \mathcal{Z}} \mathrm{Sale}_{i}^{\mu, z}(S)\right]}{\frac{1}{\beta} \operatorname{Primal}-\mathrm{S}(\beta T, \beta \boldsymbol{c})}=1
$$

which completes proof.

## Appendix B

## Technical Appendix to Chapter 2

## B. 1 Sequential Weighted Second-Price Mechanism \& Proofs of Theorems 2.4.1 and 2.5.1

In this section, we present a parameterized class of mechanisms called Sequential Weighted Second-Price (SWSP), which is denoted by $\mathcal{M}(\rho, \beta)$. Weight function $\beta: \mathbb{R} \rightarrow \mathbb{R}$ connects the bids in the first and the second rounds by manipulating the final allocation and payments in the favor of agents with higher initial valuations. Parameter $\rho \in \mathbb{R}$ specifies a lower bound on weighted bids the seller is willing to accept for the item.

As it will be more clear later, this class of mechanism includes the efficient and optimal mechanisms. We will show that the SWSP mechanisms are incentive compatible and individually rational. As a corollary of this result, we will conclude that the efficient and optimal mechanisms are IC and IR; that is, Theorems 2.4.1 and 2.5.1 hold.

We make the following assumption on function $\beta$.
Assumption 4. Weight function $\beta$ is non-decreasing and differentiable with bounded derivatives, that is, $\sup _{z}\left\{\beta^{\prime}(z)\right\}<\infty, z \in[\underline{v}, \bar{v}]$.

Note that as it will be more clear later, the non-decreasing function $\beta$ alters the social welfare and the seller's revenue by distorting the allocation via favoring agents with higher valuations. The weight function $\beta$ and parameter $\rho$ can be used to adjust the social welfare and the revenue of the mechanism. For instance, as we show later, for the efficient mechanism, function $\beta(\cdot)$ and $\rho$ are equal to 0 .

Let us start with the description of the mechanism.

Sequential Weighted Second-Price Mechanism $\mathcal{M}(\rho, \beta)$ : The selection, allocation, and payment rules are defined as follows:

- Selection: Select the following set of agents

$$
\begin{equation*}
\mathcal{S}_{\rho, \beta}\left(b_{0}\right) \in \underset{S \subseteq\{1, \cdots, n\}}{\arg \max }\left\{\Omega_{\rho, \beta}\left(b_{0}, S\right)\right\}, \tag{B.1}
\end{equation*}
$$

where $\Omega_{\rho, \beta}\left(b_{0}, S\right)$, the weighted surplus, is defined as follows

$$
\begin{equation*}
\Omega_{\rho, \beta}\left(b_{0}, S\right)=\mathrm{E}_{S}\left[\max \left\{\max _{i \in S}\left\{b_{i, 0}+\beta\left(b_{i, 0}\right)+\delta_{i}\right\}, \max _{i \notin S}\left\{b_{i, 0}+\beta\left(b_{i, 0}\right)\right\}, \rho\right\}\right]-\sum_{i \in S} c_{i} . \tag{B.2}
\end{equation*}
$$

The expectation is with respect to the second signals of the selected agents; in case of ties, we will choose one of the sets at random. Each selected agent $i$ pays $t_{i}\left(b_{0}\right)$ to the seller; see Eq. (2.2).

- Allocation and Payments: Agents participate in a "weighted second-price" auction with a reserve price $\rho$, where the mechanism allocates the item to the agent with the highest weighted bid as along as it exceeds $\rho$. More precisely, consider an agent $i^{\star} \in \operatorname{argmax}_{i}\left\{b_{i, 1}+\beta\left(b_{i, 0}\right)\right\}$. If $b_{i^{\star}, 1}+\beta\left(b_{i^{\star}, 0}\right) \geq \rho$, then the item is allocated to agent $i^{\star}$. If agent $i^{\star}$ was a selected agent, then he pays $p_{i^{\star}}=\max \left\{\max _{i \neq i^{\star}}\left\{b_{i, 1}+\beta\left(b_{i, 0}\right)\right\}, \rho\right\}-\beta\left(b_{i^{\star}, 0}\right)$. If agent $i^{\star}$ is not a selected agent, then he pays $p_{i^{\star}}=\max \left\{\max _{i \neq i^{\star}}\left\{b_{i, 1}+\beta\left(b_{i, 0}\right)\right\}, r\right\}-\beta\left(b_{i^{\star}, 0}\right)$, where $r: \mathbb{R}^{n} \rightarrow \mathbb{R}$ will be defined below.

Note that by letting $\rho=0$ and $\beta(\cdot)=0$, we can implement mechanism $\mathcal{M}^{\text {EFF. }}$. Furthermore, by setting $\rho=0$ and $\beta(\cdot)=\alpha(\cdot)$, we have mechanism $\mathcal{M}^{\text {opt }}$.

Let $\ell$ be an unselected agent with the highest weighted bid, i.e., $\ell \in \arg \max _{j \notin \mathcal{S}_{\rho, \beta}\left(b_{0}\right)}\left\{b_{j, 0}+\beta\left(b_{j, 0}\right)\right\}$. Then, if $b_{\ell, 0}+\beta\left(b_{\ell, 0}\right)<\rho$ or all agents are selected, $r=\rho$. Otherwise, $r$ solves the following equation

$$
\begin{equation*}
\int_{\max \{r, \rho\}}^{b_{\ell, 0}+\beta\left(b_{\ell, 0}\right)} \operatorname{Pr}\left[z \geq \max _{j \in \mathcal{S}_{\rho, \beta}\left(b_{0}\right)}\left\{b_{j, 0}+\delta_{j}+\beta\left(b_{j, 0}\right)\right\}\right] d z=\int_{\underline{v}}^{b_{\ell, 0}} \mathrm{E}\left[q_{\ell} \mid v_{\ell, 0}=z, v_{-\ell, 0}=b_{-\ell, 0}\right] d z . \tag{B.3}
\end{equation*}
$$

The next lemma shows that there exists an $r \in\left[\rho, b_{\ell, 0}+\beta\left(b_{\ell, 0}\right)\right]$ that solves the above equation.

Lemma B.1.1. Consider agent $\ell \in \arg \max _{j \notin \mathcal{S}_{\rho, \beta}\left(b_{0}\right)}\left\{b_{j, 0}+\beta\left(b_{j, 0}\right)\right\}$ in mechanism $\mathcal{M}(\rho, \beta)$. If $b_{\ell, 0}+$ $\beta\left(b_{\ell, 0}\right) \geq \rho$, then there exists $r \in\left[\rho, b_{\ell, 0}+\beta\left(b_{\ell, 0}\right)\right]$ that satisfies Eq. (B.3).

The following results are immediate corollaries of Lemma B.1.1.

Corollary B.1.2. Let $\ell \in \arg \max _{j \notin \mathcal{S}_{\mathrm{EFF}}\left(b_{0}\right)}\left\{b_{j, 0}\right\}$ be an unselected agent $\ell$ with the highest initial bid in mechanism $\mathcal{M}^{\text {EFF. }}$. Then, if $b_{\ell, 0} \geq 0$, then there exists $r \in\left[0, b_{\ell, 0}\right]$ that satisfies $E q$. (2.3)

Corollary B.1.3. Let $\ell \in \arg \max _{j \notin \mathcal{S}_{\text {Orr }}\left(b_{0}\right)}\left\{b_{j, 0}+\alpha\left(b_{j, 0}\right)\right\}$ be an unselected agent $\ell$ with the highest weighted bid in mechanism $\mathcal{M}^{\text {OPr } . ~ T h e n, ~ i f ~} b_{\ell, 0}+\alpha\left(b_{\ell, 0}\right) \geq 0$, then there exists $r \in\left[0, b_{\ell, 0}+\alpha\left(b_{\ell, 0}\right)\right]$ that satisfies Eq. (2.5).

We now present the main result of this section, which shows that the proposed mechanism is IC and IR.
Theorem B.1.4 (Incentive Compatibility). Suppose $\rho \geq 0$ and function $\beta$ satisfies Assumption 4. Then, the Sequential Weighted Second-Price mechanism $\mathcal{M}(\rho, \beta)$ is incentive compatible and individually rational.

The proof of the theorem is given in Appendix B.1.1.
If mechanism $\mathcal{M}(\rho, \beta)$ is incentive compatible, it maximizes the weighted surplus defined below as

$$
\Omega_{\rho, \beta}\left(v_{0}\right)=\underset{S \subseteq\{1, \ldots, n\}}{\arg \max }\left\{\Omega_{\rho, \beta}\left(v_{0}, S\right)\right\}
$$

The assumption that the derivatives of the function $\beta$ are bounded will ensure us that the weighted surplus, $\Omega_{\rho, \beta}\left(v_{0}\right)$, is absolutely continuous in the initial valuations of the agents. Note that $\Omega_{\rho=0, \beta=0}\left(v_{0}\right)$ and $\Omega_{\rho=0, \beta=\alpha}\left(v_{0}\right)$ are, respectively, the maximum social welfare and virtual revenue. Therefore, Theorem B.1.4 implies mechanisms $\mathcal{M}^{\mathrm{EFF}}$ and $\mathcal{M}^{\text {OPT }}$ are efficient and optimal, respectively.

## B.1.1 Proof of Theorem B.1.4

In this section, we prove Theorem B.1.4. We start with incentive compatibility and show that no agents would prefer to deviate from the truthful strategy, as long as all other agents are truthful. We prove this by going over the strategy of an agent in a backward manner. First, using Lemma B.1.5, we show that agents bid truthfully in the second round. Then, we prove that a selected agent obtains the additional information (Lemma B.1.6). Finally, in Lemma B.1.9 we show that agents will be better off by being truthful in the first round. We present the proof of Lemma B.1.9 in Section B.1.2 since it our key technical lemma. The proofs of other lemmas are relegated to Section B.4.

The key challenging part is to show that agents bid truthfully in the first round. The reason is that the effects of initial bids are twofold. First, they determine the set of selected agents. Second, they influence the final allocation of the item.

The following lemma shows that agents who can bid in the second round will be truthful even if they were untruthful in the first round. Precisely, we will show that

$$
\begin{equation*}
v_{i, 1}=\underset{b_{i, 1}}{\arg \max }\left\{q_{i} v_{i, 1}-p_{i}-t_{i}-c_{i} e_{i}\right\} . \tag{B.4}
\end{equation*}
$$

Note that unselected agents do not bid in the second round; that is, their initial bids are considered as their final bids.

Lemma B.1.5 (Truthfulness in the Second Round). Under mechanism $\mathcal{M}(\rho, \beta)$, for any agent that is allowed to update his bid in the second round of bidding, truthfulness is a weakly dominant strategy, even if the agent has not been truthful in the first round.

From a technical perspective, one of the aspects that differentiates our work from the previous work on dynamic mechanism design, in particular Ëso \& Szentes (2007), is that the deviation strategies of the agents, in addition to misreporting his valuations, include the decision on obtaining information. In the following lemma, we show that a selected agent $i$ will acquire the additional information when he bids truthfully in the first round, and all other agents follow the truthful strategy, i.e.,

$$
\begin{equation*}
1=\underset{e_{i} \in\{0,1\}}{\arg } \max \left\{q_{i} v_{i, 1}-p_{i}-t_{i}-e_{i} c_{i} \mid b_{0}=v_{0}, b_{-i, 1}=v_{-i, 1}, e_{j}=1 \text { if } s_{j}=1, j \neq i\right\} \tag{B.5}
\end{equation*}
$$

The conditions in the above equation imply that agent $i$ bids truthfully in the first round, and all other agents follow the truthful strategy; that is, they bid truthfully in both rounds and if they get selected, they obtain information.

Lemma B.1.6 (Obtaining Additional Information). Consider a selected agent $i$ who bids truthfully in the first round. Assuming all other agents are truthful, agent $i$ would incur cost $c_{i}$ to obtain signal $\delta_{i}$.

Lemma B.1.6 implies that a selected agent $i$ "who bids truthfully in the first round" will obtain information. However, if agent $i$ bids untruthfully in the first round is selected, he will not necessarily obtain information. We will show in Lemma B.1.9 that agent $i$ will not gain from bidding untruthfully in the first round regardless of his decision to obtain information.

The proof of Lemma B.1.6is provided in Section B.4.3. To obtain the result, we show that the incentive of the selected agent $i$ gets aligned with the selection rule. Thus, the selected agent prefers to incur $\operatorname{cost} c_{i}$ and acquire information; that is, $e_{i}=1$.

The final step is to show that an agent $i$ will bid truthfully in the first round. Let $U_{i}\left(x_{i}, \hat{x}_{i}\right)$ be the utility of agent $i$ with initial valuation $x_{i}$ when he bids $\hat{x}_{i}$ in the first round and follows the "optimal strategy" afterwards (assuming other agents are truthful). More precisely,
$U_{i}\left(x_{i}, \hat{x}_{i}\right)=\max _{b_{i, 1}, e_{i}}\left\{\mathrm{E}\left[q_{i}\left(\left(\hat{x}_{i}, v_{-i, 0}\right),\left(b_{i, 1}, v_{-i, 1}\right)\right) v_{i, 1}-e_{i} c_{i}-s_{i} t_{i}\left(\left(\hat{x}_{i}, v_{-i, 0}\right)\right)-p_{i}\left(\left(\hat{x}_{i}, v_{-i, 0}\right),\left(b_{i, 1}, v_{-i, 0}\right)\right)\right]\right\}$,
where the expectation is taken assuming that all agents except for agent $i$ are truthful. Then, considering the fact that initial valuation $v_{0}=x$, for any $j \neq i$, we have $v_{j, 1}=x_{j}+\delta_{j}$ if agent $j$ is selected and $x_{j}$ otherwise. Note that after initial bidding, agent $i$ optimizes over $\left(e_{i}, b_{i, 1}\right)$ to obtain his best (utility-maximizing) strategy.

We start with characterizing $U_{i}\left(x_{i}, x_{i}\right)$.

Lemma B.1.7. If the vector of initial valuations is given by $x$, and all agents except for agent $i$ are truthful, then the expected utility of agent $i$ who bid truthfully in the first round, denoted by $U_{i}\left(x_{i}, x_{i}\right)$, is equal to

$$
\begin{equation*}
U_{i}\left(x_{i}, x_{i}\right)=\int_{\underline{v}}^{x_{i}} \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z, \tag{B.7}
\end{equation*}
$$

In addition, $U_{i}\left(x_{i}, x_{i}\right)$ is non-decreasing in $x_{i}$.

In the proof of Lemma B.1.7, we use lemmas B.1.5 and B.1.6 where we show that if agent $i$ bids truthfully in the first round, he will prefer to follow the truthful strategy afterwards.

Lemma B.1.7 implies that the utility of truthful agent $i$ is fully determined by his allocation probability for different initial valuations. Furthermore, the higher his probability of allocation is, the more utility he earns. In fact, Eq. (B.7) is analogous to the utility of an agent in standard static incentive compatible mechanisms (see Myerson (1981)).

We now consider the utility of agent $i$ when he bids untruthfully in the first round and follows his optimal strategy thereafter, i.e., $U_{i}\left(x_{i}, \hat{x}_{i}\right)$ defined in Eq. (B.6). By Lemma B.1.5, if agent $i$ with initial bid $\hat{x}_{i} \neq x_{i}$ gets selected, he bids truthfully in the second round, i.e., $b_{i, 1}=v_{i, 1}$. But, untruthful agent $i$ will not necessarily obtain information if the mechanism selects him. In the following, we denote his best investing strategy by $e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$.

The next lemma establishes an upper bound on $U_{i}\left(x_{i}, \hat{x}_{i}\right)$.

Lemma B.1.8. Suppose the vector of initial valuations is given by $x$, and all agents except agent $i$ are truthful. We have $U_{i}\left(x_{i}, \hat{x}_{i}\right) \leq \max \left\{U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)+\int_{\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right] d z, 0\right\}$, where $\omega_{-i}$ is the maximum weighted bids of all agents but agent $i$ when he misreports $\hat{x}_{i}$ in the first round and other agents are truthful, i.e.,

$$
\begin{equation*}
\omega_{-i}=\max \left\{\max _{j \in \mathcal{S}_{\rho, \beta}\left(\hat{x_{i}}, x_{-i}\right), j \neq i}\left\{x_{j}+\beta\left(x_{j}\right)+\delta_{j}\right\}, \max _{j \notin \mathcal{S}_{\rho, \beta}\left(\hat{x}_{i}, x_{-i}\right), j \neq i}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, \rho\right\} \tag{B.8}
\end{equation*}
$$

and $e_{i}=e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$.

The term inside the integral, i.e., $\operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right]$, is the probability that agent $i$ with final weighted bid $z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right)$ wins the item when agents in set $\mathcal{S}_{\rho, \beta}\left(\hat{x}_{i}, x_{-i}\right) \backslash\{i\}$ obtain information and agent $i$ follows an investing strategy associated with $e_{i}=e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$.

Next, we show that $U_{i}\left(x_{i}, x_{i}\right) \geq U_{i}\left(x_{i}, \hat{x}_{i}\right)$; that is an agent $i$ prefers to bid truthfully in the first round. In Lemma B.1.8, we find an upper bound for $U_{i}\left(x_{i}, \hat{x}_{i}\right)$. Then, when $U_{i}\left(x_{i}, \hat{x}_{i}\right)=0$, immediately we have $U_{i}\left(x_{i}, x_{i}\right) \geq U_{i}\left(x_{i}, \hat{x}_{i}\right)=0$. Now we show that even the upper bound of $U_{i}\left(x_{i}, \hat{x}_{i}\right)$, i.e., $U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)+$ $\int_{\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right] d z$, is smaller than $U_{i}\left(x_{i}, x_{i}\right)$, where $e_{i}=e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$.

We start with defining a suboptimal selection rule $\hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)$ for any nonzero measure interval $\left[y_{1}, y_{2}\right]$ such that $y_{1} \leq y_{2}$ and $\left[y_{1}, y_{2}\right] \subseteq\left[\min \left\{x_{i}, \hat{x}_{i}\right\}, \max \left\{x_{i}, \hat{x}_{i}\right\}\right]$, where

$$
\hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)= \begin{cases}\mathcal{S}_{\rho, \beta}\left(\hat{x}_{i}, x_{-i}\right) \backslash\{i\} & \text { if } e_{i}=0 \text { and } z \in\left[y_{1}, y_{2}\right] ; \\ \mathcal{S}_{\rho, \beta}\left(\hat{x}_{i}, x_{-i}\right) & \text { if } e_{i}=1 \text { and } z \in\left[y_{1}, y_{2}\right] ; \\ \mathcal{S}_{\rho, \beta}\left(z, x_{-i}\right) & \text { otherwise },\end{cases}
$$

where $e_{i}=e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$. Note that the suboptimal selection rule $\hat{\mathcal{S}}_{y_{1}, y_{2}}$ follows the optimal selection rule everywhere except interval $\left[y_{1}, y_{2}\right.$ ]. In the interval $\left[y_{1}, y_{2}\right.$ ], the set of selected agents is the set of agents that update their valuations when all agents except for agent $i$ follow the truthful strategy, and agent $i$ with initial valuation $x_{i}$ misreports $\hat{x}_{i}$ in the first round and follows his best strategy afterward with regard to obtaining the additional information. The suboptimal selection rule captures the untruthful behavior of agent $i$ in the first round while other agents are truthful.

In the next lemma, by characterizing the difference between the weighted surplus under the selection rule of the SWSP mechanism, $\mathcal{S}_{\rho, \beta}(x)$, and suboptimal selection rule $\hat{\mathcal{S}}_{y_{1}, y_{2}}(x)$, we will show that agent $i$ prefers to bid truthfully in the first round; that is, $U_{i}\left(x_{i}, x_{i}\right) \geq U_{i}\left(x_{i}, \hat{x}_{i}\right)$.

Lemma B.1.9. For any interval $\left[y_{1}, y_{2}\right]$, consider the suboptimal selection rule $\hat{\mathcal{S}}_{y_{1}, y_{2}}$ described above. Then, for any $x \in[\underline{v}, \bar{v}]^{n}$, we have $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right)-\Omega_{\rho, \beta}\left(x, \mathcal{S}_{\rho, \beta}(x)\right) \leq 0$, and as a result, agent $i$ prefers to bid truthfully in the first round; that is, $U_{i}\left(x_{i}, x_{i}\right) \geq U_{i}\left(x_{i}, \hat{x}_{i}\right)$.

Note that by Lemma B.1.9, an agent $i$ who misreports his initial valuation will be worse off regardless of his decision to acquire information.

## B.1.2 Proof of Lemma B.1.9

Throughout the proof, to simplify our notations, we denote $\mathcal{S}_{\rho, \beta}$ by $\mathcal{S}$. Furthermore, without loss of generality, we assume that $\hat{x}_{i}<x_{i}$. A similar argument can be applied when $\hat{x}_{i}>x_{i}$.

We first show that the weighted surplus under the selection rule of the SWSP mechanism, $\mathcal{S}(x)$, i.e., $\Omega_{\rho, \beta}(x, \mathcal{S}(x))$, is less than the weighted surplus under selection rule $\hat{\mathcal{S}}_{y_{1}, y_{2}}(x)$, i.e., $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right)$. Then, by characterizing the difference between $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right)$ and $\Omega_{\rho, \beta}(x, \mathcal{S}(x))$ as a function of allocation probabilities, and using the fact that $\beta(\cdot)$ is an increasing function, we show that $U_{i}\left(x_{i}, x_{i}\right) \geq U_{i}\left(x_{i}, \hat{x}_{i}\right)$.

By definition, $\Omega_{\rho, \beta}(x, \mathcal{S}(x))=\Omega_{\rho, \beta}(x)$. Then, since the selection rule $\mathcal{S}(x)$ maximizes the weighted surplus, we have $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right)-\Omega_{\rho, \beta}(x) \leq 0$. Next, in Lemma B.1.10, we characterize $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right)-\Omega_{\rho, \beta}(x)$. The proof which uses the Envelope Theorem (cf. Milgrom \& Segal (2002)) is provided in Section B.4.4.

Lemma B.1.10. For any interval $\left[y_{1}, y_{2}\right]$, consider the suboptimal selection rule $\hat{\mathcal{S}}_{y_{1}, y_{2}}$. Then, we have

$$
\begin{aligned}
& \Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right)-\Omega_{\rho, \beta}(x) \\
& =\int_{y_{1}}^{y_{2}}\left(1+\beta^{\prime}(z)\right)\left(\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)\right]-\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \mathcal{S}\left(z, x_{-i}\right)\right]\right) d z,
\end{aligned}
$$

where $\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, S\right]$ is the probability that truthful agent $i$ with initial valuation $z$ receives the item when other agents bid truthfully and agents in set $S$ update their valuations.

In the following, using Lemma B.1.10 and the fact that $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right) \leq \Omega_{\rho, \beta}(x)$, we will show that agent $i$ prefers to bid truthfully in the first round. Precisely, we will show that

$$
\begin{equation*}
\int_{z=\hat{x}_{i}}^{x_{i}} \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z \geq \int_{z=\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right] d z \tag{B.9}
\end{equation*}
$$

where by Lemma B.1.7, the 1.h.s. is $U_{i}\left(x_{i}, x_{i}\right)-U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)$. In addition, by Lemma B.1.8, the r.h.s. is an upper bound of $U_{i}\left(x_{i}, \hat{x}_{i}\right)-U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)$. Thus the above equation implies that agent $i$ does not have any incentive to bid untruthfully in the first round.

By definition, $\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \mathcal{S}\left(z, x_{-i}\right)\right]=\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right]$. Then, by Lemma B.1.10 and the fact that $\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right) \leq \Omega_{\rho, \beta}(x)$, we have

$$
\begin{equation*}
\int_{y_{1}}^{y_{2}}\left(1+\beta^{\prime}(z)\right)\left(\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right]-\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)\right]\right) d z \geq \tag{B.10}
\end{equation*}
$$

Note that for any $y_{1} \leq z \leq y_{2}$, we have

$$
\begin{align*}
\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)\right] & =\operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta(z) \geq \omega_{-i}\right] \\
& \geq \operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right] \tag{B.11}
\end{align*}
$$

Here, again, $e_{i}=e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$, and the equality follows from the construction of the suboptimal selection rule and the definition of $\omega_{-i}$ in Eq. (B.8). In addition, the inequality holds because the weight function $\beta(\cdot)$ is non-decreasing. Applying Eq. (B.11) in Eq. (B.10), we obtain

$$
\int_{y_{1}}^{y_{2}}\left(1+\beta^{\prime}(z)\right)\left(\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right]-\operatorname{Pr}\left[z+e_{i} \delta_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right]\right) d z \geq 0 .
$$

Then, Eq. (B.9) follows from the fact that the above equation holds for any nonzero measure interval $\left[y_{1}, y_{2}\right] \subseteq\left[\hat{x}_{i}, x_{i}\right]$, and the weight function $\beta(\cdot)$ is non-decreasing. Therefore, agent $i$ prefers to bid truthfully in the first round.

## B. 2 Proof of Theorem 2.6.1

We show this result for any mechanism which selection rule maximizes the weighted surplus $\Omega_{\rho, \beta}\left(b_{0}, S\right)$, where $\Omega_{\rho, \beta}\left(b_{0}, S\right)$ is defined in Eq. (B.4), $\rho \geq 0$, and $\beta: \mathbb{R} \rightarrow \mathbb{R}$ is a non-decreasing function from an initial bid to a weight. Note that the efficient and optimal mechanisms select a set of agents that maximize $\Omega_{\rho=0, \beta=0}\left(b_{0}, S\right)$ and $\Omega_{\rho=0, \beta=\alpha}\left(b_{0}, S\right)$, respectively.

Consider two unselected agents $i, j$ such that $b_{i, 0}>b_{j, 0}$. Assume that agents in set $S$ are already selected. We will show that when the cost of information is the same for all agents, $\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)$ is greater than or equal to $\Omega_{\rho, \beta}\left(b_{0}, S \cup\{j\}\right)$; that is, the seller prefers to add agent $i$ to set $S$ rather than agent $j$.

By definition,

$$
\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)=\mathrm{E}\left[\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, b_{i, 0}+\beta\left(b_{i, 0}\right)+\delta_{i}, b_{j, 0}+\beta\left(b_{j, 0}\right), \rho\right\}\right]-c \times(|S|+1),
$$

where the expectation is with respect to second signals and $b_{k, 1}=b_{k, 0}+\delta_{k}$ if $k \in S$ and is $b_{k, 0}$ otherwise. For any realizations of second signals, let $Y_{-i}=\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, b_{j, 0}+\beta\left(b_{j, 0}\right), \rho\right\}-b_{i, 0}-$ $\beta\left(b_{i, 0}\right)$. Then, $\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)$ is given by

$$
\mathrm{E}\left[\left(b_{i, 0}+\beta\left(b_{i, 0}\right)+\delta_{i}\right) \mathbb{1}\left\{\delta_{i} \geq Y_{-i}\right\}+\left(Y_{-i}+b_{i, 0}+\beta\left(b_{i, 0}\right)\right) \mathbb{1}\left\{\delta_{i}<Y_{-i}\right\}\right]-c \times(|S|+1) .
$$

After some manipulations, it can be rewritten as

$$
b_{i, 0}+\beta\left(b_{i, 0}\right)+Y_{-i} \mathcal{G}\left(Y_{-i}\right)+\int_{Y_{-i}}^{\bar{\delta}} z d \mathcal{G}(z)-c \times(|S|+1),
$$

where $\mathcal{G}$ is the distribution of $\delta_{i}$. Likewise,

$$
\Omega_{\rho, \beta}\left(b_{0}, S \cup\{j\}\right)=b_{j, 0}+\beta\left(b_{j, 0}\right)+Y_{-j} \mathcal{G}\left(Y_{-j}\right)+\int_{Y_{-j}}^{\bar{\delta}} z d \mathcal{G}(z)-c \times(|S|+1),
$$

where $Y_{-j}=\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, b_{i, 0}+\beta\left(b_{i, 0}\right), \rho\right\}-b_{j, 0}-\beta\left(b_{j, 0}\right)$. Using integration by part, one can easily show that $\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)-\Omega_{\rho, \beta}\left(b_{0}, S \cup\{j\}\right)=b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)-\int_{Y_{-i}}^{Y-j} \mathcal{G}(z) d z$. To show the result we need to consider the following two cases.

- $Y_{-i}-Y_{-j}=b_{j, 0}+\beta\left(b_{j, 0}\right)-b_{i, 0}-\beta\left(b_{i, 0}\right)$ : In this case, $\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, \rho\right\}$ is greater than $b_{j, 0}+\beta\left(b_{j, 0}\right)$ and $b_{i, 0}+\beta\left(b_{i, 0}\right)$. That is, $Y_{-i}=\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, \rho\right\}-b_{i, 0}-\beta\left(b_{i, 0}\right)$ and $Y_{-j}=\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, \rho\right\}-b_{j, 0}-\beta\left(b_{j, 0}\right)$. Thus,
$\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)-\Omega_{\rho, \beta}\left(b_{0}, S \cup\{j\}\right) \geq b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)-\left(b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)\right)=0$,
where the inequality follows from the fact that for any $z, \mathcal{G}(z) \leq 1$.
- $Y_{-i}-Y_{-j} \neq b_{j, 0}+\beta\left(b_{j, 0}\right)-b_{i, 0}-\beta\left(b_{i, 0}\right)$ : In this case, $\max \left\{\max _{k \neq i, j}\left\{b_{k, 1}+\beta\left(b_{k, 0}\right)\right\}, \rho\right\}$ is less than $b_{i, 0}+\beta\left(b_{i, 0}\right)$. That is, $Y_{-j}=b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)$ and $Y_{-i} \geq b_{j, 0}+\beta\left(b_{j, 0}\right)-b_{i, 0}-\beta\left(b_{i, 0}\right)$. Then,
$\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)-\Omega_{\rho, \beta}\left(b_{0}, S \cup\{j\}\right) \geq b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)-\int_{b_{j, 0}+\beta\left(b_{j, 0}\right)-b_{i, 0}-\beta\left(b_{i, 0}\right)}^{b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)} \mathcal{G}(z) d z$.
Note that the upper level of the integral equals negative the lower level of the integral. Thus, by the fact the $\mathcal{G}(-z)=1-\mathcal{G}(z)$, we have $\int_{b_{j, 0}+\beta\left(b_{j, 0}\right)-b_{i, 0}-\beta\left(b_{i, 0}\right)}^{b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right.} \mathcal{G}(z) d z=b_{i, 0}+\beta\left(b_{i, 0}\right)-b_{j, 0}-\beta\left(b_{j, 0}\right)$; that is, $\Omega_{\rho, \beta}\left(b_{0}, S \cup\{i\}\right)-\Omega_{\rho, \beta}\left(b_{0}, S \cup\{j\}\right)$ is at least zero, which is the desired result.


## B. 3 Appendix to Sections 2.4.1 and 2.5.3

In this section, we first present the proof of Theorem 2.4.2. Then, we show that the single crossing conditions might not hold in the All-Access mechanism. We then provide the proof of Theorems 2.4.4 and 2.5.5.

## B.3.1 Proof of Theorem 2.4.2

To show the result, we consider the following procedure: At each round, one agent presumes that other agents are playing stationary strategies. That is, he assumes that the maximum bid of other agents, $B=\max \left\{\max _{j \neq i}\left\{v_{j, 0}+\delta_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right\}, 0\right\}$, is drawn from a stationary distribution. Then, he best responds to the strategies of other agents. We show this procedure terminates in an equilibrium. To this aim, we
establish that when an agent updates his strategy in any round, he increases a bounded potential function $\bar{\Omega}:\left\{\left(\tilde{e}_{1}, \tilde{e}_{2}, \ldots, \tilde{e}_{n}\right)\right\} \rightarrow \mathbb{R}$, defined below,

$$
\bar{\Omega}\left(\tilde{e}_{1}, \tilde{e}_{2}, \ldots, \tilde{e}_{n}\right)=\mathrm{E}\left[\max \left\{\max _{j=1,2, \ldots, n}\left\{v_{j, 0}+\delta_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right\}, 0\right\}-\sum_{j=1}^{n} c_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right],
$$

where $\tilde{e}_{j}=\tilde{e}_{j}(\cdot), j=1,2, \ldots, n$, is the investment strategy of agent $j$, and the expectation is with respect to $v_{j, 0}$ and $\delta_{j}$ for $j=1,2, \ldots, n$. Note that the potential function is the average social welfare of the auctioneer and agents when agents follow investment strategies $<\tilde{e}_{1}, \tilde{e}_{2}, \ldots, \tilde{e}_{n}>$. Then, considering the fact that $\bar{\Omega}\left(\tilde{e}_{1}, \tilde{e}_{2}, \ldots, \tilde{e}_{n}\right) \leq \mathrm{E}\left[\max _{S \subseteq\{1,2, \ldots\}}\left\{\Omega\left(v_{0}, S\right)\right\}\right]<\infty$, the potential function is bounded, and as a result, the process of updating strategies will eventually result in an equilibrium.

For any initial valuation $v_{i, 0}$, agent $i$ selects $\tilde{e}_{i}\left(v_{i, 0}\right) \in\{0,1\}$ that maximizes his utility. Let $\tilde{e}_{-i}$ be investment strategies of all agents except of agent $i$. Given $\tilde{e}_{-i}$, the utility of agent $i$ with initial valuation $v_{i, 0}$ when he does not obtain information $\left(\tilde{e}_{i}\left(v_{i, 0}\right)=0\right)$ can be written as

$$
\begin{align*}
u_{i}\left(v_{i, 0}, \tilde{e}_{i}\left(v_{i, 0}\right)=0, \tilde{e}_{-i}\right) & =\mathrm{E}_{v_{-i, 0}, \tilde{e}_{-i}}\left[\max \left\{v_{i, 0}-B, 0\right\}\right] \\
& =\mathrm{E}_{v_{-i, 0}, \tilde{e}_{-i}}\left[\max \left\{v_{i, 0}-\max \left\{\max _{j \neq i}\left\{v_{j, 0}+\delta_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right\}, 0\right\}, 0\right\}\right], \tag{B.12}
\end{align*}
$$

where $\mathrm{E}_{v_{-i, 0}, \tilde{e}_{-i}}$ denotes the expectation with respect to the initial valuations and second signals of all agents except for agent $i$ while taking into account their investment strategies, i.e., $\tilde{e}_{-i}$. Recall that the agents bid truthfully in the second price auction. Similarly, when agent $i$ obtains information, his expected utility is given by
$u_{i}\left(v_{i, 0}, \tilde{e}_{i}\left(v_{i, 0}\right)=1, \tilde{e}_{-i}\right)=\mathrm{E}_{v_{-i, 0}, \tilde{e}_{-i}}\left[\mathrm{E}\left[\max \left\{\left(v_{i, 0}+\delta_{i}\right)-\max \left\{\max _{j \neq i}\left\{v_{j, 0}+\tilde{e}_{j}\left(v_{j, 0}\right) \delta_{j}\right\}, 0\right\}, 0\right\}\right]\right]-c_{i}$,
where the inner expectation is with respect to $\delta_{i}$. In the rest of the proof, we denote all the expectations by E. Then, for any initial valuation $v_{i, 0}$, we have

$$
\begin{aligned}
\tilde{e}_{i}\left(v_{i, 0}\right) & =\arg \max _{e \in\{0,1\}}\left\{u_{i}\left(v_{i, 0}, e, \tilde{e}_{-i}\right)\right\}=\arg \max _{e \in\{0,1\}}\left\{\mathrm{E}\left[\max \left\{v_{i, 0}+\delta_{i} e-B, 0\right\}\right]-c_{i} e\right\} \\
& =\arg \max _{e \in\{0,1\}}\left\{\mathrm{E}\left[\max \left\{\left(v_{i, 0}+\delta_{i} e\right), B\right\}-B-c_{i} e\right]\right\} \\
& =\arg \max _{e \in\{0,1\}}\left\{\mathrm{E}\left[\max \left\{\left(v_{i, 0}+\delta_{i} e\right), B\right\}-c_{i} e-\sum_{j \neq i} c_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right]\right\} \\
& =\arg \max _{e \in\{0,1\}}\left\{\mathrm{E}\left[\max \left\{\left(v_{i, 0}+\delta_{i} e\right), \max \left\{\max _{j \neq i}\left\{v_{j, 0}+\delta_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right\}, 0\right\}\right\}-c_{i} e-\sum_{j \neq i} c_{j} \tilde{e}_{j}\left(v_{j, 0}\right)\right]\right\},
\end{aligned}
$$

where the last equation follows from the definition of $B$. It is easy to observe that at any round, when an agent $i$ updates his strategy, the potential function is increased. Then, by the fact that the potential function is bounded, the process of updating strategies will eventually terminate in an equilibrium point.

## B.3.2 The All-Access Mechanism and Single Crossing Conditions

The following proposition shows that in the setting in Example 2.4.3, the All-Access mechanism does not admit any equilibrium with increasing investment decisions.

Proposition B.3.1. Consider the All-Access mechanism with no reserve price and the setting in Example 2.4.3. Then, there exists no equilibrium such that both agents follow increasing investment decisions. That is, if an agent $i=1,2$ acquires information only when his initial valuation is greater than $\mu_{i} \in(0,1)$, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=1$ for $v_{i, 0} \geq \mu_{i}$, and $\tilde{e}_{i}\left(v_{i, 0}\right)=0$ for $v_{i, 0}<\mu_{i}$, then the other agent $j \neq i$ will not choose the following increasing investment decision: $\tilde{e}_{j}\left(v_{j, 0}\right)=1$ for $v_{j, 0} \geq \mu_{j}$, and $\tilde{e}_{j}\left(v_{j, 0}\right)=0$ for $v_{j, 0}<\mu_{j}$, where $\mu_{j} \in(0,1)$.

## Proof of Proposition B.3.1

To show the result, we will assume that an agent $i=1,2$ follows an increasing investment decision. That is, he only obtains the additional information when his initial valuation is greater than $\mu_{i} \in(0,1)$. Then, we will establish that agent $j \neq i$ will not follow an increasing investment decision.

Throughout the proof, for simplicity, we denote $\mu_{i}$ by $\mu$. We define $W\left(v_{j, 0}\right)$ as the difference between the utility of agent $j$ when he obtains information and the utility of agent $j$ when he does not obtain information given that agent $i \neq j$ only obtains information when his initial valuation is greater than $\mu$, i.e.,

$$
W\left(v_{j, 0}\right)=u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=1, \tilde{e}_{i}\right)-u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=0, \tilde{e}_{i}\right)
$$

where $u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=1, \tilde{e}_{i}\right)$ and $u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=0, \tilde{e}_{i}\right)$ are defined in Equations (B.13) and (B.12). Here, $\tilde{e}_{i}(v)=1$ for any $v \in[\mu, 1]$, and $\tilde{e}_{i}(v)=0$ for $v \in[0, \mu)$. We will show that $W(\cdot)$ is a unimodular function and obtains its unique maximum at $\hat{v} \in(0, \mu)$. Precisely, we will show that $\frac{\partial W\left(v_{j, 0}\right)}{\partial v_{j, 0}} \geq 0$ for $v_{j, 0} \in[0, \hat{v}]$ and $\frac{\partial W\left(v_{j, 0}\right)}{\partial v_{j, 0}} \leq 0$ for $v_{j, 0} \in[\hat{v}, 1]$. This implies that for any values of the cost, agent $j$ will not follow an increasing investment decision. To see why note that agent $j$ with initial valuation $v_{j, 0}$ updates his valuation if $W\left(v_{j, 0}\right) \geq 0$. Then, considering the fact that $W(\cdot)$ is unimodular, $\left\{v_{j, 0}: W\left(v_{j, 0}\right) \geq 0\right\}$ cannot be in the form of $\left[\mu_{j}, 1\right]$ where $\mu_{j} \in(0,1)$.

Let $H(\cdot)$ be the distribution of the maximum bid that agent $j$ competes against, i.e., $B=\max \left\{\left(v_{i, 0}+\right.\right.$ $\left.\left.\delta_{i} \tilde{e}_{i}\left(v_{i, 0}\right)\right), 0\right\}$ where $\tilde{e}_{i}\left(v_{i, 0}\right)=1$ for $v_{i, 0} \in[\mu, 1]$, and is zero otherwise. Then

$$
W\left(v_{j, 0}\right)=\mathrm{E}\left[\int_{x=0}^{\max \left\{v_{j, 0}+\delta_{j}, 0\right\}}\left(v_{j, 0}+\delta_{j}-x\right) d H(x)\right]-\int_{x=0}^{v_{j, 0}}\left(v_{j, 0}-x\right) d H(x)-c
$$

By Leibniz's integral rule, $\frac{\partial W\left(v_{j, 0}\right)}{\partial v_{j, 0}}$ can be written as

$$
\frac{\partial W_{r}\left(v_{j, 0}\right)}{\partial v_{j, 0}}=\mathrm{E}\left[H\left(\max \left\{v_{j, 0}+\delta_{j}, 0\right\}\right)\right]-H\left(v_{j, 0}\right)=\mathrm{E}\left[H\left(v_{j, 0}+\delta_{j}\right)\right]-H\left(v_{j, 0}\right)
$$

The second equality holds because $B \geq 0$ and as a result $H(x)=0$ for any $x<0$. The next Lemma characterizes the distribution $H(\cdot)$.

Lemma B.3.2. Suppose that $v_{i, 0} \sim u(0,1), \delta_{i} \sim u(-1,1), \tilde{e}_{i}\left(v_{i, 0}\right)=1$ for $v_{i, 0} \in[\mu, 1]$, and $\tilde{e}_{i}\left(v_{i, 0}\right)=0$ for $v_{i, 0} \in[0, \mu)$. Then, the distribution of $B=\max \left\{\left(v_{i, 0}+\tilde{e}_{i}\left(v_{i, 0}\right) \delta_{i}\right), 0\right\}$, denoted by $H$, is given by

$$
H(x)= \begin{cases}0 & x<0 . \\ \frac{3-\mu}{2} x+\frac{(\mu-1)^{2}}{4} & 0 \leq x<\mu \\ \frac{1-\mu}{2} x+\mu+\frac{(\mu-1)^{2}}{4} & \mu \leq x<\mu+1 \\ -\frac{x^{2}}{4}+x & \mu+1 \leq x<2 \\ 1 & x \geq 2\end{cases}
$$

The proof is straightforward. Thus, it is omitted.
To simplify our notations, in the rest of the proof, we denote $\delta_{j}$ and $v_{j, 0}$ with $\delta$ and $v$, respectively. In the following, using Lemma B.3.2, we will show that there exists $\hat{v} \in[0, \mu]$ such that for $v \in[0, \hat{v}) W(v)$ is increasing in $v$, and it is non-increasing otherwise. By Lemma B.3.2, for any $v \in[0, \mu)$, we have

$$
\begin{aligned}
\frac{\partial W(v)}{\partial v} & =\mathrm{E}[H(v+\delta)]-H(v)=\frac{1}{2} \int_{-v}^{\mu-v}\left(\frac{3-\mu}{2}(v+\delta)+\frac{(\mu-1)^{2}}{4}\right) d \delta \\
& +\frac{1}{2} \int_{\mu-v}^{1}\left(\frac{1-\mu}{2}(v+\delta)+\mu+\frac{(\mu-1)^{2}}{4}\right) d \delta-\left(\frac{3-\mu}{2} v+\frac{(\mu-1)^{2}}{4}\right) \\
& =\frac{v^{2}(1-\mu)+\left(\mu^{2}+4 \mu-9\right) v-3 \mu^{2}+5 \mu}{8}
\end{aligned}
$$

Since $\frac{\partial W(v)}{\partial v}$ is quadratic in $v$, it is easy to verify that $\frac{\partial W(v)}{\partial v}$ is decreasing in $v \in[0, \mu]$. Then, by the fact that $\left.\frac{\partial W(v)}{\partial v}\right|_{v=0}=\frac{-3 \mu^{2}+5 \mu}{8}>0$ and $\left.\frac{\partial W(v)}{\partial v}\right|_{v=\mu}=\frac{\mu(\mu-2)}{4}<0$, we can conclude that there exists $\hat{v} \in(0, \mu)$ such that $\frac{\partial W(v)}{\partial v} \geq 0$ for $v \in[0, \hat{v}]$, and $\frac{\partial W(v)}{\partial v}<0$ for $v \in(\hat{v}, \mu)$.

The last step of the proof is to show that $W(v)$ is decreasing in $v$ when $v>\mu$. Lemma B.3.2 implies that for any $v>\mu$, we have

$$
\begin{aligned}
\frac{\partial W(v)}{\partial v} & =\mathrm{E}[H(v+\delta)]-H(v)=\frac{1}{2} \int_{-v}^{\mu-v}\left(\frac{3-\mu}{2}(v+\delta)+\frac{(\mu-1)^{2}}{4}\right) d \delta \\
& +\frac{1}{2} \int_{\mu-v}^{1+\mu-v}\left(\frac{1-\mu}{2}(v+\delta)+\mu+\frac{(\mu-1)^{2}}{4}\right) d \delta \\
& +\frac{1}{2} \int_{1+\mu-v}^{1}\left(-\frac{(v+\delta)^{2}}{4}+(v+\delta)\right) d \delta-\left(\frac{1-\mu}{2} v+\mu+\frac{(\mu-1)^{2}}{4}\right) \\
& =\frac{\mu^{3}+12 \mu v-9 \mu^{2}-9 \mu}{24}+\frac{\left(3 v^{2}-3 v-v^{3}\right)}{24} .
\end{aligned}
$$

By the fact that $3 v^{2}-3 v-v^{3}$ is decreasing in $v \in(\mu, 1]$ and $v>\mu$, we have

$$
\frac{\partial W(v)}{\partial v} \leq \frac{\mu^{3}+12 \mu v-9 \mu^{2}-9 \mu}{24}+\frac{\left(3 \mu^{2}-3 \mu-\mu^{3}\right)}{24}=\frac{12 \mu v-6 \mu^{2}-12 \mu}{24} \leq \frac{-6 \mu^{2}}{24}<0,
$$

where the second inequality follows from the fact that $v \leq 1$.

## B.3.3 Proof of Theorems 2.4.4 and 2.5.5

Theorem 2.4.4 can be seen as a corollary of Theorem 2.5.5. Therefore, in the following, we will verify Theorem 2.5.5.

The proof is naturally divided into two parts. In the first part, we show that when the cost is small, there exists an equilibrium in which both agents obtain the additional information, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=1$ for $i=1,2$ and $v_{i, 0} \in[0,1]$. In the second part, we show that when the cost is large, there exists an equilibrium in which none of the agents obtain the additional information, i.e., $\tilde{e}_{i}\left(v_{i, 0}\right)=0$ for $i=1,2$ and $v_{i, 0} \in[0,1]$.

Part 1: We will show that when an agent $i=1,2$ obtains the additional information for any $v_{i, 0} \in$ $[0,1]$, then agent $j \neq i$ also has incentive to obtain information for any $v_{j, 0} \in[0,1]$ as long as cost $c \leq$ $\min \left\{\frac{4 r^{3}-3 r^{2}-6 r+5}{48}, \frac{8 r^{3}+6 r^{2}+7}{96}\right\}$.

With slightly abuse of notation, we define $W_{r}\left(v_{j, 0}\right)$ as the difference of the utility of agent $j=1,2$ with initial valuation $v_{j, 0}$ when he obtains information, and his utility when he does not obtain information given that the reserve price is $r$, and agent $i \neq j$ obtains the additional information for any initial valuation. That is,

$$
W_{r}\left(v_{j, 0}\right)=u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=1, \tilde{e}_{i}\right)-u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=0, \tilde{e}_{i}\right),
$$

where $u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=1, \tilde{e}_{i}\right)$ and $u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=0, \tilde{e}_{i}\right)$ are defined in Equations (B.13) and (B.12). Here, $\tilde{e}_{i}(v)=1$ for any $v \in[0,1]$. We will show that $W_{r}\left(v_{j, 0}\right) \geq 0$ for any $v_{j, 0} \in[0,1]$; that is, for any initial valuation, agent $j$ prefers to update his valuation. To this end, we will verify that $W_{r}(\cdot)$ achieves its minimum at either $v_{j, 0}=0$ or $v_{j, 0}=1$. Then by showing that $W_{r}(1), W_{r}(0) \geq 0$, we can conclude that $W_{r}\left(v_{j, 0}\right) \geq 0$ for any $v_{j, 0} \in[0,1]$.

Similar to the proof of Theorem B.3.1, we write $W_{r}(\cdot)$ as a function of the distribution of the competing bid $B$, where $B=\max \left\{\left(v_{i, 0}+\tilde{e}_{i}\left(v_{i, 0}\right) \delta_{i}\right), r\right\}=\max \left\{v_{i, 0}+\delta_{i}, r\right\}^{1}$. This follows because a second price auction is a truthful mechanism, and agent $i \neq j$ obtains the additional information for any value of $v_{i, 0}$. We denote the distribution of $B$ by $H_{r}$. Then

$$
W_{r}\left(v_{j, 0}\right)=\mathrm{E}\left[\int_{x=r}^{\max \left\{v_{j, 0}+\delta_{j}, r\right\}}\left(v_{j, 0}+\delta_{j}-x\right) d H_{r}(x)\right]-\int_{x=r}^{v_{j, 0}}\left(v_{j, 0}-x\right) d H_{r}(x)-c,
$$

where the expectation is with respect to $\delta_{j}$. We will show that the derivative of $W_{r}\left(v_{j, 0}\right)$ with respect to $v_{j, 0}$ is positive for small values of $v_{j, 0}$ and is non-positive for large values of $v_{j, 0}$. This implies that $W_{r}\left(v_{j, 0}\right) \geq \min \left\{W_{r}(0), W_{r}(1)\right\}$ for any $v_{j, 0} \in[0,1]$. By Leibniz's integral rule, $\frac{\partial W_{r}\left(v_{j, 0}\right)}{\partial v_{j, 0}}$ can be written as $\frac{\partial W_{r}\left(v_{j, 0}\right)}{\partial v_{j, 0}}=\mathrm{E}\left[H_{r}\left(v_{j, 0}+\delta_{j}\right)\right]-H_{r}\left(v_{j, 0}\right)$. The next lemma characterizes the distribution $H_{r}$.

Lemma B.3.3. If $v_{i, 0} \sim u(0,1)$ and $\delta_{i} \sim u(-1,1)$, the distribution of $B=\max \left\{\left(v_{i, 0}+\delta_{i}\right), r\right\}$, denoted by $H_{r}$, is given by

$$
H_{r}(x)= \begin{cases}0 & x<r \\ \frac{1}{4}+\frac{x}{2} & r \leq x \leq 1 ; \\ x-\frac{x^{2}}{4} & 1<x \leq 2 \\ 1 & x>2\end{cases}
$$

The proof is straightforward. Thus, it is omitted.
In the following, using Lemma B.3.3, we will show that $\arg \min _{v \in[0,1]}\left\{W_{r}(v)\right\}$ is either 0 or 1 . To this aim, we will show that $W_{r}(v)$ is increasing in $v$ given that $v<r$, and is decreasing in $v$ otherwise.

Observe that for any $v \in[0, r), \frac{\partial W_{r}(v)}{\partial v}=\mathrm{E}\left[H_{r}(v+\delta)\right] \geq 0$. Furthermore, by Lemma B.3.3, for any $v \in(r, 1]$, we have

$$
\begin{aligned}
\frac{\partial W_{r}(v)}{\partial v} & =\mathrm{E}\left[H_{r}(v+\delta)\right]-H_{r}(v)=\frac{1}{2} \int_{r-v}^{1-v}\left(\frac{1}{4}+\frac{v+x}{2}\right) d x+\frac{1}{2} \int_{1-v}^{1}\left((v+x)-\frac{(v+x)^{2}}{4}\right) d x-\left(\frac{1}{4}+\frac{v}{2}\right) \\
& =\frac{3 v^{2}-v^{3}-3 v}{24}+\frac{-3 r^{2}-3 r}{24} \leq \frac{3 r^{2}-r^{3}-3 r}{24}+\frac{-3 r^{2}-3 r}{24}=\frac{-r^{3}-6 r}{24} \leq 0
\end{aligned}
$$

[^32]The first inequality holds because $3 v^{2}-v^{3}-3 v$ is decreasing in $v$ and $v>r$. We established that $W_{r}(\cdot)$ gets minimized either at $v=0$ or $v=1$. Then, in the last step, we will verify $W_{r}(0), W_{r}(1) \geq 0$. By definition,

$$
\begin{aligned}
W_{r}(0) & =\mathrm{E}\left[(\delta-B)^{+}\right]-c=\operatorname{Pr}[B=r] \times \int_{r}^{1} \frac{1}{2}(\delta-r) d \delta+\frac{1}{4} \int_{r}^{1} \int_{r}^{\delta}(\delta-x) d x d \delta-c \\
& =\frac{4 r^{3}-3 r^{2}-6 r+5}{48}-c
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
W_{r}(1) & =\mathrm{E}\left[(1+\delta-B)^{+}\right]-\mathrm{E}\left[(1-B)^{+}\right]-c \\
& =\operatorname{Pr}[B=r] \times \mathrm{E}\left[(1+\delta-r)^{+}\right]+\frac{1}{4} \int_{\delta=r-1}^{1} \int_{x=r}^{\min \{1+\delta, 1\}}(1+\delta-x) d x d \delta \\
& -\operatorname{Pr}[B=r] \times(1-r)-\frac{1}{2} \int_{x=r}^{1}(1-x) d x=\frac{8 r^{3}+6 r^{2}+7}{96}-c
\end{aligned}
$$

Therefore, when $c \leq \min \left\{\frac{4 r^{3}-3 r^{2}-6 r+5}{48}, \frac{8 r^{3}+6 r^{2}+7}{96}\right\}, W_{r}(v) \geq 0$ for any $v \in[0,1]$. As a result, the agent is willing to obtain information regardless of his initial valuation.

Part 2: In this part, we will show that when the reserve price $r \leq \sqrt{2}-1$ and an agent $i=1,2$ does not obtain the additional information regardless of his initial valuation, then agent $j \neq i$ also does not have any incentive to obtain information for any $v_{j, 0} \in[0,1]$ as long as cost $c \geq \frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}$. Similarly, when $r>\sqrt{2}-1$ and the cost is greater than $\frac{3 r-r^{3}+1}{12}$, there exists an equilibrium such that none of the agents obtain the additional information.

With slightly abuse of notation, we define $W_{r}\left(v_{j, 0}\right)$ as the difference between the utility of agent $j$ when he obtains information and the utility of agent $j$ when he does not obtain information given that agent $i \neq j$ does not acquire the additional information for any initial valuation, i.e., $W_{r}\left(v_{j, 0}\right)=u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=\right.$ $\left.1, \tilde{e}_{i}\right)-u_{j}\left(v_{j, 0}, \tilde{e}_{j}\left(v_{j, 0}\right)=0, \tilde{e}_{i}\right)$, where $\tilde{e}_{i}(v)=0$ for any $v \in[0,1]$. We will show that $W_{r}\left(v_{j, 0}\right) \leq 0$ for any $v_{j, 0} \in[0,1]$ when i- $c \geq \frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}$ and $r \leq \sqrt{2}-1$, or ii- $c \geq \frac{3 r-r^{3}+1}{12}$ and $r>\sqrt{2}-1$ . To this aim, we will verify that $\max _{v \in[0,1]}\left\{W_{r}(v)\right\}=\frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}-c \leq 0$ when $r \leq \sqrt{2}-1$, and $\max _{v \in[0,1]}\left\{W_{r}(v)\right\}=\frac{3 r-r^{3}+1}{12}-c \leq 0$ when $r>\sqrt{2}-1$. This implies that the agent does not have any incentive to update his valuation for any initial valuation.

Specifically, we will show that $W_{r}(v)$ is increasing in $v$ when $v \in\left[0, \max \left\{r, \frac{1-r^{2}}{2}\right\}\right]$, and is decreasing otherwise. That is, $\arg \max _{v \in[0,1]}\left\{W_{r}(v)\right\}=\frac{1-r^{2}}{2}$ if $r \leq \sqrt{2}-1$, and $\arg \max _{v \in[0,1]}\left\{W_{r}(v)\right\}=r$ if $r>\sqrt{2}-1$. We show these statements by characterizing $\frac{\partial W_{r}(v)}{\partial v}$.

By part $1, \frac{\partial W_{r}(v)}{\partial v}=\mathrm{E}\left[H_{r}(v+\delta)\right]-H_{r}(v)$, where $H_{r}(\cdot)$ is the distribution of the bid that agent $j$ competes against. Then, considering the fact that $v_{i, 0} \sim u(0,1)$, and the other agent does not obtain information, i.e., $B=\max \left\{v_{i, 0}, r\right\}$, we have

$$
H_{r}(x)= \begin{cases}0 & x<r \\ x & x \in[r, 1] \\ 1 & x>1\end{cases}
$$

This implies that for any $v \in[0,1]$

$$
\mathrm{E}\left[H_{r}(v+\delta)\right]=\operatorname{Pr}[\delta+v \geq 1]+\frac{1}{2} \int_{r-v}^{1-v}(v+x) d x=\frac{v}{2}+\frac{1-r^{2}}{4} .
$$

Thus, for any $v<r$, we have $\frac{\partial W_{r}(v)}{\partial v}=\mathrm{E}\left[H_{r}(v+\delta)\right]=\frac{v}{2}+\frac{1-r^{2}}{4} \geq 0$ and for any $v>r$,

$$
\frac{\partial W_{r}(v)}{\partial v}=\mathrm{E}\left[H_{r}(v+\delta)\right]-H(v)=\frac{-v}{2}+\frac{1-r^{2}}{4} .
$$

We point that for any $v \geq r$, we have $\frac{\partial W_{r}(v)}{\partial v} \leq 0$ as long as $r>\sqrt{2}-1$. To see why note that

$$
\frac{\partial W_{r}(v)}{\partial v}=\frac{-v}{2}+\frac{1-r^{2}}{4} \leq \frac{-r}{2}+\frac{1-r^{2}}{4} \leq 0
$$

where the last inequality holds because $r>\sqrt{2}-1$. On the other hand, when $r \leq \sqrt{2}-1, \frac{d W_{r}(v)}{d v} \geq 0$ for $v \in\left[r, \frac{1-r^{2}}{2}\right]$ and $\frac{\partial W_{r}(v)}{\partial v} \leq 0$ for $v \geq \frac{1-r^{2}}{2}$. This implies that we have $\arg \max _{v \in[0,1]}\left\{W_{r}(v)\right\}=r$ when $r>\sqrt{2}-1$, and $\arg \max _{v \in[0,1]}\left\{W_{r}(v)\right\}=\frac{1-r^{2}}{2}$ when $r \leq \sqrt{2}-1$. The proof will be completed by showing that i- $W_{r}\left(\frac{1-r^{2}}{2}\right) \leq 0$ when $c \geq \frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}$ and $r \leq \sqrt{2}-1$, and ii- $W_{r}(r) \leq 0$ when $c \geq \frac{3 r-r^{3}+1}{12}$ and $r>\sqrt{2}-1$.

For $r \leq \sqrt{2}-1$, we have

$$
\begin{aligned}
\max _{v \in[0,1]}\left\{W_{r}(v)\right\} & =W_{r}\left(\frac{1-r^{2}}{2}\right)=\mathrm{E}\left[\left(\frac{1-r^{2}}{2}+\delta-B\right)^{+}\right]-\mathrm{E}\left[\left(\frac{1-r^{2}}{2}-B\right)^{+}\right]-c \\
& =\operatorname{Pr}[B=r] \times \mathrm{E}\left[\left(\frac{1-r^{2}}{2}+\delta-r\right)^{+}\right]+\frac{1}{2} \int_{\delta=r-\frac{1-r^{2}}{2}}^{1} \int_{x=r}^{\min \left\{\frac{1-r^{2}}{2}+\delta, 1\right\}}\left(\frac{1-r^{2}}{2}+\delta-x\right) d x d \delta \\
& -\operatorname{Pr}[B=r] \times\left(\frac{1-r^{2}}{2}-r\right)-\int_{x=r}^{\frac{1-r^{2}}{2}}\left(\frac{1-r^{2}}{2}-x\right) d x=\frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}-c \leq 0,
\end{aligned}
$$

where the inequality holds because $c \geq \frac{3 r^{4}+8 r^{3}+6 r^{2}+7}{48}$. Similarly, when $r>\sqrt{2}-1$, we have

$$
\begin{aligned}
\max _{v \in[0,1]}\left\{W_{r}(v)\right\} & =W_{r}(r)=\mathrm{E}\left[(r+\delta-B)^{+}\right]-\mathrm{E}\left[(r-B)^{+}\right]-c \\
& =\operatorname{Pr}[B=r] \times \mathrm{E}\left[(\delta)^{+}\right]+\frac{1}{2} \int_{\delta=0}^{1} \int_{x=r}^{\max \{r+\delta, 1\}}(r+\delta-x) d x d \delta-c=\frac{3 r-r^{3}+1}{12}-c \leq 0,
\end{aligned}
$$

where the inequality follows from the fact that $c \geq \frac{3 r-r^{3}+1}{12}$.

## B. 4 Delegated Proofs

## B.4.1 Proof from Section 2.5

Proof. Proof of Lemma 2.5.2 Let $x$ be the vector of the initial valuations. Given that mechanism $\mathcal{M}^{\text {ort }}$ is incentive compatible, the revenue of the seller is given by

$$
\begin{equation*}
\mathrm{E}\left[\sum_{i=1}^{n} t_{i}+p_{i}\right]=\mathrm{E}\left[\sum_{i=1}^{n} v_{i, 1} q_{i}-c_{i} \times s_{i}-U_{i}\left(x_{i}, x_{i}\right) \mid v_{0}=b_{0}=x\right], \tag{B.14}
\end{equation*}
$$

where $v_{i, 1}=x_{i}+\delta_{i}$ if $i \in \mathcal{S}_{\text {OPT }}(x)$ and $x_{i}$ otherwise, and the expectations are with respect to initial valuations and second signals. Note that the sum of the first and second terms is the social welfare of the agents and the seller. In the following, we compute the last term in the r.h.s. of the above equation, that is, $\mathrm{E}\left[U_{i}\left(x_{i}, x_{i}\right)\right]$, where the expectation is with respect to initial valuations. By Lemma B.1.7, we
have $\mathrm{E}\left[U_{i}\left(x_{i}, x_{i}\right)\right]=\int_{x_{i}=\underline{v}}^{\bar{v}} \int_{z=\underline{v}}^{x_{i}} \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z d F\left(x_{i}\right)$, where the expectation inside the integral is with respect to $v_{-i, 0}$. Changing the order of integrals, we get

$$
\int_{z=\underline{v}}^{\bar{v}} \int_{x_{i}=z}^{\bar{v}} d F\left(x_{i}\right) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z=\int_{z=\underline{v}}^{\bar{v}}(1-F(z)) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z
$$

By multiplying and dividing the r.h.s. of the equation above by the probability density $f(z)$, we obtain $\mathrm{E}\left[U_{i}\left(x_{i}, x_{i}\right)\right]=\mathrm{E}\left[\left.\frac{1-F\left(x_{i}\right)}{f\left(x_{i}\right)} q_{i} \right\rvert\, v_{0}=x\right]=\mathrm{E}\left[-\alpha\left(x_{i}\right) q_{i} \mid v_{0}=x\right]$. Substituting $\mathrm{E}\left[U_{i}\left(x_{i}, x_{i}\right)\right]$ in Eq. (B.14), the expected revenue of the seller is given by $\mathrm{E}\left[\sum_{i=1}^{n}\left(v_{i, 1}+\alpha\left(x_{i}\right)\right) q_{i}-c_{i} \times s_{i} \mid v_{0}=x\right]$. Finally, the result follows from applying the selection and allocation rules.

## Proof. Proof of Lemma 2.5.3

To find an upper bound, we consider a relaxed environment in which the seller observes the additional information of selected agents, and she can force agents to update their valuations. It is easy to see that the maximum achievable revenue in this environment is an upper bound on the revenue of the seller in the original environment.

By the revelation principle, we focus on direct incentive compatible mechanisms that consist of transfer scheme $\bar{t}_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}$, selection rule $\bar{s}_{i}: \mathbb{R}^{n} \rightarrow\{0,1\}$, and allocation rule $\bar{q}_{i}: \mathbb{R}^{2 n} \rightarrow \mathbb{R}^{+}$, where $\bar{s}_{i}$ is 1 when agent $i$ is selected and is 0 otherwise. Note that the payment and selection rules are only functions of initial bids, and the allocation rule is a function of the initial bids and the second signals observed by the seller.

To compute the upper bound on the revenue of any incentive compatible mechanism in the relaxed environment, we first need to characterize the utility of each agent $i$. Assume that agent $i$ with initial valuation $x_{i}$ reports $\hat{x}_{i}$, and other agents report truthfully. Then, his utility is given by

$$
U_{i}\left(x_{i}, \hat{x}_{i}\right)=\mathrm{E}\left[\bar{q}_{i} \times\left(x_{i}+\delta_{i} \bar{s}_{i}\right)-\bar{t}_{i}-c_{i} \bar{s}_{i} \mid b_{i, 0}=\hat{x}_{i}, v_{i, 0}=x_{i}, v_{-i, 0}=b_{-i, 0}=x_{-i}\right]
$$

where the expectation is with respect to the second signals. Incentive compatibility implies that

$$
U_{i}\left(x_{i}, x_{i}\right)-U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right) \leq U_{i}\left(x_{i}, x_{i}\right)-U_{i}\left(\hat{x}_{i}, x_{i}\right)=\left(x_{i}-\hat{x}_{i}\right) \mathrm{E}\left[\bar{q}_{i} \mid b_{i, 0}=x_{i}, b_{-i, 0}=x_{-i}\right]
$$

and similarly, $U_{i}\left(x_{i}, x_{i}\right)-U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right) \geq\left(x_{i}-\hat{x}_{i}\right) \mathrm{E}\left[\bar{q}_{i} \mid b_{i, 0}=\hat{x}_{i}, b_{-i, 0}=x_{-i}\right]$. Without loss of generality, we assume that $x_{i}>\hat{x}_{i}$. Then, using the above Equations,

$$
\mathrm{E}\left[\bar{q}_{i} \mid b_{i, 0}=\hat{x}_{i}, b_{-i, 0}=x_{-i}\right] \leq \frac{U_{i}\left(x_{i}, x_{i}\right)-U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)}{x_{i}-\hat{x}_{i}} \leq E\left[\bar{q}_{i} \mid b_{i, 0}=x_{i}, b_{-i, 0}=x_{-i}\right] .
$$

Finally by taking the limit as $\hat{x}_{i} \rightarrow x_{i}^{-}$, we get $U_{i}\left(x_{i}, x_{i}\right)=U_{i}(\underline{v}, \underline{v})+\int_{\underline{v}}^{x_{i}} \mathrm{E}\left[\bar{q}_{i} \mid v_{i, 0}=z, b_{-i, 0}=x_{-i}\right] d z$.
We are now ready to compute the upper bound of the revenue. By using the same arguments as in Lemma 2.5.2, it can be shown that for any selection rule $\bar{s}_{i}$ and allocation rule $\bar{q}_{i}$ revenue of the seller when agents are truthful is given by

$$
\begin{equation*}
\mathrm{E}\left[\sum_{i: \bar{s}_{i}(x)=1}\left(x_{i}+\alpha\left(x_{i}\right)+\delta_{i}\right) \bar{q}_{i}+\sum_{i: \bar{s}_{i}(x)=0}\left(x_{i}+\alpha\left(x_{i}\right)\right) \bar{q}_{i}-\sum_{i=1}^{n} c_{i} \bar{s}_{i}-\sum_{i=1}^{n} U_{i}(\underline{v}, \underline{v}) \mid b_{0}=x\right], \tag{B.15}
\end{equation*}
$$

where the expectation is with respect to the first and second signals. Because the mechanism should be individually rational, we set $U_{i}(\underline{v}, \underline{v})=0$ for all $i$. Then, to maximize the revenue, the item should be allocated to the agent with the highest non-negative virtual valuation, that is,

$$
\bar{q}_{i}=1 \quad \text { if } i \in \arg \max _{j}\left\{\left(x_{j}+\alpha\left(x_{j}\right)+\delta_{j} \bar{s}_{j}(x)\right)\right\}^{2}
$$

and 0 otherwise. Therefore, the expected revenue of the seller can be written as

$$
\mathrm{E}\left[\max \left\{\max _{i: \overline{\bar{s}}_{i}(x)=1}\left\{x_{i}+\alpha\left(x_{i}\right)+\delta_{i}\right\}, \max _{i: \bar{s}_{i}(x)=0}\left\{x_{i}+\alpha\left(x_{i}\right)\right\}, 0\right\}-\sum_{i=1}^{n} c_{i} \bar{s}_{i} \mid b_{0}=x\right] .
$$

So, if agents in set $\mathcal{S}_{\text {relaxed }}$, defined below, obtain the additional information, the revenue gets maximized:

$$
\mathcal{S}_{\text {relaxed }}(x)=\arg \max _{S \subset\{1, \cdots, n\}}\left\{\mathrm{E}\left[\max \left\{\max _{i \in S}\left\{x_{i}+\delta_{i}+\alpha\left(x_{i}\right)\right\}, \max _{i \notin S}\left\{x_{i}+\alpha\left(x_{i}\right)\right\}, 0\right\}\right]-\sum_{i \in S} c_{i}\right\} .
$$

Finally, the result follows by plugging the selection rule $\mathcal{S}_{\text {relaxed }}(x)$ and allocation rule $\bar{q}_{i}$ in Eq. (B.15).

## B.4.2 Proofs from Section B. 1

## Proof. Proof of Lemma B.1.1

[^33]The basic idea is to establish an upper bound on the r.h.s. of Eq. (B.3). We will show the upper bound is larger than the l.h.s. of Eq. (B.3) at $r=\rho$. This will imply that Eq. (B.3) is satisfied at some $r \geq \rho$.

We first show that the r.h.s. of Eq. (B.3), i.e., $\int_{v}^{b_{\ell, 0}} \mathrm{E}\left[q_{\ell} \mid v_{\ell, 0}=z, v_{-\ell, 0}=b_{-\ell, 0}\right] d z$, is less than or equal to $\int_{\underline{v}}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z$, where $\omega_{-\ell}=\max \left\{\max _{j \in \mathcal{S}_{\rho, \beta}\left(b_{0}\right)}\left\{b_{j, 0}+\delta_{j}+\beta\left(b_{j, 0}\right)\right\}, \rho\right\}$. Because agent $\ell$ has the largest weighted bid among unselected agents, $\omega_{-\ell}=\max \left\{\max _{j \in \mathcal{S}_{\rho, \beta}\left(b_{0}\right)}\left\{b_{j, 0}+\delta_{j}+\right.\right.$ $\left.\left.\beta\left(b_{j, 0}\right)\right\}, \max _{j \notin \mathcal{S}_{\rho, \beta}\left(b_{0}\right), j \neq \ell}\left\{b_{j, 0}+\beta\left(b_{j, 0}\right)\right\}, \rho\right\}$. Then, by Lemma B.1.8,

$$
\begin{aligned}
U_{\ell}\left(\underline{v}, b_{\ell, 0}\right) & \leq U_{\ell}\left(b_{\ell, 0}, b_{\ell, 0}\right)-\int_{\underline{v}}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z \leq U_{\ell}(\underline{v}, \underline{v}) \\
& =U_{\ell}\left(b_{\ell, 0}, b_{\ell, 0}\right)-\int_{\underline{v}}^{b_{\ell, 0}} \mathrm{E}\left[q_{\ell} \mid v_{\ell, 0}=z, v_{-\ell, 0}=b_{-\ell, 0}\right] d z,
\end{aligned}
$$

where the second inequality follows from Lemma B.1.9, and the equality follows from Lemma B.1.7. By the above equation, we can conclude that $\int_{\underline{v}}^{b_{\ell, 0}} \mathrm{E}\left[q_{\ell} \mid v_{\ell, 0}=z, v_{-\ell, 0}=b_{-\ell, 0}\right] d z$ is less than $\int_{\underline{v}}^{b_{\ell, 0}} \operatorname{Pr}[z+$ $\left.\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z$.

Next, we will show that the 1.h.s. of Eq. (B.3) at $r=\rho$ is greater than the upper bound, i.e., $\int_{\underline{v}}^{b_{e, 0}} \operatorname{Pr}[z+$ $\left.\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z$. Then, considering the fact that the upper bound is not a function of $r$, the l.h.s. is a non-increasing function of $r$, and is zero at $r=b_{\ell, 0}+\beta\left(b_{\ell, 0}\right)$, we conclude that there exists an $r \in$ [ $\left.\rho, b_{\ell, 0}+\beta\left(b_{\ell, 0}\right)\right]$ that satisfies Eq. (B.3).

By changing variable, the l.h.s. at $r=\rho$ can be written as

$$
\int_{\rho-\beta\left(b_{\ell, 0}\right)}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \max \left\{\max _{j \in \mathcal{S}_{\rho, \beta}\left(b_{0}\right)}\left\{b_{j, 0}+\delta_{j}+\beta\left(b_{j, 0}\right)\right\}, \rho\right\}\right] d z=\int_{\rho-\beta\left(b_{\ell, 0}\right)}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z,
$$

Then, because $b_{\ell, 0}+\beta\left(b_{\ell, 0}\right) \geq \rho$, we have

$$
\begin{aligned}
\int_{\rho-\beta\left(b_{\ell, 0}\right)}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq\right. & \left.\omega_{-\ell}\right] d z \geq \int_{\max \left\{\rho-\beta\left(b_{\ell, 0}\right), v\right\}}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z \\
& =\int_{\underline{v}}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z-\int_{\underline{v}}^{\max \left\{\rho-\beta\left(b_{\ell, 0}\right), \underline{v}\right\}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z \\
& =\int_{\underline{v}}^{b_{\ell, 0}} \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right] d z,
\end{aligned}
$$

where the last equality holds because $\omega_{-\ell} \geq \rho$ and for any $z \leq \rho-\beta\left(b_{\ell, 0}\right), \operatorname{Pr}\left[z+\beta\left(b_{\ell, 0}\right) \geq \omega_{-\ell}\right]=0$. The last equation shows that the 1.h.s. of Eq. (B.3) at $r=\rho$ is greater than the r.h.s. of Eq. (B.3).

## B.4.3 Proofs from Section B.1.1

Proof. Proof of Lemma B.1.5 Consider a selected agent $i$ with initial bid $b_{i, 0}$. If agent $i$ wins the item, his payment in the second round, $p_{i}$, would be equal to $\max \left\{\max _{j \neq i}\left\{b_{j, 1}+\beta\left(b_{j, 0}\right)\right\}, \rho\right\}-\beta\left(b_{i, 0}\right)$. Note that $p_{i}$ is independent of $b_{i, 1}$. Therefore, an agent cannot change his price for the item. However, the agent can change the probability of the allocation. It is easy to see that underbidding may only result in losing the item. On the other hand, over bidding may yield a negative utility. Note that overbidding can make a difference only when

$$
v_{i, 1}+\beta\left(b_{i, 0}\right) \leq \max \left\{\max _{j \neq i}\left\{b_{j, 1}+\beta\left(b_{j, 0}\right)\right\}, \rho\right\} \leq b_{i, 1}+\beta\left(b_{i, 0}\right) .
$$

In this case, the utility of agent $i$ would be non-positive:

$$
v_{i, 1}-p_{i}=v_{i, 1}+\beta\left(b_{i, 0}\right)-\max \left\{\max _{j \neq i}\left\{b_{j, 1}+\beta\left(b_{j, 0}\right)\right\}, \rho\right\} \leq v_{i, 1}+\beta\left(b_{i, 0}\right)-\left(v_{i, 1}+\beta\left(b_{i, 0}\right)\right)=0
$$

Therefore, a weekly dominate strategy of agent $i$ is to be truthful.

## Proof. Proof of Lemma B.1. 6

Consider agent $i \in \mathcal{S}_{\rho, \beta}(x)$ who bids truthfully in the first round, where $x$ is the initial valuations of agents. In the proof, to simplify the notations, we denote $\mathcal{S}_{\rho, \beta}(x)$ by $\mathcal{S}(x)$. We will show that agent $i$ will learn his second signal, i.e., $e_{i}=1$, given that other agents are truthful. To this aim, we will prove that for agent $i$, the marginal value of changing his decision to obtain information is identical to the change in the weighted surplus $\Omega_{\rho, \beta}$. More precisely, the difference between the utility of agent $i$ when he obtains information and that when he does not is equal to $\Omega_{\rho, \beta}(x, \mathcal{S}(x))-\Omega_{\rho, \beta}(x, \mathcal{S}(x) \backslash\{i\})$. Then, since the SWSP mechanism maximizes the weighted surplus, i.e., $\Omega_{\rho, \beta}(x, \mathcal{S}(x)) \geq \Omega_{\rho, \beta}(x, \mathcal{S}(x) \backslash\{i\})$, we conclude that agent $i$ prefers to update his valuation.

We first characterize the utility of agent $i$ when he does not obtain information. Let $Y$ be the random variable corresponding to the maximum weighted bid of all agents except for agent $i$, i.e., $Y=$
$\max \left\{\max _{j \in \mathcal{S}(x), j \neq i}\left\{x_{j}+\delta_{j}+\beta\left(x_{j}\right)\right\}, \max _{j \notin \mathcal{S}(x)}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, \rho\right\}$. Then, by Lemma B.1.5, when agent $i$ does not update his valuation, his utility is given by

$$
\begin{align*}
\mathrm{E}\left[\max \left\{x_{i}+\beta\left(x_{i}\right)-Y, 0\right\}\right]-t_{i} & =\mathrm{E}\left[\max \left\{Y, x_{i}+\beta\left(x_{i}\right)\right\}-Y\right]-t_{i}, \\
& =\Omega_{\rho, \beta}(x, \mathcal{S}(x) \backslash\{i\})+\sum_{j \in \mathcal{S}(x) \backslash\{i\}} c_{j}-\mathrm{E}[Y]-t_{i} \tag{B.16}
\end{align*}
$$

Note that in computing utility of agent $i$, we use the fact that other agents are truthful. In addition, the second equality holds because

$$
\begin{aligned}
& \Omega_{\rho, \beta}(x, \mathcal{S}(x) \backslash\{i\}) \\
& =\mathrm{E}\left[\max \left\{\max _{j \in \mathcal{S}(x), j \neq i}\left\{x_{j}+\beta\left(x_{j}\right)+\delta_{j}\right\}, \max _{j \notin \mathcal{S}(x)}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, x_{i}+\beta\left(x_{i}\right), \rho\right\}\right]-\sum_{j \in \mathcal{S}(x) \backslash\{i\}} c_{j} \\
& =\mathrm{E}\left[\max \left\{Y, x_{i}+\beta\left(x_{i}\right)\right\}\right]-\sum_{j \in \mathcal{S}(x) \backslash\{i\}} c_{j},
\end{aligned}
$$

Similarly, when agent $i$ obtains the additional information, his utility is equal to

$$
\begin{align*}
\mathrm{E}\left[\max \left\{x_{i}+\delta_{i}+\beta\left(x_{i}\right)-Y, 0\right\}\right]-t_{i}-c_{i} & =\mathrm{E}\left[\max \left\{Y, x_{i}+\delta_{i}+\beta\left(x_{i}\right)\right\}-Y\right]-t_{i}-c_{i} \\
& =\Omega_{\rho, \beta}(x, \mathcal{S}(x))+\sum_{j \in \mathcal{S}(x) \backslash\{i\}} c_{j}-\mathrm{E}[Y]-t_{i} . \tag{B.17}
\end{align*}
$$

The first expression follows from Lemma B. 1.5 where we show that agent $i$ bids truthfully in the second round. The second equality holds because

$$
\begin{aligned}
\Omega_{\rho, \beta}(x, \mathcal{S}(x))= & \mathrm{E}\left[\max \left\{\max _{j \in \mathcal{S}(x), j \neq i}\left\{x_{j}+\beta\left(x_{j}\right)+\delta_{j}\right\}, \max _{j \notin \mathcal{S}(x)}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, x_{i}+\delta_{i}+\beta\left(x_{i}\right), \rho\right\}\right]-\sum_{j \in \mathcal{S}(x)} c_{j} \\
& =\mathrm{E}\left[\max \left\{Y, x_{i}+\delta_{i}+\beta\left(x_{i}\right)\right\}\right]-\sum_{j \in \mathcal{S}(x)} c_{j} .
\end{aligned}
$$

In addition, note that $t_{i}$ in Eq. (B.17) is the same as $t_{i}$ in Eq. (B.16) since agent $i$ 's decision to obtain information, i.e., $e_{i}$, is not observable by the mechanism. By Equations (B.16) and (B.17), the difference between the utility of agent $i$ when he updates his valuation and his utility when he does not is $\Omega_{\rho, \beta}(x, \mathcal{S}(x))-\Omega_{\rho, \beta}(x, \mathcal{S}(x) \backslash\{i\})$. Then, considering the fact that the SWSP mechanism maximizes the
weighted surplus, i.e., $\Omega_{\rho, \beta}(x, \mathcal{S}(x)) \geq \Omega_{\rho, \beta}(x, \mathcal{S}(x) \backslash\{i\})$, we conclude that agent $i$ prefers to learn his second signal.

Proof. Proof of Lemma B.1.7 We first show that for any $x_{-i} \in[\underline{v}, \bar{v}]^{n-1}$ and any $i=1,2, \ldots, n$, $\mathrm{E}\left[q_{i} \mid v_{i, 0}=x_{i}, v_{-i, 0}=x_{-i}\right]$ is a non-decreasing function of $x_{i}$. Observe that the weighted surplus is the maximum of affine functions of $v_{i, 0}+\beta\left(v_{i, 0}\right)$. Thus, it is a convex function of $v_{i, 0}+\beta\left(v_{i, 0}\right)$. Furthermore, the weighted surplus is a continuos function of $v_{i, 0}+\beta\left(v_{i, 0}\right)$, and its derivative with respect to $v_{i, 0}+\beta\left(v_{i, 0}\right)$ at $v_{i, 0}=x_{i}$, if exists, is equal to $\mathrm{E}\left[q_{i} \mid v_{i, 0}=x_{i}, v_{-i, 0}=x_{-i}\right] .^{3}$ This implies that $\mathrm{E}\left[q_{i} \mid v_{i, 0}=x_{i}, v_{-i, 0}=x_{-i}\right]$ is a non-decreasing function of $v_{i, 0}+\beta\left(v_{i, 0}\right)$. Finally, considering the fact function $\beta$ is non-decreasing, we can conclude that $\mathrm{E}\left[q_{i} \mid v_{i, 0}=x_{i, 0}, v_{-i, 0}=x_{-i}\right]$ is a non-decreasing function of $v_{i, 0}$.

Next we show that the utility of an agent $i$ that bids truthfully in the first round and follows the optimal strategy afterwards follows from Eq. (B.7) when all other agents are truthful. By Lemma B.1.6, agent $i$ that bids truthfully in the first round will obtain information if he gets selected. Furthermore, Lemma B.1.5 implies that agent $i$ will bid truthfully in the second round if he is allowed to update his bid in the second round. Therefore, agent $i$ that bids truthfully in the first round stays truthful.

Now, we are ready to show the result. We consider the following cases. Throughout this proof, for simplicity, we drop the subscript of $\mathcal{S}_{\rho, \beta}(x)$, and denote it by $\mathcal{S}(x)$.
i) $i \in \mathcal{S}(x)$ : By lemmas B.1.5 and B.1.6, selected agent $i$ learns his second signal and reports it truthfully in the second round. Thus, his utility is given by $\mathrm{E}\left[q_{i}\left(x_{i}+\delta_{i}\right)-p_{i}-t_{i}-c_{i} \mid v_{0}=x_{0}\right]$. The claim follows from plugging $t_{i}$ from Eq. (2.2).
ii) $i \notin \mathcal{S}(x)$ and $x_{i}+\beta\left(x_{i}\right)<\max \left\{\max _{j \notin \mathcal{S}(x), j \neq i}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, \rho\right\}$ : In this case, the utility of agent $i$ and his allocation probability is 0 . By the fact that $\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right]$ is an increasing function of $z$ and $\mathrm{E}\left[q_{i} \mid v_{i, 0}=x_{i}, v_{-i, 0}=x_{-i}\right]=0$, we can write the utility of agent $i$ as $\int_{\underline{v}}^{x_{i}} \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z=0$.
iii) $i \notin \mathcal{S}(x)$ and $x_{i}+\beta\left(x_{i}\right) \geq \max \left\{\max _{j \notin \mathcal{S}(x), j \neq i}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, \rho\right\}$ : In this case, when unselected agent $i$ wins the item, he has to pay maximum of $r$ and the second highest weighted bid. Therefore,

[^34]$U_{i}\left(x_{i}, x_{i}\right)=\mathrm{E}\left[q_{i} \times\left(x_{i}+\beta\left(x_{i}\right)-\max \left\{\max _{j \in \mathcal{S}(x)}\left\{x_{j}+\delta_{j}+\beta\left(x_{j}\right)\right\}, r\right\}\right)\right]$, where $U_{i}\left(x_{i}, x_{i}\right)$ is defined in Eq. (B.6).

Let $Y=\max _{j \in \mathcal{S}(x)}\left\{x_{j}+\delta_{j}+\beta\left(x_{j}\right)\right\}$, and let $H$ be the distribution of $Y$. Then, $U_{i}\left(x_{i}, x_{i}\right)$ can be written as

$$
\begin{gathered}
\mathrm{E}\left[\left(x_{i}+\beta\left(x_{i}\right)-Y\right) \times \mathbb{1}\left\{x_{i}+\beta\left(x_{i}\right) \geq Y \geq r\right\}+\left(x_{i}+\beta\left(x_{i}\right)-r\right) \times \mathbb{1}\left\{x_{i}+\beta\left(x_{i}\right) \geq r \geq Y\right\}\right] \\
=\left(x_{i}+\beta\left(x_{i}\right)\right)\left(H\left(x_{i}+\beta\left(x_{i}\right)\right)-H(r)\right)-\int_{r}^{x_{i}+\beta\left(x_{i}\right)} z d H(z)+\left(x_{i}+\beta\left(x_{i}\right)-r\right) H(r) \\
\quad=\left(x_{i}+\beta\left(x_{i}\right)\right) H\left(x_{i}+\beta\left(x_{i}\right)\right)-r H(r)-\int_{r}^{x_{i}+\beta\left(x_{i}\right)} z d H(z)=\int_{r}^{x_{i}+\beta\left(x_{i}\right)} H(z) d z,
\end{gathered}
$$

where in the first equation, the expectation is with respect to the second signals. The last equality is followed from the integration by part. Therefore, using Eq. (B.3), we get

$$
U_{i}\left(x_{i}, x_{i}\right)=\int_{r}^{x_{i}+\beta\left(x_{i}\right)} \operatorname{Pr}\left[z \geq \max _{j \in \mathcal{S}(x)}\left\{x_{j}+\delta_{j}+\beta\left(x_{j}\right)\right\}\right] d z=\int_{\underline{v}}^{x_{i}} \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z .
$$

Proof. Proof of Lemma B.1.8 Throughout the proof, all the expectations are with respect to the second signals. Consider an untruthful agent $i$ with initial valuation $x_{i}$ who bids $\hat{x}_{i}$ in the first round. We establish an upper bound on his utility. We consider the following two cases, $s_{i}=1$ and $s_{i}=0$.
$s_{i}=1$ : When agent $i$ is selected, $s_{i}=1$, he can either obtain information or not. Given his investing decision $e_{i}=e_{i}\left(v_{i, 0}=x_{i}, b_{i, 0}=\hat{x}_{i}, t_{i}\right)$, by Lemma B.1.5, his utility can be written as

$$
\begin{align*}
U_{i}\left(x_{i}, \hat{x}_{i}\right) & =\mathrm{E}\left[\left(x_{i}+e_{i} \delta_{i}\right)-p_{i}-t_{i}-e_{i} c_{i} \mid v_{0}=x, b_{i, 0}=\hat{x}_{i}, b_{-i, 0}=x_{-i}\right]  \tag{B.18}\\
& =\mathrm{E}\left[\max \left\{x_{i}+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i}-\omega_{-i}, 0\right\}-t_{i}-e_{i} c_{i}\right],
\end{align*}
$$

where $\omega_{-i}$ is defined in Eq. (B.8), and the expectation is taken assuming the all agents except for agent $i$ are truthful. Note that for abbreviation, we omit the condition in the second equation and in the rest of the
proof. Then, by adding and subtracting $\mathrm{E}\left[\max \left\{\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i}-\omega_{-i}, 0\right\}\right]$, the utility can be rewritten as

$$
\begin{aligned}
U_{i}\left(x_{i}, \hat{x}_{i}\right)= & \mathrm{E}\left[\max \left\{\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i}-\omega_{-i}, 0\right\}-t_{i}-e_{i} c_{i}\right. \\
& \left.-\left(\max \left\{\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i}-\omega_{-i}, 0\right\}-\max \left\{x_{i}+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i}-\omega_{-i}, 0\right\}\right)\right] \\
= & \mathrm{E}\left[\max \left\{\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i}-\omega_{-i}, 0\right\}-t_{i}-e_{i} c_{i}\right]+\int_{\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left[z+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i} \geq \omega_{-i}\right] d z
\end{aligned}
$$

When $e_{i}=1$ the first term in the last line is $U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)$. Otherwise, it is the utility of selected agent $i$ with initial $\hat{x}_{i}$ who bids truthfully, gets selected, but does not learn his second signal, which is by Lemma B.1.6 is less than or equal to $U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)$. Thus, the utility is at most $U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)+\int_{\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left[z+\beta\left(\hat{x}_{i}\right)+e_{i} \delta_{i} \geq \omega_{-i}\right] d z$, which is the desired result.
$s_{i}=0: \quad$ Note that $s_{i}=0$ means agent $i$ is not selected. Then, if $\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)<$ $\max \left\{\max _{j \notin \mathcal{S}_{\rho, \beta}\left(\hat{x}_{i}, x_{-i}\right)}\left\{x_{j}+\beta\left(x_{j}\right)\right\}, \rho\right\}$, his utility is zero. If not, the utility of agent $i$ given that he stays in the game can be written as

$$
\mathrm{E}\left[q_{i} \times\left(x_{i}+\beta\left(\hat{x}_{i}\right)-\max \left\{\omega_{-i}, r\right\}\right)\right]
$$

By adding and subtracting $\mathrm{E}\left[q_{i} \times \hat{x}_{i}\right]$, and by the fact that agent $i$ receives the item if $\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)$ is greater than $\omega_{-i}$, we have

$$
\mathrm{E}\left[\max \left\{\left(\hat{x}_{i}+\beta\left(\hat{x}_{i}\right)-\max \left\{\omega_{-i}, r\right\}, 0\right\}+\left(x_{i}-\hat{x}_{i}\right) \times \mathbb{1}\left\{\hat{x}_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right\}\right]\right.
$$

The first term is $U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)$. Since $U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right) \geq 0$ and agent $i$ can exit the game if his utility gets negative, he can at most yield $\max \left\{U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)+\int_{\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left\{\hat{x}_{i}+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right\} d z, 0\right\}$, which is less than $\max \left\{U_{i}\left(\hat{x}_{i}, \hat{x}_{i}\right)+\right.$ $\left.\int_{\hat{x}_{i}}^{x_{i}} \operatorname{Pr}\left\{z+\beta\left(\hat{x}_{i}\right) \geq \omega_{-i}\right\} d z, 0\right\}$.

## B.4.4 Proofs from Section B.1.2

Proof. Proof of Lemma B.1.10 We establish the following two claims.

- Claim 1: For any set $S \subseteq\{1,2, \ldots, n\}$

$$
\Omega_{\rho, \beta}\left(\left(x_{i}, x_{-i}\right), S\right)-\Omega_{\rho, \beta}\left(\left(\hat{x}_{i}, x_{-i}\right), S\right)=\int_{\hat{x}_{i}}^{x_{i}}\left(1+\beta^{\prime}(z)\right) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, S\right] d z
$$

- Claim 2: Suppose Assumption 4 holds. Then weighted surplus is an absolutely continuous and convex function of $v_{i, 0}+\beta\left(v_{i, 0}\right)$, and it is given by

$$
\Omega_{\rho, \beta}(x)=\Omega_{\rho, \beta}\left(\left(\hat{x}_{i}, x_{-i}\right)\right)+\int_{\hat{x}_{i}}^{x_{i}}\left(1+\beta^{\prime}(z)\right) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z
$$

The proof of claims follows from Theorems 1 and 2 in Milgrom \& Segal (2002). Thus, we do not repeat it here. By Claims 1 and 2 and the fact that $\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)\right]=$ $\mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right]$ for $z<y_{1}$ and $z>y_{2}$, we have

$$
\begin{aligned}
\Omega_{\rho, \beta}\left(x, \hat{\mathcal{S}}_{y_{1}, y_{2}}(x)\right) & =\Omega_{\rho, \beta}\left(\left(\hat{x}_{i}, x_{-i}\right), \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(\hat{x}_{i}, x_{-i}\right)\right)+\int_{\hat{x}_{i}}^{y_{1}}\left(1+\beta^{\prime}(z)\right) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z \\
& +\int_{y_{1}}^{y_{2}}\left(1+\beta^{\prime}(z)\right) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}, \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(z, x_{-i}\right)\right] d z \\
& +\int_{y_{2}}^{x_{i}}\left(1+\beta^{\prime}(z)\right) \mathrm{E}\left[q_{i} \mid v_{i, 0}=z, v_{-i, 0}=x_{-i}\right] d z
\end{aligned}
$$

Then, the result follows from Claim 2 and the fact that, by construction, $\Omega_{\rho, \beta}\left(\hat{x}_{i}, x_{-i}\right)=$ $\Omega_{\rho, \beta}\left(\left(\hat{x}_{i}, x_{-i}\right), \hat{\mathcal{S}}_{y_{1}, y_{2}}\left(\hat{x}_{i}, x_{-i}\right)\right)$.

## B. 5 Numerical Experiments

In Section B.5.1, we depict the initial payments. Section B.5.2 compares the SSP mechanism with the optimal mechanism in terms of the revenue of the seller. In Section B.5.3, we study impacts of the number of agents, $n$.

## B.5.1 Payments in the First Round

Recall that the initial payment $t_{i}$ incentivizes agents to be truthful. In this section, we investigate how much the OPT and EFF mechanisms charge each agent $i$ upfront for different realizations of initial valuations. As usual, $n=2, F=N(0.5,0.5), G_{i}=N(0,0.5)$, and $c_{i}=0.05$ for $i=1,2$.

The initial payment for the first agent, $t_{1}$, in the OPT and EFF mechanisms for all realizations of $v_{1,0}$ and $v_{2,0}$ in the range of $[-1.5,2.5]$ is shown in Figures B. 1 and B.2, respectively. The x -axis is $v_{2,0}$, and the y -axis is $v_{1,0}$. Here, different shades of gray mean different initial payment as defined in the color bars next to the figures. By construction, the initial payment of the first agent is zero if he is not selected. Furthermore, when he is selected, $t_{1}$ is an increasing function of $v_{1,0}$.


Figure B.1: The payment of agent 1 in the first round, $t_{1}$, in the optimal mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2, c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.

## B.5.2 The SSP Mechanism versus the optimal Mechanism

In this section, we seek to understand how the SSP mechanism performs in compare with mechanism $\mathcal{M}^{\text {opr }}$. To this aim, we report the revenue of the SSP mechanism under four problem classes, corresponding with cost 0.02 and 0.05 , and number of agents of 2 and 3 . Here, $F=N(0.5,0.5)$ and $G_{i}=N\left(0, \sigma^{2}\right)$, where $\sigma^{2}=0.5,1,1.5$, and 2.

In Table B.1, for each problem class, we present the revenue of the SSP mechanism with revenuemaximizing $r$ as a percentage of the optimal revenue, averaged over 2000 instances in each problem class. We observe that the SSP mechanism performs better as the number of agents gets larger, $\sigma^{2}$ becomes smaller, and the additional information gets more costly. In addition, the SSP mechanism yields more than $84 \%$ of the optimal revenue.


Figure B.2: The payment of agent 1 in the first round, $t_{1}$, in the efficient mechanism for different realizations of $v_{1,0}$ and $v_{2,0}$ with $n=2, c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.

| Problem Class |  | $G_{i}=N\left(0, \sigma^{2}\right)$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $n$ | cost | $\sigma^{2}=0.5$ | $\sigma^{2}=1$ | $\sigma^{2}=1.5$ | $\sigma^{2}=2$ |
| 2 | 0.02 | 94 | 90 | 87 | 84 |
|  | 0.05 | 95 | 92 | 90 | 87 |
| 3 | 0.02 | 95 | 93 | 89 | 88 |
|  | 0.05 | 96 | 94 | 92 | 90 |

Table B.1: Revenue of the SSP mechanism (with revenue-maximizing $r$ ) as a percentage of the optimal revenue with $F=N(0.5,0.5)$. Here, the standard errors of all numbers are less than $1 \%$.

## B.5.3 More Agents

In this section, we investigate how the number of agents can affect the outcome of the OPT and EFF mechanisms. Again, $F=N(0.5,0.5), G_{i}=N(0,0.5)$, and $c_{i}=0.05$.

Figure ?? shows the average number of selected agents, revenue, and social welfare versus the number of agents, $n$. As the number of agents increases, the revenue and social welfare, and average number of selected agents in all considered mechanisms rise. However, even in the EfF mechanism, the average number of selected agents is sub-linear (concave) in $n$.


Figure B.3: Revenue in the optimal and efficient mechanisms versus number of agents, $n$, with $c=0.05$, $F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.


Figure B.4: Social welfare in the optimal and efficient mechanisms versus number of agents, $n$, with $c=$ $0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.


Figure B.5: The average number of selected agents of the optimal and efficient mechanisms versus number of agents, $n$, with $c=0.05, F=N(0.5,0.5)$, and $G_{i}=N(0,0.5)$.

## Appendix C

## Technical Appendix to Chapter 3

## C. 1 Proof of Main Results

## C.1.1 Appendix to Section 3.2

Proof. Proof of Lemma 3.2.1 The proof falls naturally into two parts. In the first part, we show that in an incentive compatible mechanism conditions in Equations (3.1) and (3.2) hold. In the second part, we show that if Equations (3.1) and (3.2) hold, the mechanism is incentive compatible.

First Part: Consider a customer with type $\theta$ that reports $\hat{\theta}$. Without loss of generality, we assume that $\theta \geq \hat{\theta}$. Then, the utility of the customer is given by $u(\theta, \hat{\theta})=V(\theta, \mathfrak{t}(\hat{\theta}))-\mathfrak{p}(\hat{\theta})$. Incentive compatibility implies that

$$
\begin{align*}
u(\theta, \theta)-u(\hat{\theta}, \hat{\theta}) & \leq u(\theta, \theta)-u(\hat{\theta}, \theta) \\
& =(V(\theta, \mathfrak{t}(\theta))-V(\hat{\theta}, \mathfrak{t}(\theta)))=\int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(\theta)) d z \tag{C.1}
\end{align*}
$$

and

$$
\begin{align*}
u(\theta, \theta)-u(\hat{\theta}, \hat{\theta}) & \geq u(\theta, \hat{\theta})-u(\hat{\theta}, \hat{\theta}) \\
& =V(\theta, \mathfrak{t}(\hat{\theta}))-V(\hat{\theta}, \mathfrak{t}(\hat{\theta}))=\int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(\hat{\theta}) d z \tag{C.2}
\end{align*}
$$

where $\partial_{1} V(\theta, t)=\frac{\partial V(\theta, t)}{\partial \theta}$. Then, using the above Equations, we have

$$
\begin{aligned}
& \frac{V(\theta, \mathfrak{t}(\hat{\theta}))-V(\hat{\theta}, \mathfrak{t}(\hat{\theta}))}{\theta-\hat{\theta}} \leq \frac{u(\theta, \theta)-u(\hat{\theta}, \hat{\theta})}{\theta-\hat{\theta}}, \\
& \frac{V(\theta, \mathfrak{t}(\theta))-V(\hat{\theta}, \mathfrak{t}(\theta))}{\theta-\hat{\theta}} \geq \frac{u(\theta, \theta)-u(\hat{\theta}, \hat{\theta})}{\theta-\hat{\theta}} .
\end{aligned}
$$

Finally by taking the limit as $\hat{\theta} \rightarrow \theta^{-}$, we get Eq. (3.1). Then, by Equations (3.1), (C.1), and (C.2), we get the second condition, given in Eq. (3.2).

Second Part: Here, we will show that if in a mechanism Equations (3.1) and (3.2) hold, the mechanism is incentive compatible. By Eq. (3.1),

$$
\begin{align*}
u(\theta, \theta)-u(\hat{\theta}, \hat{\theta})=\int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(z)) d z & \geq \int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(\hat{\theta})) d z \\
& =V(\theta, \mathfrak{t}(\hat{\theta}))-V(\hat{\theta}, \mathfrak{t}(\hat{\theta}))=u(\theta, \hat{\theta})-u(\hat{\theta}, \hat{\theta}) \tag{C.3}
\end{align*}
$$

where the inequality follows from Eq. (3.2). The final equation implies that $u(\theta, \theta) \geq u(\theta, \hat{\theta})$. Similarly,

$$
\begin{aligned}
u(\theta, \theta)-u(\hat{\theta}, \hat{\theta})=\int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(z)) d z & \leq \int_{z=\hat{\theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(\theta)) d z \\
& =V(\theta, \mathfrak{t}(\theta))-V(\hat{\theta}, \mathfrak{t}(\theta))=u(\theta, \theta)-u(\hat{\theta}, \theta)
\end{aligned}
$$

That is, $u(\hat{\theta}, \hat{\theta}) \geq u(\hat{\theta}, \theta)$. The above equation along with Eq. (C.3) imply that the mechanism is incentive compatible.

Proof. Proof of Lemma 3.2.2 Consider any incentive compatible mechanism. Then, the expected revenue of the firm from selling one unit of the item is given by

$$
\begin{equation*}
\mathrm{E}[\mathfrak{p}(\theta)-h \mathfrak{t}(\theta)-c]=\mathrm{E}[V(\theta, \mathfrak{t}(\theta))-u(\theta, \theta)-h \mathfrak{t}(\theta)-c] \tag{C.4}
\end{equation*}
$$

where the expectation is with respect to the customer type $\theta$. In the following, we compute $\mathrm{E}[u(\theta, \theta)]$. By Lemma 3.2.1

$$
\begin{align*}
\mathrm{E}[u(\theta, \theta)] & =u(\underline{\Theta}, \underline{\Theta})+\int_{\theta=\underline{\Theta}}^{\bar{\theta}} d F(\theta) \int_{z=\underline{\Theta}}^{\theta} \partial_{1} V(z, \mathfrak{t}(z)) d z \\
& =u(\underline{\Theta}, \underline{\Theta})+\int_{z=\underline{\Theta}}^{\bar{\theta}} \int_{\theta=z}^{\bar{\theta}} d F(\theta) \partial_{1} V(z, \mathfrak{t}(z)) d z \\
& =u(\underline{\Theta}, \underline{\Theta})+\int_{z=\underline{\theta}}^{\bar{\theta}}(1-F(z)) \partial_{1} V(z, \mathfrak{t}(z)) d z \\
& =u(\underline{\Theta}, \underline{\Theta})+\mathrm{E}\left[\frac{(1-F(\theta))}{f(\theta)} \partial_{1} V(\theta, \mathfrak{t}(\theta))\right] \tag{C.5}
\end{align*}
$$

By replacing Eq. (C.5) in Eq. (C.4), we get the desired result.

## C.1.2 Lower Bound on the Revenue Gain of the Dynamic Pricing Policy

Here, we compare the revenue of the optimal mechanism given in Theorem 3.3.2 with that of the optimal FP policy. Recall that in the FP policy, the firm only sells to customers with type $\theta \geq \theta_{L}$ at time zero by posting a fixed price of $\theta_{L}$.

Lemma C.1.1 (Lower Bound on the Revenue Gain of DP). Let $R_{f}$ and $R_{o p t}$ be the expected revenue of the firm under the FP policy and the optimal DP policy, respectively. Then,

$$
\frac{R_{o p t}-R_{f}}{R_{f}} \geq e^{-1} \frac{\mathrm{E}\left[\theta \mathbb{1}\left\{\theta \leq \theta_{L}\right\}\right]}{\theta_{L}\left(1-F\left(\theta_{L}\right)\right)} .
$$

Proof of Lemma C.1.1 is given at the end of this section.
Assume that the customer type $\theta$ is drawn from the Uniform distribution in the range of $[0,1]$; that is, $\theta \sim U(0,1)$. Then, $\theta_{L}=0.5$, and $R_{f}=\frac{1}{4}$. Lemma C.1.1 implies that the firm can increase his revenue by more than $100 \frac{e^{-1} \mathrm{E}\left[\theta 11\left\{\theta \leq \theta_{L}\right\}\right]}{\theta_{L}\left(1-F\left(\theta_{L}\right)\right)} \approx 18 \%$ by using DP.

Proof. Proof of Lemma C.1.1 By Lemma 3.2.2, under the FP policy,

$$
\begin{align*}
R_{f} & =\mathrm{E}\left[V(\theta, \mathfrak{t}(\theta))+\alpha(\theta) \partial_{1} V(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta}, \underline{\Theta})\right] \\
& =\mathrm{E}\left[e^{-\beta \theta t(\theta)}(\theta+\alpha(\theta)(1-\beta \theta \mathfrak{t}(\theta)))\right]=\mathrm{E}\left[(\theta+\alpha(\theta)) \mathbb{1}\left\{\theta \geq \theta_{L}\right\}\right], \tag{C.6}
\end{align*}
$$

where the last inequality holds because in the FP policy, $\mathfrak{t}(\theta)=0$ for $\theta \geq \theta_{L}$.
Similarly, under the mechanism described in Theorem 3.3.2, we have

$$
\begin{align*}
R_{o p t}= & \mathbf{E}\left[e^{-\beta \theta \mathfrak{t}_{R}(\theta)}\left(\theta+\alpha(\theta)\left(1-\beta \theta \mathfrak{t}_{R}(\theta)\right)\right)\right] \\
= & \mathrm{E}\left[(\theta+\alpha(\theta)) \times \mathbb{1}\left\{\theta \geq \theta_{H}\right\}\right.  \tag{C.7}\\
& +e^{-\beta \theta \mathfrak{t}_{R}(\theta)}\left(\theta+\alpha(\theta)\left(1-\beta \theta \mathfrak{t}_{R}(\theta)\right)\right) \times \mathbb{1}\left\{\theta \in\left(\theta_{L}, \theta_{H}\right)\right\} \\
& \left.+e^{-1} \theta \times \mathbb{1}\left\{\theta \leq \theta_{L}\right\}\right],
\end{align*}
$$

where the first equation holds because of the time of purchase in the optimal DP policy is $\mathfrak{t}_{R}(\cdot)$, given in Eq. (3.6). Then, considering the fact that $\theta \in\left(\theta_{L}, \theta_{H}\right), \mathfrak{t}_{R}(\theta)$ is the FOC solution, i.e., $\mathfrak{t}_{R}(\theta)=$ $\arg \max _{t \geq 0}\{R(\theta, t)\}$, we get

$$
R_{\text {opt }} \geq \mathrm{E}\left[(\theta+\alpha(\theta)) \times \mathbb{1}\left\{\theta \geq \theta_{L}\right\}+e^{-1} \theta \times \mathbb{1}\left\{\theta \leq \theta_{L}\right\}\right] .
$$

By the above equation and Eq. (C.6), we get $R_{\text {opt }}-R_{f} \geq e^{-1} \mathrm{E}\left[\theta \times \mathbb{1}\left\{\theta \leq \theta_{L}\right\}\right]$. Then the result follows because $R_{f}=\theta_{L}\left(1-F\left(\theta_{L}\right)\right)$.

## C.1.3 Proof of Theorem 3.4.4

We have divided the proof into a sequence of lemmas. Lemma C.1.2 characterizes the optimal mechanism when the holding cost is low. Lemmas C.1.3 and C.1.4, respectively, characterize the optimal mechanism when the holding cost is medium and high.

Lemma C.1.2 (Low Holding Cost). If Assumption 3 holds, the valuation function $V(\theta, t)=\theta e^{-\beta \theta t}$, and the holding cost $h \leq H_{l}$, then the optimal mechanism sells to the customer of type $\theta \geq \max \left\{\underline{\theta}_{L}, \underline{\theta}\right\}$ at time $\mathfrak{t}_{h}(\theta)$ and at price $\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{h}(\theta)\right)-\int_{\max \left\{\underline{\theta}_{L}, \underline{\theta}\right\}}^{\theta} e^{-\beta \mathfrak{t}_{h}(z) z}\left(1-\beta \mathfrak{t}_{h}(z) z\right) d z$ where $H_{l}, \mathfrak{t}_{h}(\cdot)$, and $\underline{\theta}_{L}$ are defined in Equations (3.7) and (3.8). For $\theta \leq \max \left\{\underline{\theta}_{L}, \underline{\theta}\right\}, \mathfrak{p}(\theta)=\infty$.

The proof of Lemma C.1.2 is given in Section C.1.4.

Lemma C.1.3 (Medium Holding Cost). If Assumption 3 holds, the valuation function $V(\theta, t)=\theta e^{-\beta \theta t}$, the holding cost $h \in\left[H_{l}, H_{h}\right]$, and $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$ has a unique solution, then the optimal mechanism
sells to the customer of type $\theta \geq \underline{\theta}_{M}$ at time $\mathfrak{t}_{h}(\theta)$ and at price $\mathfrak{p}(\theta)=V\left(\theta, \mathfrak{t}_{h}(\theta)\right)-\int_{\underline{\theta}_{M}}^{\theta} e^{-\beta \mathfrak{t}_{h}(z) z}(1-$ $\left.\beta \mathfrak{t}_{h}(z) z\right) d z$ where $H_{l}, H_{h}, \underline{\theta}_{M}, \mathfrak{t}_{h}(\cdot)$, and the FOC solution $\mathfrak{t}_{f}(\cdot)$ are defined in Equations (3.7) and (3.8). For $\theta<\underline{\theta}_{M}, \mathfrak{p}(\theta)=\infty$.

The assumption in Lemma C.1.3 is discussed in Section C.1.8, and the proof of Lemma C.1.3 is provided in Section C.1.5.

Lemma C.1.4 (High Holding Cost). If Assumption 3 holds, the valuation function $V(\theta, t)=\theta e^{-\beta \theta t}$, the holding cost $h \geq H_{h}$, and $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{h} \mathfrak{t}_{f}(\theta)=0$ has a unique solution, then the optimal mechanism sells to customers with type $\theta \geq \theta_{L}$ at time zero and at price $\mathfrak{p}(\theta)=\theta_{L}$ where $\theta_{L}$ solves $\theta_{L}+\alpha\left(\theta_{L}\right)=0$ and $H_{l}, H_{h}$, and the FOC solution $\mathfrak{t}_{f}(\cdot)$ are defined in Eq. (3.7). For customers with type $\theta<\theta_{L}, \mathfrak{p}(\theta)=\infty$.

The proof is given in Section C.1.6.

## C.1.4 Optimal Mechanism for a Low Holding Cost

In this section, we present the proof of Lemma C.1.2. Throughout the proof, for convenience, we assume that $\underline{\theta}_{L} \geq \underline{\Theta}$. We need to show that the time of allocation in the optimal mechanism is given by

$$
\mathfrak{t}^{*}(\theta):=\left\{\begin{array}{ll}
\mathfrak{t}_{h}(\theta) & \text { if } \theta \geq \max \left\{\underline{\theta}_{L}, \underline{\Theta}\right\} ;  \tag{C.8}\\
\infty & \text { if } \theta<\max \left\{\underline{\theta}_{L}, \underline{\Theta}\right\}
\end{array}= \begin{cases}0 & \text { if } \theta \geq \theta_{H}^{h} ; \\
\mathfrak{t}_{f}(\theta) & \text { if } \theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right] ; \\
\frac{1}{\beta \theta} & \text { if } \theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right] ; \\
\infty & \text { if } \theta<\underline{\theta}_{L} .\end{cases}\right.
$$

Note that $\mathfrak{t}^{*}(\theta)=\infty$ implies that the mechanism does not allocate the item to customers with type $\theta$.
To characterize the optimal mechanism, by Lemma 3.2.2, we need to solve the optimization Problem OPT-H. That is, we need to maximize the expected virtual revenue subject to IR and IC constrains. Lemma 3.2.1 shows that a mechanism is IC iff the interval and envelope conditions hold. This implies that one can replace the IC constraints with these two conditions. However, as we noted in Section 3.3, characterizing the optimal mechanism that satisfies both these two conditions in rather complicated. Thus, in the following, we relax Problem OPT-H and only consider the IR and envelope conditions. We then show that the solution of the relaxed problem also satisfies the interval condition. Thus, it is optimal.

The relaxed problem can be formulated as follows.

$$
\begin{align*}
&\{\mathfrak{t}(\theta): \theta \in[\underline{\Theta}, \bar{\Theta}]\}, u(\underline{\underline{( }}, \underline{\Theta}) \geq 0 \\
& \mathrm{E}[\max \{R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta), 0\}]  \tag{IR}\\
& \text { s.t. } u(\theta, \theta)=u(\underline{\Theta}, \underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta \mathfrak{t}(z) z}(1-\beta \mathfrak{t}(z) z) d z \geq 0 \text { for } \theta \in[\underline{\Theta}, \bar{\Theta}]
\end{align*}
$$

(OPT-H-R)
where the maximization is taken over the time of sales $\mathfrak{t}(\theta)$ and the utility of a customer with type $\underline{\Theta}$, i.e., $u(\underline{\Theta}, \underline{\Theta})$. Here, $R(\theta, t)$ is the virtual value of customer of type $\theta$ at time $t$, and is defined in Eq. (B.14).

The following lemma characterizes the optimal solution of the relaxed problem. This lemma is our main technical contribution of this section.

Lemma C.1.5. If Assumption 3 holds and the holding cost $h \leq H_{l}$, in an optimal solution of Problem OPT-H-R, time of sales is $\mathfrak{t}^{*}(\theta)$ and $u(\underline{\Theta}, \underline{\Theta})=0$ where $\mathfrak{t}^{*}(\cdot)$ is defined in Eq. (C.8).

The proof is provided in Section C.1.4. In the proof, we first show that $\mathfrak{t}^{*}(\cdot)$ is a feasible solution of the relaxed problem. Then, we show that it is optimal.

To verify that $\mathfrak{t}^{*}(\cdot)$ is an optimal solution of Problem OPT-H, we show that the interval condition specified in Lemma 3.2.1 is fulfilled. That is, for any $\hat{\theta}, \theta \in[\underline{\Theta}, \bar{\Theta}]$ such that $\hat{\theta} \leq \theta$,

$$
\begin{aligned}
\int_{\hat{\theta}}^{\theta} A\left(z, \mathfrak{t}^{*}(\hat{\theta})\right) d z & \leq \int_{\hat{\theta}}^{\theta} A\left(z, \mathfrak{t}^{*}(z)\right) d z, \\
\int_{\hat{\theta}}^{\theta} A\left(z, \mathfrak{t}^{*}(z)\right) d z & \leq \int_{\hat{\theta}}^{\theta} A\left(z, \mathfrak{t}^{*}(\theta)\right) d,
\end{aligned}
$$

where $A(z, t)=\partial_{1} V(z, t)=e^{-\beta t z}(1-\beta t z)$ and $\int_{\hat{\theta}}^{\theta} A\left(z, \mathrm{t}^{*}(z)\right) d z=u(\theta, \theta)-u(\hat{\theta}, \hat{\theta})$. To this aim, we show that for any $z \geq \hat{\theta}, A\left(z, \mathfrak{t}^{*}(\hat{\theta})\right) \leq A\left(z, \mathfrak{t}^{*}(z)\right)$ and for any $z \leq \theta, A\left(z, \mathfrak{t}^{*}(z)\right) \leq A\left(z, \mathfrak{t}^{*}(\theta)\right)$.

We will make use of the following preliminary result.

Lemma C.1.6. The FOC solution $\mathfrak{t}_{f}(\theta)$, defined in Eq. (3.7), is a decreasing function of $\theta$ as long as $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta) \geq 0$. In addition, for any $\theta \in\left[\underline{\theta}_{L}, \bar{\Theta}\right], 0 \leq A\left(\theta, \mathfrak{t}_{h}(\theta)\right) \leq 1$ where $A(z, t)=$ $e^{-\beta t z}(1-\beta t z)$.

Unless stated otherwise, the proof of all technical lemmas is given in Section C.2.

Later in Lemma C.1.12, we will show that $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta) \geq 0$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$. This result and the result in Lemma C.1.6 imply that $\mathfrak{t}_{h}(\theta)=\mathfrak{t}_{f}(\theta)$ is decreasing for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$. Then, considering the fact that $\mathfrak{t}_{h}(\theta)=0$ for $\theta \geq \theta_{H}^{h}, \mathfrak{t}_{h}(\theta)=\frac{1}{\beta \theta}$ for $\theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right], \mathfrak{t}_{h}\left(\theta_{H}^{h}\right)=0$, and $\mathfrak{t}_{h}\left(\theta_{L}^{h}\right)=\frac{1}{\beta \theta_{L}^{h}}$, we can conclude that $\mathfrak{t}_{h}(\theta)$ is decreasing in $\theta \geq \underline{\theta}_{L}$.

Now, we are ready to show that the interval conditions are satisfied.
We first note that when $\hat{\theta} \leq \underline{\theta}_{L}$, it easy to show that for any $z \geq \hat{\theta}, A\left(z, \mathrm{t}^{*}(\hat{\theta})\right) \leq A\left(z, \mathrm{t}^{*}(z)\right)$. This holds because $A\left(z, \mathfrak{t}^{*}(\hat{\theta})\right)=0$ and as shown in Lemma C.1.6, $A\left(z, \mathfrak{t}^{*}(z)\right) \geq 0$. Also, when $\theta \leq \underline{\theta}_{L}$, we have $A\left(z, \mathfrak{t}^{*}(z)\right) \leq A\left(z, \mathfrak{t}^{*}(\theta)\right)$ for any $z \leq \theta$. This follows from the fact that both $A\left(z, \mathfrak{t}^{*}(z)\right)$ and $A\left(z, \mathfrak{t}^{*}(\theta)\right)$ are zero.

Next, we assume that both $\theta$ and $\hat{\theta}$ are greater than $\underline{\theta}_{L}$. Recall that for $\theta \geq \underline{\theta}_{L}, \mathfrak{t}^{*}(\theta)=\mathfrak{t}_{h}(\theta)$. We start with showing $A\left(z, \mathfrak{t}_{h}(\hat{\theta})\right) \leq A\left(z, \mathfrak{t}_{h}(z)\right), z \geq \hat{\theta}$. We consider two cases: $1-\left(1-\beta \mathfrak{t}_{h}(\hat{\theta}) z\right) \leq 0$ and 2-$\left(1-\beta \mathfrak{t}_{h}(\hat{\theta}) z\right)>0$. Assume that $\left(1-\beta \mathfrak{t}_{h}(\hat{\theta}) z\right) \leq 0$. Then, we have

$$
e^{-\beta z \mathfrak{t}_{h}(\hat{\theta})}\left(1-\beta \mathfrak{t}_{h}(\hat{\theta}) z\right) \leq 0 \leq e^{-\beta z \mathfrak{t}_{h}(z)}\left(1-\beta \mathfrak{t}_{h}(z) z\right)
$$

where the second inequality follows from Lemma C. 1.6 where we show that $A\left(z, \mathfrak{t}_{h}(z)\right)=e^{-\beta z \mathfrak{t}_{h}(z)}(1-$ $\left.\beta \mathfrak{t}_{h}(z) z\right) \geq 0$. By the above equation, we get $A\left(z, \mathfrak{t}_{h}(\hat{\theta})\right) \leq A\left(z, \mathfrak{t}_{h}(z)\right)$.

Now, assume that $\left(1-\beta \mathfrak{t}_{h}(\hat{\theta}) z\right)>0$. Then, considering the fact that $\mathfrak{t}_{h}(\cdot)$ is decreasing, for any $z \geq \hat{\theta}$, we have $\left(1-\beta \mathfrak{t}_{h}(z) z\right) \geq\left(1-\beta \mathfrak{t}_{h}(\hat{\theta}) z\right)$, and $e^{-\beta \mathfrak{t}_{h}(z) z} \geq e^{-\beta \mathfrak{t}_{h}(\hat{\theta}) z}$. By multiplying these two Equations, we get $A\left(z, \mathfrak{t}_{h}(\hat{\theta})\right) \leq A\left(z, \mathfrak{t}_{h}(z)\right)$.

Next, we will verify that $A\left(z, \mathfrak{t}_{h}(z)\right) \leq A\left(z, \mathfrak{t}_{h}(\theta)\right)$. Given that $\mathfrak{t}_{h}(\cdot)$ is decreasing, for any $z \geq \theta$, we have

$$
0 \leq\left(1-\beta \mathfrak{t}_{h}(z) z\right) \leq\left(1-\beta \mathfrak{t}_{h}(\theta) z\right), \text { and } e^{-\beta \mathfrak{t}_{h}(z) z} \leq e^{-\beta \mathfrak{t}_{h}(\theta) z}
$$

where the first inequality holds follows from Lemma C.1.2 where we show $A\left(z, \mathfrak{t}_{h}(z)\right)=e^{-\beta z \mathfrak{t}_{h}(z)}(1-$ $\left.\beta \mathfrak{t}_{h}(z) z\right) \geq 0$. By multiplying these two Equations, we have $A\left(z, \mathfrak{t}_{h}(z)\right) \leq A\left(z, \mathfrak{t}_{h}(\theta)\right)$.

## Proof of Lemma C.1.5

Here, with some abuse of notations, we denote $\mathfrak{t}^{*}(\cdot)$ with $\mathfrak{t}_{h}(\cdot)$. Recall that $\mathfrak{t}^{*}(\theta)=\mathfrak{t}_{h}(\theta)$ when $\theta \geq \underline{\theta}_{L}$ and is $\infty$ otherwise. Also, for simplicity, we denote $u(\theta, \theta)$ by $u(\theta)$.

The proof has two parts. In the first part, we show that the solution given in Lemma C.1.5 is a feasible solution of Problem OPT-H-R. In the second part, we verify that this solution is an optimal solution of this problem.

Feasibility: To show that $\mathfrak{t}_{h}(\cdot)$ is a feasible solution of Problem OPT-H-R, we will verify that $u(\theta) \geq 0$ for any $\theta \in[\underline{\Theta}, \bar{\Theta}]$. For any $\theta \leq \theta_{L}^{h}$, it is easy to verify that $u(\theta)=u(\underline{\Theta})=0$. Thus, we only need to show that $u(\theta) \geq 0$ for any $\theta \geq \theta_{L}^{h}$. To prove that $u(\theta) \geq 0$ for $\theta \geq \theta_{L}^{h}$, we make use of Lemma C.1. 6 where we show that $e^{-\beta \mathfrak{t}_{h}(\theta) \theta}\left(1-\beta \mathfrak{t}_{h}(\theta) \theta\right) \geq 0$. This implies that $u(\theta)=\int_{\underline{\Theta}}^{\theta} e^{-\beta \mathfrak{t}_{h}(z) z}\left(1-\beta \mathfrak{t}_{h}(z) z\right) d z \geq 0$

Optimality: Here, we will show that the solution given in Lemma C.1.5, is an optimal solution of Problem OPT-H-R. To this end, we find an upper bound for the optimal value of Problem OPT-H-R by dualizing the IR constraints. Then, we will show that the solution given in Lemma C.1.5 achieves the upper bound and thus is optimal.

Upper Bound of OPT-H-R: For any time of sales $\mathfrak{t}(\cdot)$ and Lagrangian function $\lambda:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}^{+}$, we define the following function.

$$
L_{h}(\mathfrak{t}(\cdot), \lambda(\cdot), u(\underline{\Theta}))=\mathrm{E}[R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta)-u(\underline{\Theta})]+\int_{\underline{\Theta}}^{\bar{\Theta}} \lambda(z) u(z) d z
$$

where $u(z)=\int_{\underline{\Theta}}^{z} e^{-\beta \mathfrak{t}(\theta) \theta}(1-\beta \mathfrak{t}(\theta) \theta) d \theta+u(\underline{\Theta})$, and $R$ is defined in Eq. (B.14). Then, considering the fact that $\lambda(\cdot) \geq 0$, for any $(\mathfrak{t}(\cdot), u(\underline{\Theta}))$ such that $u(\theta)=u(\underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta z \mathfrak{t}(z)}(1-\beta z \mathfrak{t}(z)) \geq 0$, we have

$$
\mathrm{E}[R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta)-u(\underline{\Theta})] \leq L_{h}(\mathfrak{t}(\cdot), \lambda(\cdot), u(\underline{\Theta}))
$$

One can think of $\lambda(\theta)$ as a dual variable for the IR constraints. Therefore, for any $\lambda:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}^{+}$,

$$
\begin{equation*}
\max _{(\mathfrak{t}(\cdot), u(\underline{\Theta})) \in \mathcal{T}}\{\mathrm{E}[R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta)-u(\underline{\Theta})]\} \leq \max _{(\mathfrak{t}(\cdot), u(\underline{\Theta})) \in \mathcal{T}}\left\{L_{h}(\mathfrak{t}(\cdot), \lambda(\cdot), u(\underline{\Theta}))\right\} \tag{C.9}
\end{equation*}
$$

where $\mathcal{T}=\left\{(\mathfrak{t}(\cdot), u(\underline{\Theta})): \mathfrak{t}(\theta) \geq 0, u(\underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\beta z \mathfrak{t}(z)}(1-\beta z \mathfrak{t}(z)) \geq 0\right.$ for any $\left.\theta \in[\underline{\Theta}, \overline{\bar{\Theta}}]\right\}$ is the set of feasible solutions. In the following, we will characterize an upper bound for $\max _{(t \mathfrak{t} \cdot \cdot), u(\underline{\Theta})) \in \mathcal{T}}\{\mathrm{E}[R(\theta, \mathfrak{t}(\theta))-$ $h \mathfrak{t}(\theta)-u(\underline{\Theta})]\}$ by considering a specific Lagrangian function, defined below.

$$
\lambda_{h}(\theta)= \begin{cases}0 & \text { if } \theta>\theta_{L}^{h}  \tag{C.10}\\ \left(f(\theta)\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right)\right)^{\prime} & \text { if } \theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right] ; \\ (f(\theta)(2 \theta+\alpha(\theta)))^{\prime} & \text { if } \theta \in\left[\underline{\Theta}, \underline{\theta}_{L}\right] ;\end{cases}
$$

where $(f(\theta)(2 \theta+\alpha(\theta)))^{\prime}$ and $\left(f(\theta)\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right)\right)^{\prime}$ are respectively the derivative of $(f(\theta)(2 \theta+\alpha(\theta)))$ and $\left(f(\theta)\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right)\right)$ with respect to $\theta$. The following lemma establishes that $\lambda_{h}(\theta) \geq 0$.

Lemma C.1.7. When $h \leq H_{l}$, for any $\theta \in[\underline{\Theta}, \bar{\Theta}], \lambda_{h}(\theta)$, defined in Eq. (C.10), is non-negative.

The following claim shows that $\left(\mathfrak{t}_{h}(\cdot), u(\underline{\Theta})=0\right)$ is an optimal solution of Problem OPT-H-R.
Claim: With a slight abuse of notations, let

$$
\left(t_{\lambda}(\cdot), u_{\lambda}\right)=\arg \max _{(\mathfrak{t}(\cdot), u(\underline{\Theta})) \in \mathcal{T}}\left\{L_{h}\left(\mathfrak{t}(\cdot), \lambda_{h}(\cdot), u(\underline{\Theta})\right)\right\} .
$$

Then, $t_{\lambda}(\cdot)=\mathfrak{t}_{h}(\theta)$ for any $\theta \in[\underline{\Theta}, \bar{\Theta}]$ and $u_{\lambda}=0$. Furthermore, $L_{h}\left(\mathfrak{t}_{h}(\cdot), \lambda_{h}(\cdot), u_{\lambda}\right)=\mathrm{E}\left[R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-\right.$ $\left.h \mathfrak{t}_{h}(\theta)-u_{\lambda}\right]$.

Proof of the Claim: By definition, $\lambda_{h}(\theta)=0$ for $\theta>\theta_{L}^{h}$. Thus, we get

$$
\begin{align*}
L_{h}\left(\mathfrak{t}(\cdot), \lambda_{h}(\cdot), u(\underline{\Theta})\right) & =\int_{z=\theta_{L}^{h}}^{\bar{\Theta}}(R(z, \mathfrak{t}(z))-h \mathfrak{t}(z)) f(z) d z \\
& +\int_{z=\underline{\Theta}}^{\theta_{L}^{h}}\left((R(z, \mathfrak{t}(z))-h \mathfrak{t}(z)) f(z)+\lambda_{h}(z) u(z)\right) d z-u(\underline{\Theta}) . \tag{C.11}
\end{align*}
$$

From definition of $\lambda_{h}(z)$, the last two terms of Eq. (C.11) can be written as

$$
\begin{align*}
& \int_{\underline{\theta}_{L}}^{\theta_{L}^{h}}(R(z, \mathfrak{t}(z))-h \mathfrak{t}(z)) f(z) d z+\int_{\underline{\theta}_{L}}^{\theta_{L}^{h}} u(z) d\left(f(z)\left(z+\alpha(z)+\frac{h}{\beta z e^{-1}}\right)\right) \\
& \int_{\underline{\Theta}}^{\underline{\theta}_{L}}\left((R(z, \mathfrak{t}(z))-h \mathfrak{t}(z)) f(z) d z+\int_{\underline{\Theta}}^{\underline{\theta}_{L}} u(z) d(f(z)(2 z+\alpha(z)))-u(\underline{\Theta}) .\right. \tag{C.12}
\end{align*}
$$

We first focus on the first two terms where $z \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right]$. By integrating by part and using the definition of $R$, the first two terms can be rewritten as

$$
\begin{aligned}
& \int_{\underline{\theta}_{L}}^{\theta_{L}^{h}}\left(e^{-\beta \mathfrak{t}(z) z}(z+\alpha(z)(1-\beta \mathfrak{t}(z) z)-h \mathfrak{t}(z)) f(z) d z\right. \\
& +\left.u(z) f(z)\left(z+\alpha(z)+\frac{h}{\beta z e^{-1}}\right)\right|_{\underline{\theta}_{L}} ^{\theta_{L}^{h}}-\int_{\underline{\theta}_{L}}^{\theta_{L}^{h}} e^{-\beta \mathfrak{t}(z) z}(1-\beta \mathfrak{t}(z) z) f(z)\left(z+\alpha(z)+\frac{h}{\beta z e^{-1}}\right) d z .
\end{aligned}
$$

In the above equation, we use the fact that $\frac{d u(z)}{d z}=e^{-\beta t(z) z}(1-\beta \mathfrak{t}(z) z)$. Then, by definition of $\theta_{L}^{h}$, i.e., the fact that $\left(\theta_{L}^{h}+\alpha\left(\theta_{L}^{h}\right)+\frac{h}{\beta \theta_{L}^{h} e^{-1}}\right)=0$, the above equation is simplified as

$$
\begin{align*}
& -u\left(\underline{\theta}_{L}\right) f\left(\underline{\theta}_{L}\right)\left(\underline{\theta}_{L}+\alpha\left(\underline{\theta}_{L}\right)+\frac{h}{\beta \underline{\theta}_{L} e^{-1}}\right) \\
& +\int_{\underline{\theta}_{L}}^{\theta_{L}^{h}} f(z)\left(e^{-\beta \mathfrak{t}(z) z} z^{2} \beta \mathfrak{t}(z)-h \mathfrak{t}(z)-e^{-\beta \mathfrak{t}(z) z}(1-\beta \mathfrak{t}(z) z) \frac{h}{\beta z e^{-1}}\right) d z \tag{C.13}
\end{align*}
$$

Now, we focus on the last three terms of Eq. (C.12). Again, by integrating by part and using definition of $R$, the last two terms of Eq. (C.12) can be rewritten as

$$
\begin{align*}
& \int_{\underline{\Theta}}^{\underline{\theta}_{L}} f(z)\left(e^{-\beta \mathfrak{t}(z) z}(z+\alpha(z)(1-\beta \mathfrak{t}(z) z))-h \mathfrak{t}(z)\right) d z \\
& +\left.u(z) f(z)(2 z+\alpha(z))\right|_{\underline{\Theta}} ^{\theta_{L}}-\int_{\underline{\Theta}}^{\underline{\theta}_{L}} e^{-\beta \mathfrak{t}(z) z}(1-\beta \mathfrak{t}(z) z) f(z)(2 z+\alpha(z)) d z-u(\underline{9} \\
& =u\left(\underline{\theta}_{L}\right) f\left(\underline{\theta}_{L}\right)\left(2 \underline{\theta}_{L}+\alpha\left(\underline{\theta}_{L}\right)\right)+u(\underline{\Theta})(-1-f(\underline{\Theta})(2 \underline{\Theta}+\alpha(\underline{\Theta}))) \\
& +\int_{\underline{\Theta}}^{\underline{\theta}_{L}} f(z)\left(z e^{-\beta \mathfrak{t}(z) z}(-1+2 \beta \mathfrak{t}(z) z)-h \mathfrak{t}(z)\right) d z . \tag{C.14}
\end{align*}
$$

Note that the coefficient of $u(\underline{\Theta})$, i.e., $(-1-f(\underline{\Theta})(2 \underline{\Theta}+\alpha(\underline{\Theta})))$, can be simplified as $-2 \underline{\Theta} f(\underline{\Theta}) \leq 0$. By plugging Equations (C.13) and (C.14) into Eq. (C.11), and by using definition of $\underline{\theta}_{L}$, we get

$$
\begin{aligned}
L_{h}\left(\mathfrak{t}(\cdot), \lambda_{h}(\cdot), u(\underline{\Theta})\right) & =\int_{\theta_{L}^{h}}^{\bar{\Theta}}(R(z, \mathfrak{t}(z))-h \mathfrak{t}(z)) f(z) d z \\
& +\int_{\underline{\theta}_{L}}^{\theta_{L}^{h}} f(z)\left(e^{-\beta \mathfrak{t}(z) z} z^{2} \beta \mathfrak{t}(z)-h \mathfrak{t}(z)-e^{-\beta \mathfrak{t}(z) z}(1-\beta \mathfrak{t}(z) z) \frac{h}{\beta z e^{-1}}\right) d z \\
& +\int_{\underline{\Theta}}^{\underline{\theta}_{L}} f(z)\left(z e^{-\beta \mathfrak{t}(z) z}(-1+2 \beta \mathfrak{t}(z) z)-h \mathfrak{t}(z)\right) d z-2 \underline{\Theta} f(\underline{\Theta}) u(\underline{\Theta}) .
\end{aligned}
$$

First of all, since the coefficient of $u(\underline{\Theta})$ is negative, to maximize the above equation, we need to set $u(\underline{\Theta})$ to zero. That is, $u_{\lambda}=0$. Then, $\max _{(\mathfrak{t}(\cdot), u(\underline{\Theta})) \in \mathcal{T}}\left\{L_{h}\left(\mathfrak{t}(\cdot), \lambda_{h}(\cdot), u(\underline{\Theta})\right)\right\}$ can be upper-bounded as follows

$$
\begin{align*}
& \max _{(\mathfrak{t}(\cdot), u(\underline{\Theta})) \in \mathcal{T}}\left\{L_{h}\left(\mathfrak{t}(\cdot), \lambda_{h}(\cdot), u(\underline{\Theta})=0\right)\right\} \leq \int_{\theta_{L}^{h}}^{\bar{\Theta}} f(z) \max _{t \geq 0}\{R(z, t)-h t\} d z \\
& +\int_{\underline{\theta}_{L}}^{\theta_{L}^{h}} f(z) \max _{t \geq 0}\left\{\left(e^{-\beta t z} z^{2} \beta t-h t-e^{-\beta t z}(1-\beta t z) \frac{h}{\beta z e^{-1}}\right)\right\} d z \\
& +\int_{\underline{\Theta}}^{\underline{\theta}_{L}} f(z) \max \left\{\max _{t \geq 0}\left\{\left(z e^{-\beta t z}(-1+2 \beta t z)-h t\right)\right\}, 0\right\} d z . \tag{C.15}
\end{align*}
$$

We take advantage of the following lemma to simplify the first term of the above equation.
Lemma C.1.8. If Assumption 3 holds and the holding cost $h \leq H_{l}$, then for any $z \geq \theta_{L}^{h}$, we have $\arg \max _{t \geq 0}\{R(z, t)-h t\}=\mathfrak{t}_{h}(z)$.

Note that the optimal solution characterized in Lemma C.1.8 is the maximum of the FOC solution and zero. We now simplify the second term of Eq. (C.15). It is easy to verify that for any $z \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right]$, we have

$$
\begin{equation*}
\arg \max _{t \geq 0}\left\{\left(e^{-\beta t z} z^{2} \beta t-h t-e^{-\beta t z}(1-\beta t z) \frac{h}{\beta z e^{-1}}\right)\right\}=\frac{1}{\beta z}=\mathfrak{t}_{h}(z) \tag{C.16}
\end{equation*}
$$

Finally, the following lemma characterizes an optimal solution of the third term of Eq. (C.15).
Lemma C.1.9. If Assumption 3 holds and the holding cost $h \leq H_{l}$, for any $z \leq \underline{\theta}_{L}$, we have

$$
\max _{t \geq 0}\left\{\left(z e^{-\beta t z}(-1+2 \beta t z)-h t\right)\right\} \leq 0
$$

Lemmas C.1.8, C.1.9, and Eq. (C.16) show that $t_{\lambda}(\theta)=\mathfrak{t}_{h}(\theta)$ and $u_{\lambda}=0$. Then, the proof is completed by observing that $L_{h}\left(\mathfrak{t}_{h}(\cdot), \lambda_{h}(\cdot), 0\right)=\mathrm{E}\left[R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-h \mathfrak{t}_{h}(\theta)\right]$.

## C.1.5 Optimal Mechanism for a Medium Holding Cost

Here, we present the proof for Lemma C.1.3. We show that in the optimal solution of Problem OPT-H, the time of purchase is given by

$$
\mathfrak{t}^{*}(\theta)=\left\{\begin{array}{ll}
\mathfrak{t}_{h}(\theta) & \text { if } \theta \geq \underline{\theta}_{M} ;  \tag{C.17}\\
\infty & \text { if } \theta<\underline{\theta}_{M}
\end{array}= \begin{cases}0 & \text { if } \theta \geq \theta_{H}^{h} ; \\
\mathfrak{t}_{f}(\theta) & \text { if } \theta \in\left[\underline{\theta}_{M}, \theta_{H}^{h}\right] ; \\
\infty & \text { if } \theta<\underline{\theta}_{M} ;\end{cases}\right.
$$

and $u(\theta, \theta)=\int_{\underline{\theta}_{M}}^{\theta} e^{-\beta z \mathfrak{t}_{h}(z)}\left(1-\beta z \mathfrak{t}_{h}(z)\right) d z$. Here, $\mathfrak{t}^{*}(\theta)=\infty$ implies that customer with type $\theta$ does not purchase the item.

The proof has three main steps. In the first step, we relax the problem by ignoring both IC and IR constraints and we find an allocation rule that maximizes the virtual revenue. Then, we show that the solution of this relaxed problem can construct a mechanism that satisfy the IR and envelope conditions. Finally, we show that the aforementioned solution also satisfies the interval conditions, as a result, it is optimal.

- Maximizing virtual revenue without IC and IR constraints: Consider that following optimization problem.

$$
\begin{equation*}
\max _{\{\mathfrak{t}(\theta) \geq 0: \theta \in[\underline{\Theta}, \bar{\Theta}]\}} \mathrm{E}[\max \{R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta), 0\}] \tag{OPT-H-1}
\end{equation*}
$$

The following lemma shows that $\mathfrak{t}^{*}(\cdot)$, given in Eq. (C.17), is an optimal solution of Problem OPT-H-1.

Lemma C.1.10. The optimal solution of Problem OPT-H-1 is given by $\mathfrak{t}^{*}(\cdot)$ where $\mathfrak{t}^{*}(\cdot)$ is defined in Eq. (C.17).

The proof is similar to the proof of Lemma C.1.8; thus, it is omitted. The main idea of the proof is to show that $R(\theta, t)-h t$ as a function of $t$ has an inverted u-shape. Thus, it obtains its maximum at $\max \left\{0, \mathfrak{t}_{f}(\theta)\right\}$, where $\mathfrak{t}_{f}(\theta)$ is the FOC solution. Note that to show Lemma C.1.10, we need the
assumption that $\underline{\theta}_{M}$, i.e., the solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$, is unique. By this assumption, for any $\theta<\underline{\theta}_{M}$ we get

$$
\max _{t \geq 0}\{R(\theta, t)-h t\}=R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)<0 \quad \text { for } \theta<\underline{\theta}_{M} .
$$

This implies that it is optimal not to allocate the item to customers with type $\theta<\underline{\theta}_{M}$.

- Maximizing virtual revenue with IR and envelope constraints: Here, we show that the time of purchase $\mathfrak{t}^{*}(\cdot)$ is an optimal solution of Problem OPT-H-R. ${ }^{1}$ To this aim, we verify that

$$
u(\theta, \theta)=\int_{\underline{\theta}_{M}}^{\theta} e^{-\beta \mathfrak{t}_{h}(z) z}\left(1-\beta \mathfrak{t}_{h}(z) z\right) d z \geq 0
$$

Particularly, we show that for any $\theta \geq \underline{\theta}_{M},\left(1-\beta \theta \mathfrak{t}_{h}(\theta)\right) \geq 0$. Since $\mathfrak{t}_{h}(\theta)=0$ for $\theta \geq \theta_{H}^{h}$, it suffices to show that $\left(1-\beta \theta \mathfrak{t}_{h}(\theta)\right) \geq 0$ for any $\theta \in\left[\underline{\theta}_{M}, \theta_{H}^{h}\right]$.

Lemma C.1.11. For any $h \in\left[H_{l}, H_{h}\right]$ and $\theta \in\left[\underline{\theta}_{M}, \theta_{H}^{h}\right]$, we have $1-\beta \theta \mathfrak{t}_{h}(\theta) \geq 0$.

In the proof, we show that when $h \geq H_{l}$ and $\theta \geq \tilde{\theta}$, we have $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$. Then, we show that $\underline{\theta}_{M} \geq \tilde{\theta}$. This implies that $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$ for any $\theta \in\left[\underline{\theta}_{M}, \theta_{H}^{h}\right]$, which is the desired result.

- Maximizing virtual revenue with IR and IC constraints: Here, we need to show that the time of purchase $\mathfrak{t}^{*}(\cdot)$ and its associated payment, given in Lemma C.1.3, satisfy the interval conditions presented in Lemma 3.2.1. This part of the proof is very similar to that of Lemma C.1.2. Thus, we do not repeat it here.


## C.1.6 Optimal Mechanism for a High Holding Cost

In this section, we present the proof of Lemma C.1.4.
In the following, we show that $\max _{t \geq 0}\{R(\theta, t)-h t\}=R(\theta, 0)=\theta+\alpha(\theta) \geq 0$ for $\theta \geq \theta_{L}$, and for any $\theta<\theta_{L}, \max _{t \geq 0}\{R(\theta, t)-h t\}<0$ where $R$ is defined in Eq. (B.14). This implies that in the optimal mechanism, the firm only sells to customers with type $\theta \geq \theta_{L}$.

[^35]We first show that for any $\theta \geq \theta_{L}, \arg \max _{t \geq 0}\{R(\theta, t)-h t\}=0$. To this aim, we will verify that $\left.\frac{\partial(R(\theta, t)-h t)}{\partial t}\right|_{t=0} \leq 0$. This will give us the desired result because as we show in Lemma C.1.8, $R(\theta, t)-h t$ as a function of $t$ has an inverted $u$-shape. Therefore, if $\left.\frac{\partial(R(\theta, t)-h t)}{\partial t}\right|_{t=0} \leq 0$, we have $\arg \max _{t \geq 0}\{R(\theta, t)-$ $h t\}=0$.

By definition,

$$
\frac{\partial(R(\theta, t)-h t)}{\partial t}=-\beta \theta e^{-\beta t \theta}(\theta+\alpha(\theta)(2-\beta \theta t))-h
$$

and at $t=0$ and for any $\theta \geq \theta_{L}$, we have

$$
\begin{equation*}
\left.\frac{\partial(R(\theta, t)-h t)}{\partial t}\right|_{t=0}=-\beta \theta(\theta+2 \alpha(\theta))-h \leq 0 \tag{C.18}
\end{equation*}
$$

where the inequality holds because

$$
h \geq H_{h}=\beta \theta_{L}^{2}=-\beta \theta_{L}\left(\theta_{L}+2 \alpha\left(\theta_{L}\right)\right)=\max _{\theta \in\left[\theta_{L}, \bar{\theta}\right]}\{-\beta \theta(\theta+2 \alpha(\theta))\}
$$

The first equality follows because $\theta_{L}+\alpha\left(\theta_{L}\right)=0$ and last equality holds because $\arg \max _{\theta \in\left[\theta_{L}, \bar{\Theta}\right]}\{-\beta \theta(\theta+$ $2 \alpha(\theta))\}=\theta_{L}$. To see why the latter holds note that

$$
(-\theta(\theta+2 \alpha(\theta)))^{\prime}=-2(\theta+\alpha(\theta))-2 \theta \alpha^{\prime}(\theta) \leq 0
$$

where the inequality follows because for any $\theta \geq \theta_{L}$, we have $(\theta+\alpha(\theta)) \geq 0$.
Next, we will verify that for any $\theta<\theta_{L}, \max _{t \geq 0}\{R(\theta, t)-h t\}<0$. Note that it suffices to show that $\max _{t \geq 0}\left\{R(\theta, t)-H_{h} t\right\}<0$ considering the fact that $R(\theta, t)-h t$ is decreasing in $h$.

By Eq. (C.18), at $h=H_{h}$ we have $\mathfrak{t}_{f}\left(\theta_{L}\right)=0$, and more importantly

$$
R\left(\theta_{L}, \mathfrak{t}_{f}\left(\theta_{L}\right)\right)-H_{h} \mathfrak{t}_{f}\left(\theta_{L}\right)=\theta_{L}+\alpha\left(\theta_{L}\right)=0
$$

Then, considering the fact that $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{h} \mathfrak{t}_{f}(\theta)=0$ has unique solution, we have

$$
R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{h} \mathfrak{t}_{f}(\theta)=\max _{t \geq 0}\left\{R(\theta, t)-H_{h} t\right\}<0 \text { for any } \theta<\theta_{L}
$$

## C.1.7 Proof of Theorems 3.4.1 an 3.4.5

Here, we show the results in Theorem 3.4.5. Note that Theorem 3.4.1 is a special case of Theorem 3.4.5 with $h=0$.

The proof is similar to the proof of Theorem 3.4.4; thus we do not repeat it here. We only present a proof sketch. We need to find an optimal solution of Problem OPT-H where the objective function is replaced with $\mathrm{E}[\max \{R(\theta, \mathfrak{t}(\theta))-h \mathfrak{t}(\theta)-u(\underline{\Theta}, \underline{\Theta})-c, 0\}]$. We first relax the problem by ignoring the interval conditions. We will show that an optimal solution of the relaxed problem is $\mathfrak{t}_{h}(\theta)$ for $\theta \geq \theta_{c}$ and is $\infty$ for $\theta<\theta_{c}$, where $\theta_{c}$ solves $R\left(\theta_{c}, \mathfrak{t}_{h}\left(\theta_{c}\right)\right)-h \mathfrak{t}_{h}\left(\theta_{c}\right)=c$. This follows from Lemma C.1.12, stated at the end of this section, where we show that $R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-h \mathfrak{t}_{h}(\theta)$ is an increasing function of $\theta \geq \underline{\theta}_{h}$. Then, to complete the proof, we will show that the optimal solution of the relaxed problem satisfies the envelope conditions.

Lemma C.1.12. For any $h \geq 0, R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-h \mathfrak{t}_{h}(\theta)$ is increasing in $\theta \geq \underline{\theta}_{h}$. Furthermore, $R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-$ $h \mathfrak{t}_{h}(\theta) \geq 0$ for any $\theta \geq \underline{\theta}_{h}$.

Proof. Proof of Lemma C.1.12We show the result for $h \leq H_{l}$ where $H_{l}$ is defined in Eq. (3.7). A similar argument holds for $h>H_{l}$.

By definition, for any $h \leq H_{l}$, we have

$$
R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-h \mathfrak{t}_{h}(\theta)= \begin{cases}\theta+\alpha(\theta) & \text { if } \theta \geq \theta_{H}^{h} \\ R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta) & \text { if } \theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right] \\ e^{-1} \theta-\frac{h}{\beta \theta} & \text { if } \theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right]\end{cases}
$$

$R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-h \mathfrak{t}_{h}(\theta)$ is obviously increasing when and $\theta \geq \theta_{H}^{h}$ and $\theta \leq \theta_{L}^{h}$. Furthermore, $R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-$ $h \mathfrak{t}_{h}(\theta)$ is a continuous function of $\theta$ because $\mathfrak{t}_{h}(\theta)$ is continuous. Thus, it suffices to show that $R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-$ $h \mathfrak{t}_{h}(\theta)$ is increasing in $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$.

Recall that $\mathfrak{t}_{h}(\theta)=\mathfrak{t}_{f}(\theta)$ for $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$. That is, $\mathfrak{t}_{h}(\theta)$ is the FOC solution. Thus, by the Envelope theorem, the derivative of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)$ w.r.t. $\theta$ is given by

$$
\begin{aligned}
\frac{\partial\left(R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)\right)}{\partial \theta} & =-\beta \mathfrak{t}_{f}(\theta) e^{-\beta \mathfrak{t}_{f}(\theta) \theta}\left(\theta+\alpha(\theta)\left(2-\beta \mathfrak{t}_{f}(\theta) \theta\right)\right) \\
& +e^{-\beta \mathfrak{t}_{f}(\theta) \theta}\left(1+\alpha^{\prime}(\theta)\left(1-\beta \mathfrak{t}_{f}(\theta) \theta\right)\right) \\
& =h \frac{\mathfrak{t}_{f}(\theta)}{\theta}+e^{-\beta \mathfrak{t}_{f}(\theta) \theta}\left(1+\alpha^{\prime}(\theta)\left(1-\beta \mathfrak{t}_{f}(\theta) \theta\right)\right) \geq 0
\end{aligned}
$$

where the inequality holds because, as we show in Lemma C.1.6, $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$, and the second equality follows from the FOC, i.e., by the fact that

$$
\left.\frac{\partial R(\theta, t)}{\partial t}\right|_{t=\mathfrak{t}_{f}(\theta)}-h=-\beta \theta e^{-\beta \mathfrak{t}_{f}(\theta) \theta}\left(\theta+\alpha(\theta)\left(2-\beta \theta \mathfrak{t}_{f}(\theta)\right)\right)-h=0
$$

Finally, since $R\left(\underline{\theta}_{L}, \mathfrak{t}_{h}\left(\underline{\theta}_{L}\right)\right)-h \mathfrak{t}_{h}\left(\underline{\theta}_{L}\right)=0$, we have $R\left(\theta, \mathfrak{t}_{h}(\theta)\right)-h \mathfrak{t}_{h}(\theta) \geq 0$ for $\theta \geq \underline{\theta}_{L}$; see definition of $\underline{\theta}_{L}$ in Eq. (3.7).

## C.1.8 Discussing the Assumption in Theorem 3.4.4

In this section, we discuss the assumption in Theorem 3.4.4. This assumption requires that the solution of equation $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$ to be unique.

The following lemma shows that for any $h \in\left[H_{l}, H_{h}\right], R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$ has a unique solution if the solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)=0$ is unique. In addition, it shows that $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)=0$ has a unique solution when $\alpha^{\prime}(\theta)$ is small enough.

Lemma C.1.13. If the solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)=0$ is unique, then, for any $h \in\left[H_{l}, H_{h}\right]$, $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)=0$ has a unique solution. Furthermore, the solution of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)=0$ is unique if $\alpha^{\prime}(\theta) \leq \frac{(\sqrt{5}+1)^{2}}{2} \approx 5.2$ for any $\theta \leq \tilde{\theta}$ where $\tilde{\theta}$ solves $2 \tilde{\theta}+\alpha(\tilde{\theta})=0$ and $H_{l}$ and the FOC solution $\mathfrak{t}_{f}(\cdot)$ are defined in Eq. (3.7).

The proof of Lemma C.1.13 is given at the end of this section. Note that for the Uniform and Exponential distributions, we have $\alpha^{\prime}(\theta) \leq 5.2$. In fact, for the Uniform distribution $U(a, b)$, we have $\alpha^{\prime}(\theta)=1$ for any $\theta \in[a, b]$ where $a<b$ and $a, b \in \mathbb{R}$. For the Exponential distribution with rate $\lambda \geq 0, \alpha^{\prime}(\theta)=0$ for
any $\theta \geq 0$. Furthermore, for a truncated Normal distribution with mean $\mu$, standard deviation $\sigma$, and cut-off greater than $\mu-\sigma$, we have $\alpha^{\prime}(\theta) \leq 4.48$ for any $\theta \geq(\mu-\sigma)$. Note that the domain of the truncated Normal distribution with cut-off $\mathcal{C}$ is $[\mathcal{C}, \infty)$.

## Proof. Proof of Lemma C.1.13

First, we show that if the solution of Eq. (C.19) is unique at $h=H_{l}$, then this equation has a unique solution for any $h \in\left[H_{l}, H_{h}\right]$.

$$
\begin{equation*}
R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathbf{t}_{f}(\theta)=0 \tag{C.19}
\end{equation*}
$$

By Lemma C.1.3, $\tilde{\theta}$ solves $R\left(\tilde{\theta}, \mathfrak{t}_{f}(\tilde{\theta})\right)-H_{l} \mathfrak{t}_{f}(\tilde{\theta})=0$ where $2 \tilde{\theta}+\alpha(\tilde{\theta})=0$ and $1-\beta \mathfrak{t}_{f}(\tilde{\theta}) \tilde{\theta}=0$. Then, by our assumption, $\tilde{\theta}$ is the unique solution of Eq. (C.19) at $h=H_{l}$. This assumption and the proof of Lemma C.1.11 imply that for any $h>H_{l}$, any solutions of Eq. (C.19) satisfy the following property: $1-\beta \theta \mathfrak{t}_{f}(\theta) \geq 0$.

Next, we use this property to show that for any $h \in\left[H_{l}, H_{h}\right]$, there is only one solution to Eq. (C.19). Let $\theta_{0}$ solve Eq. (C.19). By the Envelope theorem, the derivative of $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)$ w.r.t. $\theta$ at $\theta_{0}$ is given by

$$
\begin{align*}
\left.\frac{\partial\left(R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)\right)}{\partial \theta}\right|_{\theta=\theta_{0}} & =-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) e^{-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}}\left(\theta_{0}+\alpha\left(\theta_{0}\right)\left(2-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}\right)\right) \\
& +e^{-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}}\left(1+\alpha^{\prime}\left(\theta_{0}\right)\left(1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}\right)\right) \\
& =h \frac{\mathfrak{t}_{f}\left(\theta_{0}\right)}{\theta_{0}}+e^{-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}}\left(1+\alpha^{\prime}\left(\theta_{0}\right)\left(1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}\right)\right)>0 \tag{C.20}
\end{align*}
$$

where the second equality follows from the FOC, i.e., $\left.\frac{\partial\left(R\left(\theta_{0}, t\right)-h t\right)}{\partial t}\right|_{\mathfrak{t}_{f}\left(\theta_{0}\right)}=0$ and the inequality holds because $\left(1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}\right) \geq 0$. By the above equation the derivative of $R\left(\theta_{0}, \mathfrak{t}_{f}\left(\theta_{0}\right)\right)-h \mathfrak{t}_{f}\left(\theta_{0}\right)$ w.r.t. $\theta_{0}$ is always positive. This implies that Eq. (C.19) has a unique solution.

Next, we show that at $h=H_{l}$, the solution of Eq. (C.19) is unique if for any $\theta \leq \tilde{\theta}$, $\alpha^{\prime}(\theta) \leq \frac{(\sqrt{5}+1)^{2}}{2} \approx 5.2$.

We first argue that any $\theta>\tilde{\theta}$ cannot solve Eq. (C.19). To this end, we use the proof of Lemma C.1.3 where we show $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$ for any $\theta \geq \tilde{\theta}$. The fact that $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$ for any $\theta \geq \tilde{\theta}$ implies that
$\frac{\partial\left(R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-h \mathfrak{t}_{f}(\theta)\right)}{\partial \theta}>0$; see Eq. (C.20). Then, considering the fact that $R\left(\tilde{\theta}, \mathfrak{t}_{f}(\tilde{\theta})\right)-H_{l} \mathfrak{t}_{f}(\tilde{\theta})=0$, we have $R\left(\theta, \mathfrak{t}_{f}(\theta)\right)-H_{l} \mathfrak{t}_{f}(\theta)>0$ for any $\theta \geq \tilde{\theta}$.

Next, we show that any $\theta<\tilde{\theta}$ cannot solve Eq. (C.19). Let $\zeta=\beta \theta \mathfrak{t}_{f}(\theta)$. Then, Eq. (C.19) at $h=H_{l}$ can be written as

$$
G(\theta, \zeta):=\beta \theta e^{-\zeta}(\theta+\alpha(\theta)(1-\zeta))-H_{l} \zeta=0 .
$$

We assume, contrary to our result, that there exists $\theta_{0}<\tilde{\theta}$ that solves Eq. (C.19). Then, we show that we have $\left.\frac{\partial G}{\partial \theta}\right|_{\theta=\theta_{0}}>0$ and $\left.\frac{\partial G}{\partial \theta}\right|_{\theta=\tilde{\theta}}>0$. This implies that there cannot exist $\theta_{0}<\tilde{\theta}$ that solves Eq. (C.19).

We consider the following two cases: i- $1-\zeta \geq 0$ and ii- $1-\zeta<0$.
Case i- By the FOC, we have $\frac{\partial G}{\partial \zeta}=0$. This implies that

$$
\begin{align*}
\left.\frac{\partial G}{\partial \theta}\right|_{\theta=\theta_{0}} & =\beta e^{-\zeta}\left(\theta_{0}+\alpha\left(\theta_{0}\right)(1-\zeta)\right)+\beta \theta_{0} e^{-\zeta}\left(1+\alpha^{\prime}\left(\theta_{0}\right)(1-\zeta)\right) \\
& =\frac{H_{l} \zeta}{\theta_{0}}+\beta \theta_{0} e^{-\zeta}\left(1+\alpha^{\prime}\left(\theta_{0}\right)(1-\zeta)\right) \geq 0 \tag{C.21}
\end{align*}
$$

where the second equation holds because $G\left(\theta_{0}, \zeta\right)=0$ and the inequality holds because $1-\zeta \geq 0$. Note that the above equation also implies that $\left.\frac{\partial G(\theta, \zeta)}{\partial \theta}\right|_{\theta=\tilde{\theta}}>0$ considering the fact that at $\theta=\tilde{\theta}$, we have $1-\zeta=1-\beta \tilde{\theta} \mathfrak{t}_{f}(\tilde{\theta})=0$.

Case ii- Next we focus on the case of $1-\zeta<0$. In the following, we show when $1-\zeta<0$ and $\alpha^{\prime}(\theta) \leq \frac{(\sqrt{5}+1)^{2}}{2} \approx 5.2$ for any $\theta \leq \tilde{\theta}$, we get $\left.\frac{\partial G}{\partial \theta}\right|_{\theta=\theta_{0}}>0$. By defintion,

$$
\begin{align*}
\left.\frac{\partial G}{\partial \theta}\right|_{\theta=\theta_{0}} & =\beta e^{-\zeta}\left(2 \theta_{0}+\left(\theta_{0} \alpha^{\prime}\left(\theta_{0}\right)+\alpha\left(\theta_{0}\right)\right)(1-\zeta)\right) \\
& \left.\geq \beta e^{-\zeta}\left(2 \theta_{0}+\left(\theta_{0} \alpha^{\prime}\left(\theta_{0}\right)-2 \theta_{0}\right)(1-\zeta)\right)\right) \\
& =\beta e^{-\zeta} \theta_{0}\left(\alpha^{\prime}\left(\theta_{0}\right)(1-\zeta)+2 \zeta\right) . \tag{C.22}
\end{align*}
$$

The inequality holds because $1-\zeta<0$ and $\theta_{0} \leq \tilde{\theta}$. Note that for any $\theta_{0} \leq \tilde{\theta}, \alpha\left(\theta_{0}\right) \leq-2 \theta_{0}$. To complete the proof, we show that $\left(\alpha^{\prime}\left(\theta_{0}\right)(1-\zeta)+2 \zeta\right) \geq 0$ when $\alpha^{\prime}\left(\theta_{0}\right) \leq \frac{(\sqrt{5}+1)^{2}}{2}$.

First assume that $\alpha^{\prime}\left(\theta_{0}\right) \leq 2$. Then, we get

$$
\left(\alpha^{\prime}\left(\theta_{0}\right)(1-\zeta)+2 \zeta\right) \geq 2(1-\zeta)+2 \zeta=2>0,
$$

where the first inequality holds because $1-\zeta<0$.

Now, assume that $\alpha^{\prime}\left(\theta_{0}\right) \in\left[2, \frac{(\sqrt{5}+1)^{2}}{2}\right]$. We make use of the following claim.
Claim: Let $\theta_{0}<\tilde{\theta}$ solve Eq. (C.19). Then, $\zeta=1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0} \leq \frac{1+\sqrt{5}}{2}$.

The proof of the claim is given at the end of the proof of this lemma.
Given that $\alpha^{\prime}\left(\theta_{0}\right) \in\left[2, \frac{(\sqrt{5}+1)^{2}}{2}\right]$, then $\zeta \mapsto\left(\alpha^{\prime}(\theta)(1-\zeta)+2 \zeta\right)$ is decreasing. Then, by the claim, we get

$$
\alpha^{\prime}(\theta)(1-\zeta)+2 \zeta \geq \alpha^{\prime}(\theta)\left(1-\frac{1+\sqrt{5}}{2}\right)+2\left(\frac{1+\sqrt{5}}{2}\right) \geq 0
$$

where the second inequality holds because $\alpha^{\prime}(\theta) \leq \frac{(\sqrt{5}+1)^{2}}{2} \approx 5.2$.

Proof of Claim: Since $\theta_{0}$ solves Eq. (C.19) and $\mathfrak{t}_{f}\left(\theta_{0}\right)$ is the FOC solution, we get

$$
\beta \theta_{0} e^{-\zeta}\left(\theta_{0}+\alpha\left(\theta_{0}\right)(1-\zeta)\right)=H_{l} \zeta \quad \text { and } \quad-\beta \theta_{0} e^{-\zeta}\left(\theta_{0}+\alpha\left(\theta_{0}\right)(2-\zeta)\right)=H_{l}
$$

where the second equation implies that $\zeta \leq 2$. By dividing these two Equations, we get $-\frac{\left(\theta_{0}+\alpha\left(\theta_{0}\right)(1-\zeta)\right)}{\left(\theta_{0}+\alpha\left(\theta_{0}\right)(2-\zeta)\right)}-$ $\zeta=0$. This can be simplified as

$$
\theta_{0}+\alpha\left(\theta_{0}\right) \frac{1+\zeta-\zeta^{2}}{1+\zeta}=0
$$

where for any $\zeta \in[1,2], \zeta \mapsto \frac{1+\zeta-\zeta^{2}}{1+\zeta}$ is decreasing, $\frac{1+\zeta-\zeta^{2}}{1+\zeta}$ crosses zero at $\frac{1+\sqrt{5}}{2}$, and at $\zeta=1$, $\left.\frac{1+\zeta-\zeta^{2}}{1+\zeta}\right|_{\zeta=1}=\frac{1}{2}$. We note that $\theta_{0}+\left.\alpha\left(\theta_{0}\right) \frac{1+\zeta-\zeta^{2}}{1+\zeta}\right|_{\zeta=1}=\theta_{0}+\frac{1}{2} \alpha\left(\theta_{0}\right)<0$ for any $\theta_{0}<\tilde{\theta}$, and $\theta_{0}+\left.\alpha\left(\theta_{0}\right) \frac{1+\zeta-\zeta^{2}}{1+\zeta}\right|_{\zeta=\frac{1+\sqrt{5}}{2}}=\theta_{0} \geq 0$. Then, we can conclude that $\zeta$ that solves $\theta_{0}+\alpha\left(\theta_{0}\right) \frac{1+\zeta-\zeta^{2}}{1+\zeta}=0$ should be less than $\frac{1+\sqrt{5}}{2}$.

## C.1.9 Proof of Theorem 3.5.1

Throughout the proof, for simplicity, we denote $u(\theta, \theta)$ by $u(\theta)$. Here, we characterize the optimal mechanism under valuation function $V(\theta, t)=\theta e^{-g(\theta) t}$ where $g(\cdot)$ satisfies the following assumption.

Assumption 5. For any $\theta \in[\underline{\Theta}, \bar{\Theta}], g(\theta) \geq 0$ and $g^{\prime}(\theta) \geq 0$. Furthermore, $\frac{g^{\prime}(\theta)}{g(\theta)}$ and $t_{g}(\theta)$, defined in Eq. (3.9), are decreasing in $\theta \in[\underline{\Theta}, \bar{\Theta}]$.

Under the valuation function $V(\theta, t)=\theta e^{-g(\theta) t}$, the expected virtual revenue can be written as $\mathrm{E}\left[R_{g}(\theta, t)\right]-u(\underline{\Theta})$ where

$$
\begin{equation*}
R_{g}(\theta, t)=e^{-g(\theta) t}\left(\theta+\alpha(\theta)\left(1-g^{\prime}(\theta) \theta t\right)\right) . \tag{C.23}
\end{equation*}
$$

Then, to characterize the optimal mechanism, one needs to find an optimal solution of the following optimization problem.

$$
\begin{align*}
& \max _{\{\mathfrak{t}(\theta), \mathfrak{p}(\theta):} \max _{\underline{\Theta}, \bar{\Theta}], u(\underline{\Theta}) \geq 0\}} \mathrm{E}\left[\max \left\{R_{g}(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta}), 0\right\}\right] \\
& \text { s.t. } \text { IC and IR constraints, }, \tag{OPT-G}
\end{align*}
$$

where by Lemma 3.2.1, the IC constraints hold if the envelope and interval conditions are satisfied. Similar to the proof of Theorem 3.3.2, we first ignore the interval conditions and consider the following relaxed problem.

$$
\begin{gather*}
\max _{\{\mathfrak{t}(\theta): \theta \in[\underline{\Theta}, \bar{\Theta}], u(\underline{\Theta}) \geq 0\}} \mathrm{E}\left[\max \left\{R_{g}(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta}), 0\right\}\right] \\
\text { s.t. } u(\theta) \geq 0 \quad \theta \in[\underline{\Theta}, \bar{\Theta}], \tag{OPT-G-R}
\end{gather*}
$$

where by the envelope condition given in Lemma 3.2.1, $u(\theta)=u(\underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-\mathfrak{t}(z) g(z)}\left(1-\mathfrak{t}(z) g^{\prime}(z) z\right) d z$.
In the following, we first show that $\mathfrak{t}_{g}(\cdot)$, given in Eq. (3.9), and $u(\underline{\Theta})=0$ solve the above optimization problem. Then, we show that this solution satisfies the interval condition. Therefore, it is an optimal solution of Problem OPT-G.

Lemma C.1.14. If Assumptions 3 and 5 hold, in an optimal solution of the Problem OPT-G-R, time of sales is $\mathfrak{t}_{g}(\theta)$ and $u(\underline{\Theta})=0$ where $\mathfrak{t}_{g}(\theta)$ is given in Eq. (3.9).

We provide the proof of Lemma C.1.14 in Section C.1.14. The proof has two parts. In the first part, we need to show that $\mathfrak{t}_{g}(\theta)$ and $u(\underline{\Theta})=0$ construct a feasible solution of Problem OPT-G-R. To this aim, we use our assumption that $g(\cdot)$ is log-concave; that is, $\frac{g^{\prime}(\theta)}{g(\theta)}$ is decreasing in $\theta$. In the second part, we show that the solution is optimal. To this end, similar to the proof of Theorem 3.3.2, we find an upper bound on the relaxed problem by dualizing the IR constraints. Then we show that $\left(\mathfrak{t}_{g}(\theta), u(\underline{\Theta})=0\right)$ achieves the upper bound.

The last step is to show the optimal solution of the relaxed problem satisfies the interval conditions. That is, for any $\hat{\theta}, \theta \in[\underline{\Theta}, \bar{\Theta}]$ such that $\hat{\theta} \leq \theta$ :

$$
\begin{aligned}
\int_{\hat{\theta}}^{\theta} A_{g}\left(z, \mathfrak{t}_{g}(\hat{\theta})\right) d z & \leq \int_{\hat{\theta}}^{\theta} A_{g}\left(z, \mathfrak{t}_{g}(z)\right) d z, \\
\int_{\hat{\theta}}^{\theta} A_{g}\left(z, \mathfrak{t}_{g}(z)\right) d z & \leq \int_{\hat{\theta}}^{\theta} A_{g}\left(z, \mathfrak{t}_{g}(\theta)\right) d z,
\end{aligned}
$$

where $A_{g}(z, t)=e^{-g(z) t}\left(1-g^{\prime}(z) t z\right)$. To this aim, we show that for any $z \geq \hat{\theta}, A_{g}\left(z, \mathfrak{t}_{g}(\hat{\theta})\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(z)\right)$ and for any $z<\theta, A_{g}\left(z, \mathfrak{t}_{g}(z)\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(\theta)\right)$.

We start by showing that $A_{g}\left(z, \mathfrak{t}_{g}(\hat{\theta})\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(z)\right)$ for any $z \geq \hat{\theta}$. We consider two cases: 1-$\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\hat{\theta}) z\right) \leq 0$ and $2-\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\hat{\theta}) z\right)>0$.

Assume that $\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\hat{\theta}) z\right) \leq 0$. Then,

$$
e^{-g(z) \mathfrak{t}_{g}(\hat{\theta})}\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\hat{\theta}) z\right) \leq 0 \leq e^{-g(z) \mathfrak{t}_{g}(z)}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right),
$$

where the second inequality follows from Lemma C.1.14. There, we show that for any $z \in[\underline{\Theta}, \bar{\Theta}]$, we have $e^{-g(z) \mathfrak{t}_{g}(z)}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) \geq 0$. The above equation implies that $A_{g}\left(z, \mathfrak{t}_{g}(\hat{\theta})\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(z)\right)$.

Now, assume that $\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\hat{\theta}) z\right)>0$. By Assumption $5, \mathfrak{t}_{g}(\cdot)$ is decreasing. This leads to

$$
\left(1-g^{\prime}(z) \mathfrak{t}_{g}(z) z\right) \geq\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\hat{\theta}) z\right) \geq 0, \text { and } e^{-g(z) \mathfrak{t}_{g}(z)} \geq e^{-g(z) \mathfrak{t}_{g}(\hat{\theta})} .
$$

By the above Equations, we have $A_{g}\left(z, \mathfrak{t}_{g}(\hat{\theta})\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(z)\right)$.
Next, we will verify that $A_{g}\left(z, \mathfrak{t}_{g}(z)\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(\theta)\right)$, for $z<\theta$. Since $\mathfrak{t}_{g}(\cdot)$ is decreasing, for any $z<\theta$, we get

$$
0 \leq\left(1-g^{\prime}(z) \mathfrak{t}_{g}(z) z\right) \leq\left(1-g^{\prime}(z) \mathfrak{t}_{g}(\theta) z\right), \text { and } e^{-g(z) \mathfrak{t}_{g}(z)} \leq e^{-g(z) \mathfrak{t}_{g}(\theta)}
$$

where the first inequality holds because as we show in Lemma C.1.14, for any $z \in[\underline{\Theta}, \bar{\Theta}]$, we have $e^{-g(z) \mathfrak{t}_{g}(z)}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) \geq 0$. The above Equations show that $A_{g}\left(z, \mathfrak{t}_{g}(z)\right) \leq A_{g}\left(z, \mathfrak{t}_{g}(\theta)\right)$.

## Proof of Lemma C.1.14

We first show that the allocation $\mathfrak{t}_{g}(\cdot)$ and $u(\underline{\Theta})=0$ construct a feasible solution of Problem OPT-G-R. To this aim, we show that $u(\theta)=\int_{\underline{\Theta}}^{\theta} e^{-g(z) \mathfrak{t}_{g}(z)}\left(1-g^{\prime}(z) \mathfrak{t}_{g}(z) z\right) d z=0$ for any $\theta \leq \theta_{L}^{g}$, and $u(\theta)>0$ otherwise. The former is easy to verify as $\mathfrak{t}_{g}(\theta)=\frac{1}{\theta g^{\prime}(\theta)}$ for $\theta \leq \theta_{L}^{g}$. To verify that $u(\theta)>0$ for any $\theta>\theta_{L}^{g}$, we make use of the following lemma.

Lemma C.1.15. For any $\theta \in[\underline{\Theta}, \bar{\Theta}]$, we have $\left(1-g^{\prime}(\theta) \mathfrak{t}_{g}(\theta) \theta\right) \geq 0$.
To show Lemma C.1.15, we use assumption 5 that $\frac{g^{\prime}(\theta)}{g(\theta)}$ is decreasing.
Next, we show that $\left(\mathfrak{t}_{g}(\cdot), u(\underline{\Theta})=0\right)$ is an optimal solution of Problem OPT-G-R. To this aim, we find an upper bound for the optimal value of Problem OPT-G-R using the weak duality theorem. Then, we will show that $\left(\mathfrak{t}_{g}(\cdot), u(\underline{\Theta})=0\right)$ achieves the upper bound, thus it is optimal.

In the proof, with a slight abuse of notations, we denote the optimal value of Problem OPT-G-R with OPT-G-R.

Upper Bound of OPT-G-R: For any allocation time $\mathfrak{t}(\cdot)$, and Lagrangian function $\lambda:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}^{+}$, we define the following function.

$$
L_{g}(\mathfrak{t}(\cdot), \lambda(\cdot), u(\underline{\Theta}))=\mathrm{E}\left[R_{g}(\theta, \mathfrak{t}(\theta))\right]+\int_{\underline{\Theta}}^{\bar{\Theta}} \lambda(z) u(z) d z-u(\underline{\Theta})
$$

where $u(\theta)=u(\underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-g(z) \mathfrak{t}(z)}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) d z$. Then, considering the fact that $\lambda(\theta) \geq 0$, for any $(\mathfrak{t}(\cdot), u(\underline{\Theta}))$ such that $u(\theta)=u(\underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-g(z) \mathfrak{t}(z)}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) d z \geq 0$, we have we have

$$
\mathrm{E}\left[R_{g}(\theta, \mathfrak{t}(\theta))-u(\underline{\Theta})\right] \leq L_{g}(\mathfrak{t}(\cdot), \lambda(\cdot), u(\underline{\Theta}))
$$

Therefore, for any $\lambda:[\underline{\Theta}, \bar{\Theta}] \rightarrow \mathbb{R}^{+}$,

$$
\begin{equation*}
\max _{(\mathfrak{t}(\cdot), u(\underline{(\underline{)})) \in \mathcal{T}}}\left\{\mathrm{E}\left[R_{g}(\theta, \mathfrak{t}(\theta))\right]-u(\underline{\Theta})\right\} \leq \max _{(\mathfrak{t}(\cdot), u(\underline{(\underline{)}}) \in \mathcal{T}}\left\{L_{g}(\mathfrak{t}(\cdot), \lambda(\cdot), u(\underline{\Theta}))\right\}, \tag{C.24}
\end{equation*}
$$

where

$$
\mathcal{T}=\left\{(\mathfrak{t}(\cdot), u(\underline{\Theta})): \mathfrak{t}(\theta) \geq 0, u(\underline{\Theta})+\int_{\underline{\Theta}}^{\theta} e^{-g(z) \mathfrak{t}(z)}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) d z \geq 0 \text { for any } \theta \in[\underline{\Theta}, \bar{\Theta}]\right\}
$$

is the set of feasible solution. In the following, we will characterize an upper bound for OPT-G-R by evaluating the r.h.s. of the above equation for a specific Lagrangian function $\lambda(\cdot)$, defined below.

$$
\lambda_{g}(\theta)= \begin{cases}\left(f(\theta)\left(\frac{g(\theta)}{g^{\prime}(\theta)}+\alpha(\theta)\right)\right)^{\prime} & \text { if } \theta \leq \theta_{L}^{g}  \tag{C.25}\\ 0 & \text { otherwise }\end{cases}
$$

Lemma C.1.16. If Assumptions 3 and 5 hold, for any $\theta \in[\underline{\Theta}, \bar{\Theta}], \lambda_{g}(\theta) \geq 0$.

In the proof of Lemma C.1.16, we use our assumption that $g^{\prime}(\theta) \geq 0$ and $\frac{g(\theta)}{g^{\prime}(\theta)}$ is increasing in $\theta$.
Then, the result follows by the following claim.

Claim: With a slight abuse of notations, let $\left(t_{\lambda}(\cdot), u_{\lambda}\right)=\arg \max _{(t(\cdot), u(\underline{\Theta})) \in \mathcal{T}}\left\{L_{g}\left(\mathfrak{t}(\cdot), \lambda_{g}(\cdot), u(\underline{\Theta})\right)\right\}$. Then, $t_{\lambda}(\theta)=\mathfrak{t}_{g}(\theta)$ for any $\theta \in[\underline{\Theta}, \bar{\Theta}]$ and $u_{\lambda}=0$. Furthermore, $L\left(\mathfrak{t}_{g}(\cdot), \lambda_{g}(\cdot), u(\underline{\Theta})=0\right)=$ $\mathrm{E}\left[R_{g}\left(\theta, \mathfrak{t}_{g}(\theta)\right)\right]$, where the expectation is taken w.r.t. $\theta$.

Proof of the Claim: By definition,

$$
\begin{equation*}
L_{g}\left(\mathfrak{t}(\cdot), \lambda_{g}(\cdot), u(\underline{\Theta})\right)=\int_{z=\theta_{L}^{g}}^{\bar{\Theta}} R_{g}(z, \mathfrak{t}(z)) f(z) d z+\int_{\underline{\Theta}}^{\theta_{L}^{g}}\left(R_{g}(z, \mathfrak{t}(z)) f(z)+\lambda_{g}(z) u(z)\right) d z-u(\underline{\Theta}) \tag{C.26}
\end{equation*}
$$

where the second term can be rewritten as

$$
\begin{align*}
& \int_{\underline{\Theta}}^{\theta_{L}^{g}} R_{g}(z, \mathfrak{t}(z)) f(z) d z+\int_{\underline{\Theta}}^{\theta_{L}^{g}} u(z) d\left(f(z)\left(\frac{g(z)}{g^{\prime}(z)}+\alpha(z)\right)\right) \\
& =\int_{\underline{\Theta}}^{\theta_{L}^{g}} f(z) e^{-g(z) z}\left(z+\alpha(z)\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right)\right) d z \\
& +\left.u(z) f(z)\left(\frac{g(z)}{g^{\prime}(z)}+\alpha(z)\right)\right|_{\underline{\Theta}} ^{\theta_{L}^{g}}-\int_{\underline{\Theta}}^{\theta_{L}^{g}} e^{-g(z) z}\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) f(z)\left(\frac{g(z)}{g^{\prime}(z)}+\alpha(z)\right) d z \\
& =\int_{\underline{\Theta}}^{\theta_{L}^{g}} f(z) e^{-g(z) \mathfrak{t}(z)}\left(z-\left(1-g^{\prime}(z) t(z) z\right) \frac{g(z)}{g^{\prime}(z)}\right) d z-u(\underline{\Theta}) f(\underline{\Theta})\left(\frac{g(\underline{\Theta})}{g^{\prime}(\underline{\Theta})}+\alpha(\underline{\Theta})\right) \tag{C.27}
\end{align*}
$$

where the second equation follows from Eq. (C.23) and integrating by part, and the last equation follows from definition of $\theta_{L}^{g}$; that is, $\frac{g\left(\theta_{L}^{g}\right)}{g^{\prime}\left(\theta_{L}^{g}\right)}+\alpha\left(\theta_{L}^{g}\right)=0$. By plugging Eq. (C.27) in Eq. (C.26), we have

$$
\begin{aligned}
L_{g}\left(\mathfrak{t}(\cdot), \lambda_{g}(\cdot), u(\underline{\Theta})\right) & =\int_{\theta_{L}^{g}}^{\bar{\Theta}} R_{g}(z, \mathfrak{t}(z)) f(z) d z \\
& +\int_{\underline{\Theta}}^{\theta_{L}^{g}} f(z) e^{-g(z) \mathfrak{t}(z)}\left(z-\left(1-g^{\prime}(z) \mathfrak{t}(z) z\right) \frac{g(z)}{g^{\prime}(z)}\right) d z-u(\underline{\Theta}) f(\underline{\Theta}) \frac{g(\underline{\Theta})}{g^{\prime}(\underline{\Theta})} .
\end{aligned}
$$

Considering the fact that the coefficient of $u(\underline{\Theta})$, i.e., $-f(\underline{\Theta}) \frac{g(\underline{\Theta})}{g^{\prime}(\underline{\Theta})} \leq 0$, to maximize the above equation, we set $u(\underline{\Theta})=0$. That is, $u_{\lambda}=0$. Then,

$$
\begin{aligned}
\max _{(\mathfrak{t}(\cdot), u(\underline{\Theta})=0) \in \mathcal{T}}\left\{L_{g}\left(\mathfrak{t}(\cdot), \lambda_{g}(\cdot), u(\underline{\Theta})=0\right)\right\} & \leq \int_{\theta_{L}^{g}}^{\bar{\Theta}} f(z) \max _{t \geq 0}\left\{R_{g}(z, t)\right\} d z \\
& +\int_{\underline{\Theta}}^{\theta_{L}^{g}} f(z) \max _{t \geq 0}\left\{e^{-g(z) t}\left(z-\left(1-g^{\prime}(z) t z\right) \frac{g(z)}{g^{\prime}(z)}\right)\right\} d z
\end{aligned}
$$

In Lemma C.1.17, we show that for any $\theta \geq \theta_{L}^{g}$, we get $\arg \max _{t \geq 0}\left\{R_{g}(z, t)\right\}=\mathfrak{t}_{g}(z)$. Then the result follows because for any $z \leq \theta_{L}^{g}$, we have

$$
\arg \max _{t \geq 0}\left\{e^{-g(z) t}\left(z-\left(1-g^{\prime}(z) t z\right) \frac{g(z)}{g^{\prime}(z)}\right)\right\}=\frac{1}{g^{\prime}(z) z}
$$

Lemma C.1.17. For any $\theta \geq \theta_{L}^{g}$, we have $\arg \max _{t \geq 0}\{R(\theta, t)\}=\mathfrak{t}_{g}(\theta)$.

## C. 2 Technical Proofs

## C.2.1 Proof of Lemma C.1.6

- $\mathfrak{t}_{f}(\theta)$ is decreasing in $\theta$ : Since $\mathfrak{t}_{f}(\theta)$ is the FOC solution, we have

$$
\left.\frac{\partial(R(\theta, t)-h t)}{\partial t}\right|_{t=\mathfrak{t}_{f}(\theta)}=-\beta \theta e^{-\beta \mathfrak{t}_{f}(\theta) \theta}\left(\theta+\alpha(\theta)\left(2-\beta \theta \mathfrak{t}_{f}(\theta)\right)\right)-h=0 .
$$

Define $W(\theta, t):=\frac{\partial(R(\theta, t)-h t)}{\partial t}=-\beta \theta e^{-\beta t \theta}(\theta+\alpha(\theta)(2-\beta \theta t))-h$. Then, the FOC implies that $W\left(\theta, \mathfrak{t}_{f}(\theta)\right)=0$. Thus,

$$
\frac{\partial \mathfrak{t}_{f}(\theta)}{\partial \theta}=-\frac{W_{\theta}\left(\theta, \mathfrak{t}_{f}(\theta)\right)}{W_{t}\left(\theta, \mathfrak{t}_{f}(\theta)\right)}
$$

where $W_{\theta}(\theta, t)=\left.\frac{\partial W(\theta, t)}{\partial \theta}\right|_{t=\mathfrak{t}_{f}(\theta)}$ and $W_{t}\left(\theta, \mathfrak{t}_{f}(\theta)\right)=\left.\frac{\partial W(\theta, t)}{\partial t}\right|_{t=\mathfrak{t}_{f}(\theta)}$. In the following, we will show that both $W_{\theta}(\theta, t)$ and $W_{t}(\theta, t)$ are non-positive. This implies that $\frac{\partial \mathrm{t}_{f}(\theta)}{\partial \theta} \leq 0$. Throughout the proof, for simplicity, we denote $\mathfrak{t}_{f}(\theta)$ by $t$.

By definition, we get

$$
\begin{aligned}
W_{\theta}(\theta, t) & =-\beta(1-\beta t \theta) e^{-\beta t \theta}(\theta+\alpha(\theta)(2-\beta t \theta))-\beta \theta e^{-\beta t \theta}\left(1+\alpha^{\prime}(\theta)(2-\beta t \theta)-\beta t \alpha(\theta)\right) \\
& =(1-\beta t \theta) \frac{h}{\theta}-\beta \theta e^{-\beta t \theta}\left(1+\alpha^{\prime}(\theta)(2-\beta t \theta)-\beta t \alpha(\theta)\right),
\end{aligned}
$$

where the second equality follows because $W(\theta, t)=0$. Again, by the fact that $W(\theta, t)=0$, we can replace $-\beta \theta e^{-\beta t \theta}$ by $\frac{h}{(\theta+\alpha(\theta)(2-\beta \theta t))}$. Then,

$$
\begin{align*}
W_{\theta} & =(1-\beta t \theta) \frac{h}{\theta}+\frac{h}{\theta+\alpha(\theta)(2-\beta t \theta)}-\beta \theta e^{-\beta t \theta}\left(\alpha^{\prime}(\theta)(2-\beta t \theta)-\beta t \alpha(\theta)\right) \\
& =h(2-\beta t \theta) \frac{\theta+\alpha(\theta)(1-\beta t \theta)}{\theta(\theta+\alpha(\theta)(2-\beta t \theta))}-\beta \theta e^{-\beta t \theta}\left(\alpha^{\prime}(\theta)(2-\beta t \theta)-\beta t \alpha(\theta)\right) \leq 0 . \tag{C.28}
\end{align*}
$$

The inequality holds because by the FOC condition, i.e., $W(\theta, t)=0$, we have $2-\beta t \geq 0$ and $(\theta+\alpha(\theta)(2-\beta t \theta)) \leq 0$, and by our assumption that $R(\theta, t)-h t \geq 0$, we have $\theta+\alpha(\theta)(1-\beta t \theta) \geq 0$. Note that since $R(\theta, t)-h t=e^{-\beta t \theta}(\theta+\alpha(\theta)(1-\beta t \theta))-h t \geq 0$, we get $\theta+\alpha(\theta)(1-\beta t \theta) \geq 0$. Next, we show that $W_{t}(\theta, t) \leq 0$. By definition,

$$
W_{t}(\theta, t)=\beta^{2} \theta^{2} e^{-\beta t \theta}(\theta+\alpha(\theta)(2-\beta t \theta))+\beta^{2} \theta^{2} \alpha(\theta) e^{-\beta t \theta} \leq 0,
$$

where the inequality holds because by the FOC condition, $(\theta+\alpha(\theta)(2-\beta t \theta)) \leq 0$. The above equation along with Eq. (C.28) imply that $\mathfrak{t}_{f}(\cdot)$ is decreasing.

- $\boldsymbol{A}\left(\boldsymbol{\theta}, \mathfrak{t}_{\boldsymbol{h}}(\boldsymbol{\theta})\right) \in[\mathbf{0}, \mathbf{1}]$ : Note that $A\left(\theta, \mathfrak{t}_{h}(\theta)\right)=0$ for $\theta \leq \theta_{L}^{h}$ and is 1 for $\theta \geq \theta_{H}^{h}$. Thus, it suffices to show that $A\left(\theta, \mathfrak{t}_{h}(\theta)\right) \in[0,1]$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$. To this end, in Lemma C.2.1, we will show that $1-\beta \mathfrak{t}_{h}(\theta) \theta \geq 0$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$.

Lemma C.2.1. When $h \leq H_{l}$, then $\left(1-\beta \mathfrak{t}_{h}(\theta) \theta\right) \geq 0$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$.

## Proof of Lemma C.2.1

Here, we show that set $\left\{\theta: \theta>\theta_{L}^{h}\right.$, and $\left.1-\beta \mathfrak{t}_{f}(\theta) \theta=0\right\}$ is empty. That is, there does not exists any $\theta>\theta_{L}^{h}$ with $1-\beta \mathfrak{t}_{f}(\theta) \theta=0$. Then, by the fact that $1-\beta \mathfrak{t}_{f}\left(\theta_{H}^{h}\right) \theta_{H}^{h}=1$ and $1-\beta \mathfrak{t}_{f}\left(\theta_{L}^{h}\right) \theta_{L}^{h}=0$, we have $1-\beta \theta \mathfrak{t}_{f}(\theta) \geq 0$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$.

Assume, contrary to our result, that there exists $\theta_{0}>\theta_{L}^{h}$ that solves $1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}=0$. Then, we show that this cannot happen.

Let $\theta \in\left\{\theta_{L}^{h}, \theta_{0}\right\}$. Since $\mathfrak{t}_{f}(\theta)$ is the FOC solution, we have $\left.\frac{\partial(R(\theta, t)-h t)}{\partial t}\right|_{\mathfrak{t}_{f}(\theta)}=0$. This condition can be rewritten as

$$
W(\theta, \zeta, h):=-\beta \theta e^{-\zeta}(\theta+\alpha(\theta)(2-\zeta))-h=0,
$$

where $\zeta=\beta \theta \mathfrak{t}_{f}(\theta)$. In the following, we will show that for $\theta \in\left\{\theta_{L}^{h}, \theta_{0}\right\}$, we have $\frac{\partial \zeta}{\partial \theta}=-\frac{W_{\theta}}{W_{\zeta}} \leq 0$, where $W_{\theta}:=\frac{\partial W(\theta, \zeta, h)}{\partial \theta}$ and $W_{\zeta}:=\frac{\partial W(\theta, \zeta, h)}{\partial \zeta}$. This implies that there does not exists $\theta_{0}>\theta_{L}^{h}$ that solves $1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}=0$.

To show $\frac{\partial \zeta}{\partial \theta} \leq 0$, we will verify that $W_{\theta} \leq 0$ and $W_{\zeta} \leq 0$. By definition,

$$
W_{\zeta}=\beta \theta e^{-\zeta}(\theta+\alpha(\theta)(2-\zeta))+\beta \theta \alpha(\theta) e^{-\zeta} \leq 0,
$$

where the inequality follows from the FOC, i.e., the fact that $W(\theta, \zeta, h)=0$. To make it more clear, by the FOC, $(\theta+\alpha(\theta)(2-\zeta))<0$ and as a result, $W_{\zeta} \leq 0$.

Next, we show that $W_{\theta} \leq 0$ for $\theta \in\left\{\theta_{L}^{h}, \theta_{0}\right\}$. By definition,

$$
W_{\theta}=-\beta e^{-\zeta}\left(2 \theta+\left(\alpha(\theta)+\theta \alpha^{\prime}(\theta)\right)(2-\zeta)\right) .
$$

By the fact that for $\theta \in\left\{\theta_{L}^{h}, \theta_{0}\right\}$, we have $1-\beta \theta \mathfrak{t}_{f}(\theta)=0$, and thus $\zeta=1$. This shows that

$$
W_{\theta}=-\beta e^{-1}\left(2 \theta+\alpha(\theta)+\theta \alpha^{\prime}(\theta)\right) \leq 0,
$$

where the inequality holds because $\theta_{L}^{h} \geq \tilde{\theta}$ and as a result $2 \theta+\alpha(\theta) \geq 0$ for $\theta \in\left\{\theta_{L}^{h}, \theta_{0}\right\}$. Recall that $\theta_{0}>\theta_{L}^{h}$.

## C.2.2 Proofs of Lemmas in Sections C.1.4 and C.1.5

Proof. Proof of Lemma C.1.7 The proof is naturally divided into two parts. In the first part, we show that $\lambda_{h}(\theta) \geq 0$ for any $\theta<\underline{\theta}_{L}$ and in the second part, we show that $\lambda_{h}(\theta) \geq 0$ for any $\theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right]$.

First Part: By Eq. (C.10), for any $\theta \leq \underline{\theta}_{L}$, we have

$$
\begin{equation*}
\lambda_{h}(\theta)=f^{\prime}(\theta)(2 \theta+\alpha(\theta))+f(\theta)\left(2+\alpha^{\prime}(\theta)\right) \tag{C.29}
\end{equation*}
$$

We note that by definition, we have $\beta e^{-1}\left(\underline{\theta}_{L}\right)^{2}=h$. Thus, given that $h \leq H_{l}=\beta \tilde{\theta}^{2} e^{-1}$, we have $\underline{\theta}_{L} \leq \tilde{\theta}$. This implies that for any $\theta \leq \underline{\theta}_{L}$, we have $2 \theta+\alpha(\theta) \leq 0$. Then, if $f^{\prime}(\theta) \leq 0$, we have $\lambda_{h}(\theta) \geq 0$. Now, assume that $f^{\prime}(\theta)>0$. Then,

$$
\begin{equation*}
\lambda_{h}(\theta) \geq f^{\prime}(\theta) \alpha(\theta)+f(\theta) \alpha^{\prime}(\theta)=(f(\theta) \alpha(\theta))^{\prime}=(F(\theta)-1)^{\prime} \geq 0 \tag{C.30}
\end{equation*}
$$

Second Part: By definition, for any $\theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right]$, we have

$$
\begin{aligned}
\lambda_{h}(\theta) & =f^{\prime}(\theta)\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right)+f(\theta)\left(1+\alpha^{\prime}(\theta)-\frac{h}{\beta \theta^{2} e^{-1}}\right) \\
& \geq f^{\prime}(\theta)\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right)+f(\theta)\left(1+\alpha^{\prime}(\theta)-\frac{h}{\beta\left(\underline{\theta}_{L}\right)^{2} e^{-1}}\right) \\
& =f^{\prime}(\theta)\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right)+f(\theta) \alpha^{\prime}(\theta)
\end{aligned}
$$

where the inequality holds because $\theta \geq \underline{\theta}_{L}$, and the last equation follows from definition of $\underline{\theta}_{L}$. We consider the following two cases.

Case 1: $f^{\prime}(\theta) \leq 0$ : To show $\lambda_{h}(\theta) \geq 0$, we use the fact that for $\theta \geq \underline{\theta}_{L}$, function $\theta \mapsto \theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}$ is increasing in $\theta$. Then, considering the fact that $\theta_{L}^{h}+\alpha\left(\theta_{L}^{h}\right)+\frac{h}{\beta \theta_{L}^{h} e^{-1}}=0$, we have $\left(\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}\right) \leq 0$ for $\theta \in\left[\underline{\theta}_{L}, \theta_{L}^{h}\right]$. This implies that $\lambda_{h}(\theta) \geq 0$ when $f^{\prime}(\theta) \leq 0$.

The derivative of $\theta+\alpha(\theta)+\frac{h}{\beta \theta e^{-1}}$ w.r.t. $\theta$ is given by

$$
1+\alpha^{\prime}(\theta)-\frac{h}{\beta \theta^{2} e^{-1}} \geq 1+\alpha^{\prime}(\theta)-\frac{h}{\beta\left(\underline{\theta}_{L}\right)^{2} e^{-1}}=\alpha^{\prime}(\theta) \geq 0
$$

where the first inequality holds because $\theta \geq \underline{\theta}_{L}$, and the second inequality follows from the definition of $\underline{\theta}_{L}$.

Case 1: $f^{\prime}(\theta)>0$ : In this case, we have

$$
\lambda_{h}(\theta) \geq f^{\prime}(\theta) \alpha(\theta)+f(\theta) \alpha^{\prime}(\theta)=(f(\theta) \alpha(\theta))^{\prime}=(F(\theta)-1)^{\prime} \geq 0
$$

The last inequality completes the proof.

## Proof. Proof of Lemma C.1.8:

Here, we will show that for any $\theta$, the objective function, $R(\theta, t)-h t$ is a unimodular function of $t$ and achieves its maximum at the FOC solution, denoted by $\mathfrak{t}_{f}(\cdot)$. Then, we show that $\arg \max _{t \geq 0}\{R(\theta, t)-$ $h t\}=\max \left\{\mathfrak{t}_{f}(\theta), 0\right\}=\mathfrak{t}_{h}(\theta)$.

To show that the objective function is unimodular, we will make the following observations: 1- The derivative of the objective function w.r.t. $t$ at $t=\frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$ is negative, at $t=-\infty$ is $\infty$, and at $t=\infty$ is negative. 2- For any $t \leq \frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$, the objective function is a concave function of $t$, and for any $t>\frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$, the objective function is a convex function of $t$. These two observations imply that for any given $\theta, R(\theta, t)$ is a unimodular function of $t$, and achieves its maximum at $t<\frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$.

First Part: The derivative of the objective function with respect to $t$ is given by

$$
\begin{equation*}
\frac{\partial R(\theta, t)}{\partial t}-h=-\beta \theta e^{-\beta t \theta}(\theta+\alpha(\theta)(2-\beta \theta t))-h \tag{C.31}
\end{equation*}
$$

Note that as $t$ approaches $-\infty$, the derivative of the objective function with respect to $t$ converges to $\infty$. Furthermore, as $t$ converges to $\infty$, the derivative goes to $-h$. In addition, one can easily show that the derivative is negative at $t=\frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$.

Second Part: The second derivative of the objective function with respect to $t$ is given by

$$
\begin{equation*}
(\beta \theta)^{2} e^{-\beta t \theta}(\theta+\alpha(\theta)(3-\beta \theta t)) . \tag{C.32}
\end{equation*}
$$

It is easy to observe that the second derivative is negative for any $t<\frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$, and is non-negative otherwise. This implies that the objective function is concave for any $t \leq \frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$ and it is convex for any $t>\frac{\theta+3 \alpha(\theta)}{\beta \theta \alpha(\theta)}$.

So far, we established that $R(\theta, t)-h t$ is a unimodular function of $t$ and achieves its maximum at the FOC solution, denoted by $\mathfrak{t}_{f}(\cdot)$. By Lemma C.1.6, the FOC solution is decreasing in $\theta$. This and the fact that $\mathfrak{t}_{f}\left(\theta_{H}^{h}\right)=0$ lead to $\max \left\{\mathfrak{t}_{f}(\theta), 0\right\}=0$ for any $\theta \geq \theta_{H}^{h}$ and $\max \left\{\mathfrak{t}_{f}(\theta), 0\right\}=\mathfrak{t}_{f}(\theta)=\mathfrak{t}_{h}(\theta)$ for any $\theta \in\left[\theta_{L}^{h}, \theta_{H}^{h}\right]$.

Proof. Proof of Lemma C.1.9 Let $G(z, t)=z e^{-\beta t z}(-1+2 \beta t z)-h t$. We show that for any $z \leq \underline{\theta}_{L}$, we have $\max _{t \geq 0}\{G(z, t)\} \leq 0$. First observe that $G(z, t=0)=-z \leq 0$ and $G(z, t=\infty)=-\infty$. Then, to show that $\max _{t \geq 0}\{G(z, t)\} \leq 0$, we will verify that $G(z, t) \leq 0$ at the FOC solution, i.e., at the solution that solves

$$
\frac{\partial G(z, t)}{\partial t}=e^{-\beta t z} \beta z^{2}(3-2 \beta t z)-h=0
$$

We denote the FOC solution by $t_{F}(z)$. We show that $G\left(z, t_{F}(z)\right) \leq 0$.
To this aim, we show that i- $\frac{\partial G\left(z, t_{F}(z)\right)}{\partial z} \geq 0$ when $\left(-1+2 \beta t_{F}(z) z\right) \geq 0$, ii- $\beta z t_{F}(z)$ is increasing in $z$, and iii- $G\left(\underline{\theta}_{L}, t_{F}\left(\underline{\theta}_{L}\right)\right)=0$. The fact that $\beta z t_{F}(z)$ is increasing in $z$ implies that there exists $\hat{z} \in\left[\underline{\Theta}, \underline{\theta}_{L}\right)$ such that $\left(-1+2 \beta t_{F}(z) z\right)>0$ for any $z>\hat{z}$ and $\left(-1+2 \beta t_{F}(z) z\right) \leq 0$. Then, since $G\left(\underline{\theta}_{L}, t_{F}\left(\underline{\theta}_{L}\right)\right)=0$ and $\frac{\partial G\left(z, t_{F}(z)\right)}{\partial z} \geq 0$ when $\left(-1+2 \beta t_{F}(z) z\right) \geq 0$, for any $z>\hat{z}$ we have $G\left(z, t_{F}(z)\right) \leq G\left(\underline{\theta}_{L}, t_{F}\left(\underline{\theta}_{L}\right)\right)=0$, which is the desired result. Furthermore, for any $z \leq \hat{z}, G\left(z, t_{F}(z)\right) \leq 0$ as for this range of $z$, we have $\left(-1+2 \beta t_{F}(z) z\right) \leq 0$.

- Claim i: $\frac{\partial G\left(z, t_{F}(z)\right)}{\partial z} \leq 0$ when $\left(-1+2 \beta t_{F}(z) z\right)>0$. By the envelope theorem, we get

$$
\frac{\partial G\left(z, t_{F}(z)\right)}{\partial z}=e^{-\beta t_{F}(z) z}\left(-1+\beta t_{F}(z) z\left(5-2 \beta t_{F}(z) z\right)\right) \geq 0
$$

where the inequality holds because $\left(-1+2 \beta t_{F}(z) z\right)>0$ and by the FOC $\left(3-2 \beta t_{F}(z) z\right) \geq 0$. To see why note that $x \mapsto-1+x(5-2 x)$ is positive when $x \in\left[\frac{1}{2}, \frac{3}{2}\right]$.

- Claim ii: $z \mapsto\left(\beta z t_{F}(z)\right)$ is an increasing function. Define $\zeta=\beta z t_{F}(z)$. By the FOC, we have $W(z, \zeta):=e^{-\zeta} \beta z^{2}(3-2 \zeta)-h=0$. Then,

$$
\frac{\partial \zeta}{\partial z}=-\frac{\frac{\partial W(z, \zeta)}{\partial z}}{\frac{\partial W(z, \zeta)}{\partial \zeta}}=\frac{e^{-\zeta} \beta z^{2}(5-2 \zeta)}{2 z e^{-\zeta} \beta(3-2 \zeta)} \geq 0
$$

where the inequality holds because by the FOC $3-2 \zeta \geq 0$.

- Claim iii: $G\left(\underline{\theta}_{L}, t_{F}\left(\underline{\theta}_{L}\right)\right)=0$. Note that $t_{F}\left(\underline{\theta}_{L}\right)=\frac{1}{\beta \theta_{L}}$ and as a result,

$$
G\left(\underline{\theta}_{L}, t_{F}\left(\underline{\theta}_{L}\right)\right)=\underline{\theta}_{L} e^{-1}-\frac{h}{\underline{\theta}_{L}}=0,
$$

where the last equation follows from definition of $\underline{\theta}_{L}$.

Proof. Proof of Lemma C.1.11 The proof has two parts. In the first part, we show that when $h \geq H_{l}$ and $\theta \geq \tilde{\theta}$, we have $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$. Then, in the second part of the proof, we show that $\underline{\theta}_{M} \geq \tilde{\theta}$. This implies that $1-\beta \mathfrak{t}_{f}(\theta) \theta \geq 0$ for any $\theta \in\left[\underline{\theta}_{M}, \theta_{H}^{h}\right]$, which is the desired result.

First Part: Here, we show that any solution of $1-\beta \mathfrak{t}_{f}(\theta) \theta=0$, denoted by $\theta_{0}$, is less than equal to $\tilde{\theta}$. Let $\bar{\theta}_{0}$ be the maximum of such solution; that is $\bar{\theta}_{0}=\max \left\{\theta: 1-\beta \mathfrak{t}_{f}(\theta) \theta=0\right\}$. Then, considering the fact that $\bar{\theta}_{0} \leq \tilde{\theta}, 1-\beta \theta_{H}^{h} \mathfrak{t}_{f}\left(\theta_{H}^{h}\right)=1$, and $1-\beta \bar{\theta}_{0} \mathfrak{t}_{f}\left(\bar{\theta}_{0}\right)=0$, we can conclude that $1-\beta \theta \mathfrak{t}_{f}(\theta)>0$ for any $\theta \in\left[\tilde{\theta}, \theta_{H}^{h}\right]$.

Suppose, contrary to our claim, that there exists $\theta_{0}>\tilde{\theta}$ that solves $1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}=0$. By the FOC, we have

$$
\left.\frac{\partial R\left(\theta_{0}, t\right)}{\partial t}\right|_{t=\mathfrak{t}_{f}\left(\theta_{0}\right)}=-\beta \theta_{0} e^{-\beta \theta_{0} \mathfrak{t}_{f}\left(\theta_{0}\right)}\left(\theta_{0}+\alpha\left(\theta_{0}\right)\left(2-\beta \theta_{0} \mathfrak{t}_{f}\left(\theta_{0}\right)\right)\right)-h=0
$$

Since $\theta_{0}$ solves $1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}=0$, we get

$$
\begin{equation*}
\left.\frac{\partial R\left(\theta_{0}, t\right)}{\partial t}\right|_{t=\mathfrak{t}_{f}\left(\theta_{0}\right)}=-\beta \theta_{0} e^{-1}\left(\theta_{0}+\alpha\left(\theta_{0}\right)\right)=h \tag{C.33}
\end{equation*}
$$

We note that $\theta_{0} \mapsto-\beta \theta_{0} e^{-1}\left(\theta_{0}+\alpha\left(\theta_{0}\right)\right)$ is decreasing in $\theta_{0}$. This holds because

$$
\frac{d\left(-\theta_{0}\left(\theta_{0}+\alpha\left(\theta_{0}\right)\right)\right)}{d \theta_{0}}=-\left(2 \theta_{0}+\alpha\left(\theta_{0}\right)\right)-\theta_{0} \alpha^{\prime}\left(\theta_{0}\right) \leq 0
$$

where the inequality follows because $\theta_{0}>\tilde{\theta}$. This implies that

$$
\max _{\theta_{0} \geq \tilde{\theta}}\left\{-\beta \theta_{0} e^{-1}\left(\theta_{0}+\alpha\left(\theta_{0}\right)\right)\right\}=-\beta \tilde{\theta} e^{-1}(\tilde{\theta}+\alpha(\tilde{\theta}))=\beta \tilde{\theta}^{2} e^{-1}=H_{l} .
$$

Then, by Eq. (C.33), we can conclude that when $h>H_{l}$, there does not exists any $\theta_{0}>\tilde{\theta}$ such that $1-\beta \mathfrak{t}_{f}\left(\theta_{0}\right) \theta_{0}=0$.

Second Part: Here, we show that $\underline{\theta}_{M} \geq \tilde{\theta}$. To this aim, we show that $\frac{\partial \theta_{M}}{\partial h} \geq 0$ when $1-\beta \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right) \underline{\theta}_{M} \geq$ 0 . This verifies that $\underline{\theta}_{M}$ increases as we increase $h$ from $H_{l}$. The reason is that at $h=H_{l}$, we have $\underline{\theta}_{M}=\tilde{\theta}$ and $1-\beta \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right) \underline{\theta}_{M}=0$. This implies at $h=H_{l}$, when $h$ is increased, we have $\underline{\theta}_{M} \geq \tilde{\theta}$. Then, by the first part of the lemma, we know that $1-\beta \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right) \underline{\theta}_{M} \geq 0$ when we increase $h$. This allows us to repeat this procedure to show that $\frac{\partial \theta_{M}}{\partial h} \geq 0$ for any $h \geq H_{l}$.

Let $\theta=\underline{\theta}_{M}$ and $\zeta=\beta \mathfrak{t}_{f}(\theta) \theta$. Then, by definition, we have

$$
\begin{aligned}
G(\theta, \zeta, h) & :=\beta \theta e^{-\zeta}(\theta+\alpha(\theta)(1-\zeta))-h \zeta=0 \\
W(\theta, \zeta, h) & :=-\beta \theta e^{-\zeta}(\theta+\alpha(\theta)(2-\zeta))-h=0
\end{aligned}
$$

The first equation follows from the fact that at $\theta=\underline{\theta}_{M}, R\left(\underline{\theta}_{M}, \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right)\right)-h \mathfrak{t}_{f}\left(\underline{\theta}_{M}\right)=0$ and the second equation follows from the FOC, i.e., $\left.\frac{\partial(R(\theta, t)-h t)}{\partial t}\right|_{t=t_{f}(\theta)}=0$. In the following, we show that $\frac{\partial \theta}{\partial h} \geq 0$ when $1-\zeta \geq 0$.

The aforementioned Equations imply that

$$
\begin{gathered}
\frac{\partial G}{\partial \theta} \frac{\partial \theta}{\partial h}+\frac{\partial G}{\partial \zeta} \frac{\partial \zeta}{\partial h}=\zeta \\
\frac{\partial W}{\partial \theta} \frac{\partial \theta}{\partial h}+\frac{\partial W}{d \zeta} \frac{\partial \zeta}{\partial h}=1
\end{gathered}
$$

This leads to

$$
\frac{\partial \theta}{\partial h}=\frac{\left|\begin{array}{cc}
\zeta & \frac{\partial G}{\partial \zeta} \\
1 & \frac{\partial W}{\partial \zeta}
\end{array}\right|}{\left|\begin{array}{cc}
\frac{\partial G}{d \theta} & \frac{\partial G}{d \zeta} \\
\frac{\partial W}{d \theta} & \frac{\partial W}{d \zeta}
\end{array}\right|}=\frac{\zeta \frac{\partial W}{\partial \zeta}-\frac{\partial G}{\partial \zeta}}{\frac{\partial G}{\partial \theta} \frac{\partial W}{\partial \zeta}-\frac{\partial G}{\partial \zeta} \frac{\partial W}{\partial \theta}}
$$

It is easy to observe that $\frac{\partial G}{\partial \zeta}=W(\theta, \zeta, h)=0$. Thus, $\frac{\partial \theta}{\partial h}=\frac{\zeta}{\frac{\partial G}{\partial \theta}}$. In the following, we will show that $\frac{\partial \theta}{\partial h} \geq 0$ by verifying $\frac{\partial G}{\partial \theta} \geq 0$. By definition,

$$
\frac{\partial G}{\partial \theta}=\beta e^{-\zeta}(\theta+\alpha(\theta)(1-\zeta))+\beta \theta e^{-\zeta}\left(1+\alpha^{\prime}(\theta)(1-\zeta)\right)
$$

We note that the first term, i.e., $(\theta+\alpha(\theta)(1-\zeta))$, is non-negative because $G(\theta, \zeta, h)=\beta \theta e^{-\zeta}(\theta+\alpha(\theta)(1-$ $\zeta))-h \zeta=0$. Also, the second term is positive as $1-\zeta \geq 0$. This gives us $\frac{\partial G}{\partial \theta} \geq 0$ and thus $\frac{\partial \theta}{\partial h} \geq 0$.

## C.2.3 Proof of Lemmas in Section C.1.9

Proof. Proof of Lemma C.1.15It is easy to verify that $\left(1-g^{\prime}(\theta) \mathfrak{t}_{g}(\theta) \theta\right)=1$ for any $\theta>\theta_{H}^{g}$, and it is zero for any $\theta \leq \theta_{L}^{g}$. Thus, it suffices to show that $\left(1-g^{\prime}(\theta) \mathfrak{t}_{g}(\theta) \theta\right) \geq 0$ when $\theta \in\left[\theta_{L}^{g}, \theta_{H}^{g}\right]$.

By definition, for any $\theta \in\left[\theta_{L}^{g}, \theta_{H}^{g}\right]$, we have

$$
\left(1-g^{\prime}(\theta) \mathfrak{t}_{g}(\theta) \theta\right)=-\theta\left(\frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)}\right) .
$$

Since $\theta>0$, to show $\left(1-g^{\prime}(\theta) \mathfrak{t}_{g}(\theta) \theta\right)>0$, it suffices to verify that $\frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)} \leq 0$. To that end, we show that $\frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)}$ is decreasing in $\theta$. Then by the fact that $\frac{1}{\alpha\left(\theta_{L}^{g}\right)}+\frac{g^{\prime}\left(\theta_{L}^{g}\right)}{g\left(\theta_{L}^{g}\right)}=0$, we have $\frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)} \leq 0$.

The derivative of $\frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)}$ w.r.t. $\theta$ is given by

$$
\frac{-\alpha^{\prime}(\theta)}{\alpha(\theta)^{2}}+\left(\frac{g^{\prime}(\theta)}{g(\theta)}\right)^{\prime} \leq 0 .
$$

The inequality holds because by Assumption 3, we have $\alpha^{\prime}(\theta) \geq 0$, and by Assumption 5, $\frac{g^{\prime}(\theta)}{g(\theta)}$ is decreasing in $\theta$.

Proof. Proof of Lemma C.1.16 To show the result, we will verify that $\lambda_{g}(\theta) \geq 0$ for any $\theta \leq \theta_{L}^{g}$.
By Eq. (C.25), for any $\theta \leq \theta_{L}^{g}$,

$$
\left.\lambda_{g}(\theta)=f^{\prime}(\theta)\left(\frac{g(\theta)}{g^{\prime}(\theta)}+\alpha(\theta)\right)\right)+f(\theta)\left(\left(\frac{g(\theta)}{g^{\prime}(\theta)}\right)^{\prime}+\alpha^{\prime}(\theta)\right) .
$$

We consider the following cases:
1- $f^{\prime}(\theta) \leq 0$ : Observe that the first term of $\lambda_{g}(\theta)$ is non-negative. This is the case because by Assumption $5, \frac{g(\theta)}{g^{\prime}(\theta)}+\alpha(\theta) \leq 0$ for any $\theta \leq \theta_{L}^{g}$. To see why note that $\frac{g(\theta)}{g^{\prime}(\theta)}+\alpha(\theta)$ is increasing in $\theta$, and
$\frac{g\left(\theta_{L}^{g}\right)}{g^{\prime}\left(\theta_{L}^{g}\right)}+\alpha\left(\theta_{L}^{g}\right)=0$. Also, note that the second term of $\lambda_{g}(\theta)$, i.e., $f(\theta)\left(\left(\frac{g(\theta)}{g^{\prime}(\theta)}\right)^{\prime}+\alpha^{\prime}(\theta)\right)$, is greater than or equal to zero. This holds because by Assumption 5, $\left(\frac{g(\theta)}{g^{\prime}(\theta)}\right)^{\prime} \geq 0$.

1- $f^{\prime}(\theta) \geq 0$ : Since $g^{\prime}(\theta) \geq 0$, we have

$$
\lambda_{g}(\theta) \geq f^{\prime}(\theta) \alpha(\theta)+f(\theta) \alpha^{\prime}(\theta)=(f(\theta) \alpha(\theta))^{\prime} \geq 0
$$

The last inequality holds because $f(\theta) \alpha(\theta)=F(\theta)-1$ is increasing in $\theta$.

Proof. Proof of Lemma C.1.17 Here, we show that for any $\theta \geq \theta_{L}^{g}, \arg \max _{t \geq 0}\left\{R_{g}(\theta, t)\right\}=\mathfrak{t}_{g}(\theta)$. To this end, we show that the objective function, i.e., $R_{g}(\theta, t)$, has an inverted u-shape in $t$. Then, we show that for $\theta \geq \theta_{H}^{g}, R_{g}(\theta, t)$ achieves its maximum at $t=0$, and for $\theta \in\left[\theta_{L}^{g}, \theta_{H}^{g}\right], R_{g}(\theta, t)$ gets maximized at $\mathfrak{t}_{g}(\theta)$, where $\mathfrak{t}_{g}(\theta)$ solves the FOC.

Following the steps in Lemma C.1.8, one can show that $R_{g}(\theta, t)$ has an inverted u-shape and gets maximized at the FOC solution where the FOC solution solves

$$
\begin{equation*}
\frac{\partial R_{g}(\theta, t)}{\partial t}=e^{-g(\theta) t}\left(-g(\theta) \theta-\alpha(\theta) g(\theta)\left(1-g^{\prime}(\theta) \theta t\right)-\alpha(\theta) g^{\prime}(\theta) \theta\right)=0 \tag{C.34}
\end{equation*}
$$

This implies that the FOC solution, denoted by $\mathfrak{t}_{f}(\theta)$, is given by $\frac{g(\theta) \theta+\alpha(\theta) g^{\prime}(\theta) \theta+\alpha(\theta) g(\theta)}{\alpha(\theta) g(\theta) g^{\prime}(\theta) \theta}$.
Next, we show that for any $\theta \geq \theta_{H}^{g}$, the FOC solution is negative. Then, by the fact that $R_{g}(\theta, t)$ has an inverted u-shape, we can conclude that for any $\theta \leq \theta_{H}^{g}, R_{g}(\theta, t)$ gets maximized at $\mathfrak{t}_{g}(\theta)=0$.

Since $g^{\prime}(\theta) \geq 0$, the FOC solution is negative if

$$
g^{\prime}(\theta) \mathfrak{t}_{f}(\theta)=\frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)}+\frac{1}{\theta} \leq 0 .
$$

By Assumptions 3 and $5, \frac{1}{\alpha(\theta)}+\frac{g^{\prime}(\theta)}{g(\theta)}+\frac{1}{\theta}$ is decreasing in $\theta$. Then, considering this and the fact that $g^{\prime}\left(\theta_{H}^{g}\right) \mathfrak{t}_{f}\left(\theta_{H}^{g}\right)=0$, we have $g^{\prime}(\theta) \mathfrak{t}_{f}(\theta) \leq 0$ for any $\theta \geq \theta_{H}^{g}$. This implies that $\mathfrak{t}_{f}(\theta) \leq 0$ for any $\theta \geq \theta_{H}^{g}$. A similar argument can be used to show that the FOC solution is positive for any $\theta<\theta_{H}^{g}$.


[^0]:    ${ }^{1}$ The potential customers in each location correspond to the people in the location who bought one of the top 200 DVDs with the highest national sales volumes during the summer of 2005. We choose 200 DVDs as the cutoff because they account for a large proportion of the sales volumes.

[^1]:    ${ }^{2}$ For each future customer and an assortment, the clairvoyant algorithm knows the probability that the customer would purchase a product from that assortment, but does not know the exact choice that the customer would make. In other words, the clairvoyant algorithm knows the choice model but does not know the realization of the random choices of a future customer.

[^2]:    ${ }^{3}$ The analysis of Liu \& van Ryzin (2008a) and Jasin \& Kumar (2012) extends to the environment with time-varying demand if the demand varies slowly overtime, but not when the demand is volatile.

[^3]:    ${ }^{4}$ This information is communicated to the website that the user is visiting.
    ${ }^{5}$ This information may be identified through each customer's IP address or (opt-in) cellphone's GPS signals; see Steel \& Angwin (2010).

[^4]:    ${ }^{6}$ Note that we use the upper-bound linear program Primal $\left(\left\{z_{t}\right\}_{t=1}^{T}\right)$ as the proxy for the revenue of the optimal clairvoyant algorithm. This can be interpreted as giving even more power to the clairvoyant algorithm since it can now respect inventory constraints only in expectation. However, such additional power would be negligible with large inventory levels which we believe are the more interesting and realistic instances of the problem.

[^5]:    ${ }^{7}$ Bernstein et al. (2011) show that (under certain assumptions) the behavior of the optimal dynamic program is similar to this intuition.

[^6]:    ${ }^{8}$ On the other hand, it is not difficult to show that the Myopic policy obtains at least $\frac{1}{2}$ of the revenue of the benchmark revenue. Hence, ratio $\frac{1}{2}$ is tight.
    ${ }^{9}$ This is because, under the random utility maximization (cf. Talluri \& van Ryzin 2004c), $\phi_{i}^{z}(S)=\operatorname{Pr}\left\{U_{i}^{z} \geq \max _{\ell \in S \cup\{0\}} U_{\ell}^{z}\right\}$ where $\left(U_{0}^{z}, U_{1}^{z}, \ldots, U_{n}^{z}\right)$ is the random utility vector that a customer of type $z$ assigns to each product. The random variables $\left(U_{0}^{z}, U_{1}^{z}, \ldots, U_{n}^{z}\right)$ may be correlated and can have arbitrary distributions. If $j \neq i$, then

    $$
    \phi_{i}^{z}(S \cup\{j\})=\operatorname{Pr}\left\{U_{i}^{z} \geq \max _{\ell \in S \cup\{j\} \cup\{0\}} U_{\ell}^{z}\right\} \leq \operatorname{Pr}\left\{U_{i}^{z} \geq \max _{\ell \in S \cup\{0\}} U_{\ell}^{z}\right\}=\phi_{i}^{z}(S) .
    $$

[^7]:    ${ }^{10}$ See Gallego et al. (2004), Liu \& van Ryzin (2008a) for a similar linear programming formulation in the context of choice-based network revenue management.

[^8]:    ${ }^{11}$ For the purpose of analysis, assume that no product is shown to this customer that is $S^{t}=\{0\}$ for $t \leq \epsilon T$.
    ${ }^{12}$ The dual-heuristic is $1-O(\epsilon)$-competitive for real-time assortment optimization problem, with high probability, if: i) $\max \left\{\frac{r_{i}}{\operatorname{Primal}\left(\left\{z_{t}\right\}_{t=1}^{T}\right)}\right\} \leq \frac{\epsilon}{(n+1)\left(\ln (T)+\ln \left(2^{n}\right)\right)}$ ii) $\frac{1}{c_{\text {MIN }}} \leq \frac{\epsilon^{3}}{(n+1)\left(\ln (T)+\ln \left(2^{n}\right)\right)}$.

[^9]:    ${ }^{13}$ Note that $\bar{\phi}_{i}^{z}(S)$ does not depend on the mechanism, however, $\phi_{i}^{z_{t}}(S)$ is a function of the mechanism since the estimations of the mechanism for each customer type depend on the assortments offered in the past.

[^10]:    ${ }^{14}$ For example, William Poundstone, in his book Priceless, documented the following case: "Williams-Sonoma added a $\$ 429$ breadmaker next to their $\$ 279$ model: Sales of the cheaper model doubled even though practically nobody bought the $\$ 429$ machine." Thompson (2012)

[^11]:    ${ }^{15}$ Two of the names of the DVDs in Table 1.2 are shortened; "Lost" and "The Simpsons", respectively, correspond to "Lost - The Complete First Season" and "The Simpsons - The Complete Sixth Season".

[^12]:    ${ }^{16}$ For a $k$-dimensional Dirichlet distribution $\boldsymbol{\beta}$ with parameters $\alpha_{1}, \ldots, \alpha_{k}$, the mean and variance of $\beta_{i}$ for $i=1,2, \ldots, k$ are respectively given by $\frac{\alpha_{i}}{\alpha_{0}}$ and $\frac{\alpha_{i}\left(\alpha_{0}-\alpha_{i}\right)}{\alpha_{0}^{2}\left(\alpha_{0}+1\right)}$ where $\alpha_{0}=\sum_{j=1}^{k} \alpha_{j}$. Given that $\alpha_{1}=\alpha_{2}=\ldots=\alpha_{k}=\alpha$, the coefficient of variation is $\sqrt{\frac{k-1}{k \alpha+1}}$. The parameter of the Dirichlet distribution, $\alpha$, is chosen so that $\mathrm{E}\left[\beta_{i}\right]=0.1$ for all $i$ and $\beta_{i}$ has the desired coefficient of variation.
    ${ }^{17}$ We get similar results when each type is not equally likely.
    ${ }^{18}$ The Myopic policy corresponds to an Inventory-Balancing algorithm with penalty function $\Psi(x)=\mathbb{1}[x>0]$.
    ${ }^{19}$ When $T$ is known, $\mathrm{E}[T]$ is replaced by $T$.

[^13]:    ${ }^{20}$ Note that when the length of the horizon is known, $\mathrm{E}[T-t+1 \mid T \geq t]$ is simply $T-t+1$.
    ${ }^{21}$ We do not present the results for the Hybrid algorithm that combines LIB solution with LP resolving heuristics since its performance is very close to the one we considered here.

[^14]:    ${ }^{22}$ We choose the interval $[\underline{T}, \bar{T}]$ by first computing the expected number of customers $\mathrm{E}[T]$ from the loading factor and the total initial inventory of each problem class. Then, we set $\underline{T}=0.5 \mathrm{E}[T]$ and $\bar{T}=1.5 \mathrm{E}[T]$.
    ${ }^{23}$ When the degree of uncertainty in the number of customers reduces, the LPR heuristic performs better, e.g., when $\bar{T}=$ $1.3 \times E[T], \underline{T}=0.7 \times \mathrm{E}[T], \mathrm{LF}=1.4$ and $\mathrm{CV}=1$, the $\mathrm{LPR}^{500}$ obtains $94 \%$ of the upper bound. Even in this case IB algorithms beat the LPR heuristic by $3 \%$.

[^15]:    ${ }^{24}$ Policy $A$ stochastically dominates policy $B$ if for any $x, \bar{F}_{A}(x) \geq \bar{F}_{B}(x)$ where $\bar{F}_{A}$ and $\bar{F}_{B}$ is the complementary cumulative distribution function of revenue (as \% of the Upper Bound) for policy $A$ and $B$, respectively. Note that when policy $A$ dominates policy $B$, the worst-case performance of policy $A$ is also better (higher) than that of policy $B$.

[^16]:    ${ }^{25}$ Note that we did not use the optimal solution because it is rather challenging to solve the corresponding linear programs with constrained assortment size. When there is no restriction on the size of the assortment, though the LP has exponentially many variables, we can solve it efficiently using the techniques from Topaloglu (2013).

[^17]:    ${ }^{1}$ Examples include BlueKai, eXelate, Acxiom, PulsePoint, LiveRamp, Neustar, DataLogix, and OpenTracker.

[^18]:    ${ }^{2}$ Major ad exchanges, such as Google AdX, allow publishers to run private auctions (Google DoubleClick Documentation 2016) where they can control the information that is provided to the advertisers.

[^19]:    ${ }^{3}$ In this work, we consider two stage mechanisms where all the agents are selected at the same time (in the first stage) and the item is allocated in the second stage. See the discussion on adaptive selection rules at the end of Section 2.4.

[^20]:    ${ }^{4}$ This assumption to a large extent is without loss of generality. In general, a mechanism can charge any agents in the first round independent of being selected or not. We are not putting any restrictions on the payment in the second round; therefore, any such payment in the first round can be added to the payment in the second round.

[^21]:    ${ }^{5}$ We note that when the cost is less than $\frac{7}{96}$, the All-Access mechanism has an interval-based equilibrium with $a_{1}=a_{2}=0$ and $\beta_{1}=\beta_{2}=1$. Similarly, when the cost is greater than $\frac{7}{48}$, the All-Access mechanism has an interval-based equilibrium with $a_{1}=\beta_{1}$ and $a_{2}=\beta_{2}$.

[^22]:    ${ }^{6}$ The optimality is defined among all two-stage mechanisms. As discussed in the previous section, a mechanism with an adaptive selection rule may obtain higher revenue.

[^23]:    ${ }^{7}$ See Appendix B.5.3 regarding the impact of the number of agents.

[^24]:    ${ }^{8}$ We also relax an assumption in Lu \& Ye (2014) where the mechanism can verify whether or not a selected agent has invested in obtaining information.

[^25]:    ${ }^{1}$ To characterize the optimal mechanism, using the revelation principle, it suffices to focus only on mechanisms in which customers have an incentive to participate, that is, the individual rationality constraints hold, and customers are willing to reveal their private information to the mechanism designer, that is, the incentive compatibility constraints hold (see Myerson 1981).

[^26]:    ${ }^{2}$ Note that the allocation rule $\mathfrak{t}$ for a customer is only a function of the type of customer, and does not depend on the type of other customers, because in our model, each customer is infinitesimal and there is no inventory constraint.

[^27]:    ${ }^{1}$ Consider the sample space of sequences of customer types in I.I.D. model and divide it into groups such that in each group the number of customers of each type is the same for every sequence. Since each group includes all the equally likely permutations of some sequence, every sequence of customers (and their types) in the I.I.D. model can be mapped to a sequence of customers in the random arrival model.

[^28]:    ${ }^{3}$ Note that due to discretization, we only measure the amount of increase in $\chi$ at $\frac{1}{\epsilon}$ points.

[^29]:    ${ }^{6}$ The proof is built upon ideas from Mehta et al. (2007). Our analysis is different, more rigorous, and applies to smaller number of products. For instance, theirs omits the corresponding proof of Lemma A.4.1 which we establish via induction, using the dynamic programming formulation of the problem.

[^30]:    ${ }^{7}$ We do not investigate that a randomized algorithm would be able to outperform the aforementioned policy, but we are not studying that questions in this work.

[^31]:    ${ }^{8}$ This is with abuse of notation and done for the sake of economy of notation, we previously used $\eta^{z}$ as the expected number of customers of type $z$, not the probability.

[^32]:    ${ }^{1}$ We can assume that the seller is one of the opponents with submitted bid of $r$.

[^33]:    ${ }^{2}$ In case of ties, we choose one of them randomly.

[^34]:    ${ }^{3}$ To see that note $\mathrm{E}\left[q_{i} \mid v_{i, 0}=x_{i}, v_{-i, 0}=x_{-i}\right]$ is equal to $\operatorname{Pr}\left[v_{i, 1}+\beta\left(x_{i}\right) \geq \max \left\{\max _{j \in \mathcal{S}(x)}\left\{x_{j}+\delta_{j}+\right.\right.\right.$ $\left.\left.\left.\beta\left(x_{j, 0}\right)\right\}, \max _{j \notin \mathcal{S}(x)}\left\{x_{j}+\beta\left(x_{j, 0}\right)\right\}, \rho\right\}\right]$, where $v_{i, 1}=x_{i}+\delta_{i}$ if $i \in S(x)$ and it is $x_{i}$ otherwise.

[^35]:    ${ }^{1}$ It is easy to observe that in an optimal solution problem OPT-H-R, we need to set $u(\underline{\Theta}, \underline{\Theta})$ to zero.

