

Essays on Online Decisions, Model Uncertainty and Learning

by

Van Vinh Nguyen

Business Administration
Duke University

Date: _____

Approved:

Robert Nau, Co-Supervisor

Bora Keskin, Co-Supervisor

Saša Pekeč

Peng Sun

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Business Administration
in the Graduate School of Duke University
2017

ABSTRACT

Essays on Online Decisions, Model Uncertainty and Learning

by

Van Vinh Nguyen

Business Administration
Duke University

Date: _____

Approved:

Robert Nau, Co-Supervisor

Bora Keskin, Co-Supervisor

Saša Pekeč

Peng Sun

An abstract of a dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in Business Administration
in the Graduate School of Duke University
2017

Copyright © 2017 by Van Vinh Nguyen
All rights reserved except the rights granted by the
Creative Commons Attribution-Noncommercial Licence

Abstract

This dissertation examines optimal solutions in complex decision problems with one or more of the following components: online decisions, model uncertainty and learning. The first model studies the problem of online selection of a monotone subsequence and provides distributional properties of the optimal objective function. The second model studies the robust optimization approach to the decision problem of an auction bidder who has imperfect information about rivals' bids and wants to maximize her worst-case payoff. The third model analyzes the performance of a myopic Bayesian policy and one of its variants in the dynamic pricing problem of a monopolistic insurer who sells a business interruption insurance product over a planning horizon.

To my parents, Nghiem Nguyen and Hua Pham, and my family, for their unconditional love and support.

Contents

Abstract	iv
List of Tables	ix
List of Figures	x
Acknowledgements	xii
1 Introduction	1
2 Optimal Online Selection of a Monotone Subsequence	4
2.1 Introduction	4
2.1.1 Background and Overview	4
2.1.2 Related Literature	6
2.2 Analytic Properties of Value Functions	9
2.2.1 Structure of the Value Functions	9
2.2.2 Intermezzo: Possibilities for De-Poissonization	13
2.2.3 Smoothness of the Value and Threshold Functions	15
2.2.4 Spending Symmetry: Curvature of the Value Functions	15
2.3 Martingale Relations and Analysis	17
2.3.1 Optimality Martingale	17
2.3.2 Conditional Variances	18
2.3.3 Inferences from the Uniform Model	20
2.3.4 A Cauchy-Schwarz Argument	20

2.3.5	Inferences from the Exponential Model	22
2.3.6	Combined Inferences: Variance Bounds in General	25
2.4	The Central Limit Theorem	25
2.5	Concluding Remarks	27
3	Robust Bidding Policies	29
3.1	Introduction	29
3.1.1	Background and Overview	29
3.1.2	Related Literature	34
3.2	Model	35
3.2.1	Discriminatory Auction	36
3.2.2	Generalized Second-price Auction	37
3.2.3	Core-selecting Auction	37
3.2.4	Robust Optimization Formulation	41
3.3	Robust Bidding with Unit Demand	42
3.3.1	Discriminatory Auction	43
3.3.2	Generalized Second-price Auction	45
3.4	Robust Bidding with Multiple Demand	52
3.4.1	Relationship between VCG and the Core	53
3.4.2	Optimal Bidding Policies under Perfect Information	55
3.4.3	Single-minded Bidder	58
3.4.4	Double-minded Bidder	65
3.4.5	Beyond Double-minded Bidders: the Role of Minimax Equality	69
3.5	Concluding Remarks	75
4	Economic Order Quantity with Model Uncertainty and Business Interruption Insurance	77
4.1	Introduction	77

4.1.1	Background and Overview	77
4.1.2	Related Literature	79
4.2	Problem Formulation and Preliminaries	80
4.2.1	Basic Model Elements	80
4.2.2	Single-period Problems	83
4.2.3	Learning Problems	89
4.2.4	Induced Probability Measures and Performance Metrics	92
4.2.5	The Myopic Bayesian Policy	94
4.3	Analysis	95
4.3.1	Preliminary Results	95
4.3.2	Belief Convergence and Regret in Learning ρ	98
4.3.3	Belief Convergence and Regret in Learning ϕ	99
4.4	Concluding Remarks	102
A	Proofs in Chapter 2	104
A.1	Proofs in Section 2.2	104
A.2	Proofs in Section 2.3	113
B	Proofs in Chapter 3	120
B.1	Proof in Section 3.3	120
B.2	Proofs in Section 3.4	123
C	Proofs in Chapter 4	138
C.1	Proofs in Section 4.3	138
	Bibliography	158
	Biography	165

List of Tables

3.1	Summary of settings for studied auction formats.	31
3.2	Bidder valuations under the $L \setminus L \setminus G$ valuation structure.	60
3.3	Bidder valuations under the $LG \setminus L \setminus G$ valuation structure.	67
3.4	Bidder valuations, Example 3.7.	73

List of Figures

3.1	Illustration of Example 3.1 – Comparison of b_1^{RO} and b_1^{EM}	45
3.2	Illustration of Example 3.1 – Bidder 1’s payoff function under b_1^{RO} and b_1^{EM}	45
3.3	Illustration of Lemma 3.1 – The supremum of $b^{(k)}$ is increasing in x on $[l(k)u(k))$, constant on $[u(k), u(k - 1 k))$ and decreasing in x on $[u(k - 1 k), u(k - 1))$	47
3.4	Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have identical click-through rates.	47
3.5	Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have different click-through rates satisfying $\alpha_1(v_1 - u(1 2)) \leq \alpha_2(v_1 - u(2))$	48
3.6	Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have different click-through rates satisfying $\alpha_2(v_1 - u(2 1)) \leq \alpha_1(v_1 - u(1))$	49
3.7	Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have different click-through rates satisfying $\alpha_2(v_1 - u(2)) < \alpha_1(v_1 - u(1 2))$ and $\alpha_1(v_1 - u(1)) < \alpha_2(v_1 - u(2 1))$	49
3.8	Illustration of Example 3.3 – Bidder 1’s payoff function under b_1^{RO} and b_1^{EM} when $v_1 = 4.5$	51
3.9	Illustration of Example 3.3 – Bidder 1’s payoff function under b_1^{RO} and b_1^{EM} when $v_1 = 11.5$	52
3.10	Illustration of Example 3.4 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and truthful policy $b_1^{TR} = (10, 10)$	62
3.11	Illustration of Example 3.4 – Comparison of bidder 1’s payoffs under robust policies $b_1^{RO,1} = (6, 13)$ and $b_1^{RO,2} = (6, 6)$	62

3.12	Illustration of Example 3.4 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and perfect-information policy $b_1^{PI} = (0, 10)$	63
3.13	Illustration of Example 3.5 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and expected-payoff maximizing policy $b_1^{EM} = (4, 12)$ when $v_1 = (10, 10)$	64
3.14	Illustration of Example 3.5 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and expected-payoff maximizing policy $b_1^{EM} = (3, 11)$ when $v_1 = (20, 20)$	65
3.15	Worst-case payoff functions in different cases under LG\L\G setting. (a) $a' \leq a + \bar{b} - \epsilon_b$, winning one item always yields better worst-case payoff. (b) $a + \bar{b} + \epsilon_b < a'$, winning two items always yields better worst-case payoff. (c) and (d) $a + \bar{b} - \epsilon_b < a' \leq a + \bar{b} + \epsilon_b$, two worst-case payoff functions intersects.	70
3.16	Illustration of Example 3.6: Non-convexity of the upper level set of in b_1 of $\pi_1(b_1, b_{-1})$ with $\lambda = 1$	72
3.17	Illustration of Example 3.7: Worst-case payoff functions corresponding to winning one, two and three items.	75
4.1	Illustration of the EOQ model for one selling season: (A) Shipment arrives before demand interval, (B) Shipment arrives during demand interval.	85

Acknowledgements

First of all, I would like to express my sincere gratitude to my advisors, Bora Keskin and Robert Nau, for their continuous support during my PhD study. I have benefited intellectually and professionally from working with both of them. Bora teaches me so much about academic research and work ethic. He is always generous in sharing his career advice and believes in me even when I struggle with making progress. Bob has been supportive and has given me the freedom to pursue my research interests. I have learned a lot and had so much fun working with him on his statistical software RegressIt.

I would also like to thank other members of my dissertation committee, Saša Pekeč and Peng Sun, for their valuable suggestions and extensive time commitment. My special thank goes to Saša who has provided me with tremendous support, both as an advisor for my second year paper and as a PhD program coordinator. He has offered me valuable career advice, as well as constructive feedback when I did not work hard enough to keep up with the PhD program's expectation. Peng has taught me so much in his dynamic programming class and provided me with various useful suggestions for my dissertation.

I also thank many faculty members in decision sciences area and operations management area at Fuqua who have either taught me in their classes or helped me in research. In particular, I'm grateful to Alessandro Arlotto and Michael Steele, whom I had a chance to work with in my first year. I also thank David Brown and Ozan

Candogan who served as committee members of my preliminary exam. My further thank goes to Alex Belloni, Bob Winkler, Robert Swinney and other faculty members who have provided me with valuable career advice throughout the PhD study. I thank the Fuqua School of Business, Duke University, for the generous doctoral fellowship. My sincere gratitude goes to Ms. Bobbie Clinkscales for her support in numerous issues related to the graduate school requirements.

I would also like to thank my great friends at Duke, Asa Palley, Charlotte Agger, Noah Snyder, Kelsey Lucca, Ben McCartney, Tres Lambert, Joowon Kim, Jae Yun Kim, Troy Campbell, Gürkan Yardimci, Jing Hu, Hüseyin Gürkan, Yunke Mai, Sajad Modaresi, Gabriel Bonilla, Hieu Bui, Quang Tran, An Pham, Hien Nguyen and many others. Their friendship, encouragement and laughter make my graduate life a memorable experience.

Last but not least, this dissertation could not be completed without the unconditional love and support from my family. I'm thankful to my parents, Nghiem Nguyen and Hua Pham, for their upbringing, guidance and support. I also thank my brother Hien Nguyen and his family for always being by my side. I'm thankful to my niece Elena Nguyen, who is the strongest girl in the world and is my source of inspiration.

1

Introduction

Models of complex decision making in operations management often involve one or more of the following components: online decisions, model uncertainty and learning. *Online decisions* refer to the choices that a decision maker makes over time based on the observations of data that are revealed sequentially in order to optimize certain objective functions. *Model uncertainty* is related to parameters that characterize a decision maker's modeling of decision problems but the exact values of which are unknown at the time of making decisions. *Learning* is a process by which a decision maker updates her belief about the uncertain model parameters from obtained data. There are three questions on these topics that pique our interest. In an online decision problem, given that one has established an optimal policy, what can we say about distributional properties of the objective function under such a policy? When model uncertainty is present but the problem is one-time nature, what are the relevant objective functions that represent a decision maker's aversion toward such uncertainty and how does one optimize those objectives? Finally, when model uncertainty is present in an online decision problem, how does one make decisions that strike the balance between optimizing an objective function and learning about the

uncertain parameters? Subsequent chapters of this dissertation attempt to address these questions in several contexts.

In Chapter 2, we study the problem of online selection of a monotone subsequence. In particular, we consider a sequence of n independent random variables with a common continuous distribution F , and consider the task of choosing an increasing subsequence where the observations are revealed sequentially and where an observation must be accepted or rejected when it is first revealed. There is a unique selection policy that is optimal in the sense that it maximizes the expected value of the number of selected observations. We investigate the distribution of the number of selected observations under the optimal policy; in particular, we obtain a central limit theorem for this quantity and a detailed understanding of its mean and variance for large n .

Chapter 3 studies the decision problem of an auction bidder who has imperfect information about the rivals' bids and wants to maximize her worst-case payoff. This information is modeled via an uncertainty set consisting of all possible realizations of rivals' bids. Maximizing the bidder's worst-case payoff over this set yields robust bidding policies that do not depend on distributional assumptions. We study robust bidding policies for three auction formats: discriminatory, generalized second-price and core-selecting auctions. In these settings, establishing the classical minmax equality yields the construction of optimal robust bidding policies. These robust bidding policies could provide better payoff than truthful bidding. Furthermore, compared to expected-payoff maximizing policies, they could result in less allocation risks and provide higher payoff under adversarial realizations of rivals' bids.

Chapter 4 examines the dynamic pricing problem of a monopolistic insurer who sells a business interruption insurance product to a collection of supply chain firms that arrive sequentially. Each supply chain customer faces an Economic Order Quantity model with disruption risks. Due to information asymmetry, the insurer employs

vertical differentiation by providing a menu of contracts with different coverage levels. With regard to model uncertainty in this context, we consider two settings of learning: (a) the insurer is uncertain about the probability of disruption of a shipment; and (b) the insurer is uncertain about the distribution of customer's type. We analyze the performance of a myopic Bayesian policy and its variant in these settings in terms of regret, i.e., the profit loss, with respect to a clairvoyant who observes the underlying value of uncertain parameters. In the first setting, we show that under the MBP, the insurer's belief converges to the true hypothesis at an exponential rate and the regret is bounded by a constant. In the second setting, we show that under a variant of the MBP in which the separability of menus of contracts is enforced in every decision epoch, the insurer's belief converges to the true hypothesis at an exponential rate and the regret is bounded by a constant.

Optimal Online Selection of a Monotone Subsequence

2.1 Introduction

2.1.1 Background and Overview

In the problem of *online* selection of a *monotone increasing* subsequence, a decision maker observes a sequence of independent non-negative random variables $\{X_1, X_2, \dots, X_n\}$ with common continuous distribution F , and the task is to select a subsequence $\{X_{\tau_1}, X_{\tau_2}, \dots, X_{\tau_j}\}$ such that

$$X_{\tau_1} \leq X_{\tau_2} \leq \dots \leq X_{\tau_j}$$

where the indices $1 \leq \tau_1 < \tau_2 < \dots < \tau_j \leq n$ are stopping times with respect to the σ -fields $\mathcal{F}_i = \sigma\{X_1, X_2, \dots, X_i\}$, $1 \leq i \leq n$. In other words, at time i when the random variable X_i is first observed, the decision maker has to choose to accept X_i as a member of the monotone increasing sequence that is under construction, or to reject X_i from any further consideration.

This chapter is based on joint work with Alessandro Arlotto and Michael Steele. The results presented here are also in the paper Arlotto et al. (2015), published in the journal *Stochastic Processes and their Applications*.

We call such a sequence of stopping times a *feasible policy*, and we denote the set of all such policies by $\Pi(n)$. For any $\pi \in \Pi(n)$, we then let $L_n(\pi)$ be the random variable that counts the number of selections made by policy π for the realization $\{X_1, X_2, \dots, X_n\}$; that is,

$$L_n(\pi) = \max\{j : X_{\tau_1} \leq X_{\tau_2} \leq \dots \leq X_{\tau_j} \text{ where } 1 \leq \tau_1 < \tau_2 < \dots < \tau_j \leq n\}.$$

Samuels and Steele (1981) found that for each $n \geq 1$ there is a unique policy $\pi_n^* \in \Pi(n)$ such that

$$\mathbb{E}[L_n(\pi_n^*)] = \sup_{\pi \in \Pi(n)} \mathbb{E}[L_n(\pi)], \quad (2.1)$$

and for such optimal policies one has

$$\mathbb{E}[L_n(\pi_n^*)] \sim (2n)^{1/2} \quad \text{as } n \rightarrow \infty. \quad (2.2)$$

Bruss and Robertson (1991) and Gnedin (1999) showed that one actually has the crisp upper bound

$$\mathbb{E}[L_n(\pi_n^*)] \leq (2n)^{1/2} \quad \text{for all } n \geq 1, \quad (2.3)$$

and, as corollaries of related work, Rhee and Talagrand (1991), Gnedin (1999) and Arlotto and Steele (2011) all found that there is an asymptotic error rate for the lower bound

$$(2n)^{1/2} - O(n^{1/4}) \leq \mathbb{E}[L_n(\pi_n^*)] \quad \text{as } n \rightarrow \infty. \quad (2.4)$$

Here, our main goal is to show that $L_n(\pi_n^*)$ satisfies a central limit theorem.

Theorem 2.1 (Central Limit Theorem for Optimal Online Monotone Selections).

For any continuous distribution F one has for $n \rightarrow \infty$ that

$$(2n)^{1/2} - O(\log n) \leq \mathbb{E}[L_n(\pi_n^*)] \leq (2n)^{1/2}, \quad (2.5)$$

$$\frac{1}{3}\mathbb{E}[L_n(\pi_n^*)] - O(1) \leq \text{Var}[L_n(\pi_n^*)] \leq \frac{1}{3}\mathbb{E}[L_n(\pi_n^*)] + O(\log n), \quad (2.6)$$

and one has the convergence in distribution

$$\frac{3^{1/2}\{L_n(\pi_n^*) - (2n)^{1/2}\}}{(2n)^{1/4}} \implies N(0, 1). \quad (2.7)$$

2.1.2 Related Literature

Two connections help to put this result in context. First, it is useful to recall the analogous problem of offline (or full information) subsequence selection, for which there is a remarkably rich literature. If one knows all of the values $\{X_1, X_2, \dots, X_n\}$ at the time the selections begin, then decision maker can select a maximal increasing subsequence with length

$$L_n = \max\{k : X_{i_1} \leq X_{i_2} \leq \dots \leq X_{i_k} \text{ where } 1 \leq i_1 < i_2 < i_3 < \dots < i_k \leq n\}. \quad (2.8)$$

This *full information* or *offline* length L_n has been studied extensively.

The question of determining the distribution of L_n was first raised by Ulam (1972), but the analysis of L_n was taken up in earnest by Hammersley (1972), Kingman (1973), Logan and Shepp (1977), and Veršik and Kerov (1977) who established in steps that

$$\mathbb{E}[L_n] \sim 2n^{1/2} \quad \text{as } n \rightarrow \infty.$$

Much later H. Kesten conjectured (cf. Aldous and Diaconis, 1999, p. 416) that there should be positive constants α and β such that

$$\mathbb{E}[L_n] = 2n^{1/2} - \alpha n^{1/6} + o(n^{1/6}) \quad \text{and} \quad \{\text{Var}[L_n]\}^{1/2} = \beta n^{1/6} + o(n^{1/6}). \quad (2.9)$$

After subtle progress by Pilpel (1990), Bollobás and Brightwell (1992), Kim (1996), Bollobás and Janson (1997), and Odlyzko and Rains (2000) this conjecture was settled affirmatively by Baik et al. (1999) who proved moreover that $n^{-1/6}(L_n - 2n^{1/2})$ converges in distribution to the famous Tracy-Widom law which had emerged just a bit earlier from the theory of random matrices. The recent monograph of Romik (2014) gives a highly readable account of this development.

One distinction between the online and the offline problems is that, while the means are of the same order in each case, the variances are not of the same order. The standard deviation for offline selection is of order $n^{1/6}$, but by (2.6) the standard deviation for the online selection is of order $n^{1/4}$. Intuitively this difference reflects greater uncertainty in the online selection problem than in the offline problem, but it is harder to imagine why moving to the online formulation would drive one all of the way from the Tracy-Widom law to the Gaussian law.

Second, there are closely related results of Bruss and Delbaen (2001, 2004) that deal with sequential selection where the number of values to be seen is random with a Poisson distribution. Consider the problem of sequential selection of a monotone increasing subsequence from $\{X_1, X_2, \dots, X_{N_\nu}\}$ where N_ν is a Poisson random variable with mean ν that is independent of the sequence $\{X_1, X_2, \dots\}$. Just as in (2.1) there is a unique sequential policy that maximizes the expected number of selections that are made. If we denote this optimal policy by $\pi_{N_\nu}^*$ then as before $L_{N_\nu}(\pi_{N_\nu}^*)$ is the number of selections from $\{X_1, X_2, \dots, X_{N_\nu}\}$ that are made by the policy $\pi_{N_\nu}^*$.

Bruss and Delbaen (2001) proved that, as $\nu \rightarrow \infty$, one has the mean estimate

$$\mathbb{E}[L_{N_\nu}(\pi_{N_\nu}^*)] = (2\nu)^{1/2} + O(\log \nu), \quad (2.10)$$

and the variance estimate

$$\text{Var}[L_{N_\nu}(\pi_{N_\nu}^*)] = \frac{1}{3}(2\nu)^{1/2} + O(\log \nu).$$

Moreover, Bruss and Delbaen (2004) proved that, as $\nu \rightarrow \infty$, one has the convergence in distribution

$$\frac{3^{1/2}\{L_{N_\nu}(\pi_{N_\nu}^*) - (2\nu)^{1/2}\}}{(2\nu)^{1/4}} \implies N(0, 1).$$

One needs to ask if it is possible to “de-Poissonize” these results to get Theorem 2.1, either in whole or in part. We show in Section 2.2.2 that the lower half of (2.5)

can be obtained from (2.10) by an easy de-Poissonization argument; in fact, this is the only proof we know of this bound. In Section 2.2.2 we also explain as best we can, why no further parts of Theorem 2.1 can be obtained by de-Poissonization.

One can further ask if it might be possible to adapt the *methods* of Bruss and Delbaen (2001, 2004) to prove Theorem 2.1. The major benefit of a Poisson horizon is that it gives access to the tools of continuous time Markov processes such as the infinitesimal generator and Dynkin’s martingales. Moreover, in this instance the associated value function $V(t, x)$ can be written as a function of one variable by the space-time transformation $V(t, x) = \bar{V}(t(1 - x))$.

Here we lack these benefits. We work in discrete time with a known finite horizon, and our value function $v_k(s)$ permanently depends on the state s and the time to the horizon k . This puts one a long way from the world of Bruss and Delbaen (2001, 2004). Still, in Section 2.3.3 we give a brief proof of the well-known upper bound (2.3) that echoes an argument of Bruss and Delbaen (2004, pp. 291–292). This seems to be the only instance of an overlap of technique.

Organization of the Analysis

The proof of our central limit theorem has two phases. In the first phase, we investigate the analytic properties of the value functions given by framing the selection problem as a Markov decision problem. Section 2.2.1 addresses the monotonicity and the submodularity of the value functions. We also obtain that the map $n \mapsto \mathbb{E}[L_n(\pi_n^*)]$ is concave, and this is used in Section 2.2.2 to prove the lower half of (2.5); this is our only de-Poissonization argument.

Sections 2.2.3 and 2.2.4 develop smoothness and curvature properties of the value functions. In particular, we find that in the uniform model the value functions are concave as a function of the state variable, but, for the exponential model, they are convex. This broken symmetry is surprisingly useful even though the distribution of

$L_n(\pi_n^*)$ does not depend on the model distribution F .

The second phase of the proof deals with a natural martingale that one obtains from the value functions. This martingale is defined in Section 2.3, and it is used in Sections 2.3.3, 2.3.5 and 2.3.6 to estimate the conditional variances of $L(\pi_n^*)$. These estimates and a martingale central limit theorem are then used in Section 2.4 to complete the proof of Theorem 2.1. Finally, in Section 2.5 we comment briefly on two open problems and the general nature of the methods developed here.

2.2 Analytic Properties of Value Functions

2.2.1 Structure of the Value Functions

We now let $v_k(s)$ denote the expected value of the number of monotone increasing selections under the optimal policy when (i) there are k observations that remain to be seen and (ii) the value of the most recently selected observation is equal to s . The functions $\{v_k : 1 \leq k < \infty\}$ are called the *value functions*, and they can be determined recursively. Specifically, we have the *terminal condition*

$$v_0(s) = 0 \quad \text{for all } s \geq 0,$$

and if we set $F(s) = P(X_i \leq s)$ then for all $k \geq 1$ and $s \geq 0$ we have the recursion

$$v_k(s) = F(s)v_{k-1}(s) + \int_s^\infty \max\{v_{k-1}(s), 1 + v_{k-1}(x)\} dF(x). \quad (2.11)$$

To see why this equation holds, note that with probability $F(s)$ one is presented at time $i = n - k + 1$ with a value X_i that is less than the previously selected value s . In this situation, we do not have the opportunity to select X_i . This leaves us with $k - 1$ observations to be seen and with the value of the last selected observation, s , unchanged. This possibility contributes the term $F(s)v_{k-1}(s)$ to our equation.

Now, if the newly presented value satisfies $s \leq X_i$ then we have the *option* to select or reject $X_i = x$. If we select $X_i = x$, then the sum of our present reward and

expected future reward is $1 + v_{k-1}(x)$. On the other hand, if we choose not to select $X_i = x$, then we have no present reward and the expected future reward is $v_{k-1}(s)$ since the value of the running maximum is not changed. Since X_i has distribution F , the expected optimal contribution is given by the second term of equation (2.11).

The identity (2.11) is called the *Bellman recursion* for the sequential selection problem. In principle, it tells us everything there is to know about the value functions; in particular, it determines

$$\mathbb{E}[L_n(\pi_n^*)] = v_n(0) \quad \text{for all } n \geq 1.$$

Qualitative information can also be extracted from the recursion (2.11). For example, it is immediate from (2.11) that the value functions are always continuous. More refined properties of the value functions may depend on F , and here it is often useful to consider a special subclass of distributions.

Definition 2.1 (Admissible Distribution). A distribution F is said to be *admissible* if there is an open interval $\mathcal{I} \subseteq [0, \infty)$ such that

- (i) F is differentiable on \mathcal{I} ,
- (ii) $F'(x) = f(x) > 0$ for all $x \in \mathcal{I}$, and
- (iii) $\int_{\mathcal{I}} f(x) dx = 1$.

The next lemma illustrates how admissibility can be used. The result is largely intuitive, but the formal proof via (2.11) suggests that some care is needed.

Lemma 2.1 (Monotonicity of Value Functions). *For any distribution F the value functions are non-increasing. Moreover, if F is admissible, then the value functions are strictly decreasing on \mathcal{I} .*

The monotonicity of v_{k-1} tells us that the integrand in (2.11) equals the right maximand $\{1 + v_{k-1}(x)\}$ on a certain initial segment of $[s, \infty)$, and it equals the left

maximand $v_{k-1}(s)$ on the rest of the segment. This observation leads to a useful reformulation of the Bellman recursion; specifically, if we set

$$h_k(s) = \sup\{x \in [s, \infty) : F(x) < 1 \text{ and } v_{k-1}(s) \leq 1 + v_{k-1}(x)\}, \quad (2.12)$$

then the Bellman recursion (2.11) can be written as

$$v_k(s) = \{1 - F(h_k(s)) + F(s)\}v_{k-1}(s) + \int_s^{h_k(s)} \{1 + v_{k-1}(x)\} dF(x). \quad (2.13)$$

The functions $\{h_k : 1 \leq k < \infty\}$ defined by (2.12) are called the *optimal threshold functions*.

If $v_{k-1}(s) \leq 1$, the characterization (2.12) has an informative policy interpretation. Namely, if $v_{k-1}(s) \leq 1$, then the optimal strategy for the decision maker is the greedy strategy where one accepts any arriving observation that is as large as s . On the other hand, if $v_{k-1}(s) > 1$, the optimal decision maker needs to act more conservatively; when k observations remain to be seen, one only accepts the newly arriving observation if it falls in the interval $[s, h_k(s)]$.

When F is admissible, we have the strict monotonicity of v_{k-1} , and this allows a second characterization of the threshold function:

$$h_k(s) \text{ uniquely satisfies } v_{k-1}(s) = 1 + v_{k-1}(h_k(s)) \quad \text{if } v_{k-1}(s) > 1. \quad (2.14)$$

The value $h_k(s)$ of the threshold function thus marks the point of indifference between the optimal acceptance region and the optimal rejection region. The characterization (2.14) also motivates a definition.

Definition 2.2 (Critical Value). If F is admissible, then the unique solution of the equation $v_k(s) = 1$ is called the *critical value*, and it is denoted by s_k^* .

The analytical character of h_k changes at s_k^* , and one has to be attentive to the differing behavior of h_k above and below s_k^* . We will not need this distinction until Section 2.2.3, but it is critical there.

We complete this section by recording two simple (but useful) bounds on the time-difference of the value function. These bounds follow from the characterization (2.12) for the optimal threshold h_k and the monotonicity of the value function v_{k-1} .

Lemma 2.2. *For $s \geq 0$ and $1 \leq k < \infty$, we have the inequalities*

$$0 \leq v_k(s) - v_{k-1}(s) \leq F(h_k(s)) - F(s) \leq 1. \quad (2.15)$$

From a modeler's perspective, this inequality is intuitive since $F(h_k(s)) - F(s)$ can be interpreted as the probability that one selects the next observation when k observations remain to be seen. A formal confirmation of (2.15) illustrates the handiness of the second form (2.13) of the Bellman equation.

If one increases the number k of observations yet to be seen, then the decision maker faces a richer set of future possibilities. This in turn suggests that the decision maker may want to act more conservatively, keeping more powder dry for future action. Specifically, one might guess that $h_{k+1}(s) \leq h_k(s)$ for all $s \in [0, \infty)$ and all $1 \leq k < \infty$. We confirm this guess as a corollary of the next proposition which gives us a pivotally useful property of the value functions.

Proposition 2.1 (Submodularity of the Value Functions). *The sequence of value functions $\{v_k : 1 \leq k < \infty\}$ determined by the Bellman recursion (2.11) is submodular in the sense that for all $1 \leq k < \infty$ one has*

$$v_{k-1}(s) - v_{k-1}(t) \leq v_k(s) - v_k(t) \quad \text{for all } 0 \leq s \leq t < \infty. \quad (2.16)$$

The submodularity guaranteed by Proposition 2.1 is more powerful than one might expect. In particular, it delivers three basic corollaries.

Corollary 2.1 (Monotonicity of Optimal Thresholds). *For the threshold functions characterized by (2.12) we have*

$$h_{k+1}(s) \leq h_k(s) \quad \text{for } 0 \leq s < \infty. \quad (2.17)$$

Corollary 2.2 (Concavity in k of the Value Functions.). *The value functions are concave as functions of k ; that is, for each $s \in [0, \infty)$ and all $k \geq 1$, one has*

$$v_{k+1}(s) - 2v_k(s) + v_{k-1}(s) \leq 0.$$

Corollary 2.3 (Concavity in n of the Expected Length). *For any continuous F , the map $n \mapsto \mathbb{E}[L_n(\pi_n^*)]$ is concave in n .*

This is just a special case of Corollary 2.2 (where one just takes $s = 0$ and $k = n$), but, as we will see in Section 2.2.2, this concavity carries noteworthy force.

Remark 2.1 (Further Context: An Offline Open Problem). It is not known if the corresponding concavity holds for the *offline* monotone subsequence problem. That is, we do not know if the map $n \mapsto \mathbb{E}[L_n]$ is concave where L_n is defined by (2.8). In this case we do know $\mathbb{E}[L_n] = 2n^{1/2} - \alpha n^{1/6} + o(n^{1/6})$ so concavity does seem like a highly plausible conjecture.

2.2.2 Intermezzo: Possibilities for De-Poissonization

If N is an integer valued random variable, then one can consider the problem of sequential selection of a monotone increasing subsequence from the random length sequence $\mathcal{S} = \{X_1, X_2, \dots, X_N\}$. Here, as usual, the elements of the sequence are independent with a common continuous distribution F , and they are also independent of N . We also assume that the decision maker knows F and the distribution of N , but the decision maker does not know the value of N until the sequence \mathcal{S} has been exhausted. We let $L_N(\pi)$ denote the number of selections that are made when one follows a policy π for sequential selection from \mathcal{S} .

Proposition 2.2 (Information Lower Bound). *If $\mathbb{E}[N] = n$ for some $n \in \mathbb{N}$, then*

$$\mathbb{E}[L_N(\pi)] \leq \mathbb{E}[L_n(\pi_n^*)]. \tag{2.18}$$

Proof. The policy π is determined before the realization of N is known, and, for any given k , the policy π is suboptimal when it is used for sequential selection from the sequence $\{X_1, X_2, \dots, X_k\}$. Thus, if we condition on $N = k$, we then have

$$E[L_N(\pi) \mid N = k] \leq \mathbb{E}[L_k(\pi_k^*)]. \quad (2.19)$$

Now, if we take $\phi : [0, \infty) \rightarrow [0, \infty)$ to be the piecewise linear extension of the map $j \mapsto \mathbb{E}[L_j(\pi_j^*)]$, then by Corollary 2.3 we have that ϕ is also concave. Finally, by the suboptimality (2.19), the definition of ϕ , and Jensen's inequality we obtain

$$\mathbb{E}[L_N(\pi)] \leq \sum_{j=0}^{\infty} \mathbb{E}[L_j(\pi_j^*)] \mathbb{P}(N = j) = \mathbb{E}[\phi(N)] \leq \phi(\mathbb{E}[N]) = \mathbb{E}[L_n(\pi_n^*)]. \quad \square$$

The next corollary establishes one of the five assertions of Theorem 2.1. It is an immediate consequence of Proposition 2.2 and the lower half of the mean bound (2.10) from Bruss and Delbaen (2001).

Corollary 2.4. *For any continuous F we have as $n \rightarrow \infty$ that*

$$(2n)^{1/2} - O(\log n) \leq E[L_n(\pi_n^*)]. \quad (2.20)$$

This is a notable improvement over the bound (2.4) that had been established by several earlier investigations; it improves a $O(n^{1/4})$ error bound all the way down to $O(\log n)$. For the central limit theorem (2.7), one could still get along with a lower bound as weak as $(2n)^{1/2} - o(n^{1/4})$.

De-Poissonization and Decision Problems

We get the bound (2.20) by a *de-Poissonization argument* in the sense that a “fixed n ” fact is extracted from a “Poisson N ” fact. Such arguments are common in computer science, combinatorics and analysis; one finds many examples in (Jacquet and Szpankowski, 1998), Flajolet and Sedgewick (2009, Subsection VIII.5.3), and

Korevaar (2004, Chapter 6). Still, Proposition 2.2 is our only instance of a de-Poissonization argument, and the proof of the proposition suggests in part why one may be hard-pressed to find more.

Decision problems are unlike the classical examples mentioned above. The Poisson N problem and the fixed n problem have different optimal policies, and this mismatch forestalls the kind of direct analytical connection one has in the classical examples. Conditioning on $N = k$ does engage the problem, but the suboptimality of the mismatched policy leads only to one-sided relations such as (2.18) and (2.19).

2.2.3 Smoothness of the Value and Threshold Functions

We need to show that the value functions associated with an admissible distribution F are continuously differentiable on \mathcal{I} . As preliminary step, we consider the differentiability of the threshold functions in a region determined by the critical values s_k^* that were defined in Section 2.2.1.

Lemma 2.3 (Differentiability of the Threshold Functions). *Take F to be admissible and take $k > 1$. If v_{k-1} is differentiable on \mathcal{I} and $s \in \mathcal{I} \cap [0, s_{k-1}^*)$, then h_k is differentiable at s , and one has*

$$h'_k(s) = \frac{v'_{k-1}(s)}{v'_{k-1}(h_k(s))} \geq 0. \quad (2.21)$$

Proposition 2.3 (Continuous Differentiability of the Value Functions). *If F is admissible, then for each $1 \leq k < \infty$ the value function $s \mapsto v_k(s)$ is continuously differentiable on \mathcal{I} , and we have*

$$v'_k(s) = -f(s) + \{1 - F(h_k(s)) + F(s)\} v'_{k-1}(s) \quad \text{for } s \in \mathcal{I}. \quad (2.22)$$

2.2.4 Spending Symmetry: Curvature of the Value Functions

For any continuous F the distribution of $L_n(\pi_n^*)$ is the same; this is an invariance property — or a *symmetry*. When one chooses a particular F , say the uniform

distribution, there is a sense in which one *spends symmetry*.

All earlier analyses of $L_n(\pi_n^*)$ passed directly to the uniform distribution without any apparent thought about what might be lost or gained by the transition. Still, it does make a difference how one spends this symmetry. The distribution of $L_n(\pi_n^*)$ is insensitive to F , but the value functions are not.

Specifically, for the uniform distribution the value functions are concave, but for the exponential distribution the value functions are convex. This change of curvature gives one access to different estimates. Over the next several sections we see how specialization of the driving distribution has a big influence on the estimation of variances and conditional variances.

We first break symmetry in the conventional way and take F to be the uniform distribution on $[0, 1]$. Specialization of the Bellman recursion (2.11) defines the sequence of value functions $\{v_k^u : 1 \leq k < \infty\}$, and specialization of the characterization (2.12) defines the sequence of threshold functions $\{h_k^u : 1 \leq k < \infty\}$. Here, we have by (2.12) that $h_k^u(s) \leq 1$ for all $s \in [0, 1]$ and $1 \leq k < \infty$.

Lemma 2.4 (Concavity of the Uniform Value Functions). *For each $1 \leq k < \infty$ the value function $v_k^u : [0, 1] \rightarrow \mathbb{R}^+$ is concave.*

We now break the symmetry in a second way. We take $F(x) = 1 - e^{-x}$ for $x \geq 0$, and we let v_k^e and h_k^e denote the corresponding value and threshold functions. We will shortly find that v_k^e is convex on $[0, \infty)$ for all $k \geq 1$, but we need a preliminary lemma.

Lemma 2.5. *For $1 \leq k < \infty$ and $s \in (0, \infty)$ one has*

$$- \{1 - e^{-h_{k+1}^e(s)+s}\}^{-1} \leq (v_k^e)'(s). \quad (2.23)$$

We now have the main result of this section.

Lemma 2.6 (Convexity of the Exponential Value Functions). *For each $1 \leq k < \infty$, the value function $v_k^e : [0, \infty) \rightarrow \mathbb{R}^+$ is convex on $[0, \infty)$.*

2.3 Martingale Relations and Analysis

2.3.1 Optimality Martingale

One can represent $L_n(\pi_n^*)$ as a sum of functionals of a time non-homogeneous Markov chain. To see how this goes, we first set $M_0 = 0$ and then we define M_i recursively by

$$M_i = \begin{cases} M_{i-1} & \text{if } X_i \notin [M_{i-1}, h_{n-i+1}(M_{i-1})] \\ X_i & \text{if } X_i \in [M_{i-1}, h_{n-i+1}(M_{i-1})], \end{cases} \quad (2.24)$$

so, less formally, M_i is the maximum value of the elements of the subsequence that have been selected up to and including time i . Since we accept X_i if and only if $X_i \in [M_{i-1}, h_{n-i+1}(M_{i-1})]$ and since $L_n(\pi_n^*)$ counts the number of the observations X_1, X_2, \dots, X_n that we accept, we have

$$L_n(\pi_n^*) = \sum_{i=1}^n \mathbb{1}(X_i \in [M_{i-1}, h_{n-i+1}(M_{i-1})]). \quad (2.25)$$

It is also useful to set $L_0(\pi_n^*) = 0$ and to introduce the shorthand,

$$L_i(\pi_n^*) \stackrel{\text{def}}{=} \sum_{j=1}^i \mathbb{1}(X_j \in [M_{j-1}, h_{n-j+1}(M_{j-1})]), \quad \text{for } 1 \leq i \leq n.$$

We now come to a martingale that is central the rest of our analysis.

Proposition 2.4 (Optimality Martingale). *The process $\{Y_i : i = 0, 1, \dots, n\}$ defined by setting*

$$Y_i = L_i(\pi_n^*) + v_{n-i}(M_i) \quad \text{for } 0 \leq i \leq n, \quad (2.26)$$

is a martingale with respect to the filtration $\mathcal{F}_i = \sigma\{X_1, X_2, \dots, X_i\}$, $1 \leq i \leq n$.

Proof. Obviously Y_i is \mathcal{F}_i -measurable and bounded. Moreover, by the definition of $v_{n-i}(s)$ we have $v_{n-i}(M_i) = \mathbb{E}[L_n(\pi_n^*) - L_i(\pi_n^*) \mid \mathcal{F}_i]$, so

$$Y_i = L_i(\pi_n^*) + \mathbb{E}[L_n(\pi_n^*) - L_i(\pi_n^*) \mid \mathcal{F}_i] = \mathbb{E}[L_n(\pi_n^*) \mid \mathcal{F}_i]. \quad \square$$

Since the martingale $\{Y_i : 1 \leq i \leq n\}$ is capped by $L_n(\pi_n^*)$, we also have the explicit identity

$$\mathbb{E}[L_n(\pi_n^*) \mid \mathcal{F}_i] = L_i(\pi_n^*) + v_{n-i}(M_i), \quad (2.27)$$

which is often useful.

2.3.2 Conditional Variances

In (2.26), the term $v_{n-i}(M_i) = \mathbb{E}[L_n(\pi_n^*) - L_i(\pi_n^*) \mid \mathcal{F}_i]$ tells us the expected number of selections that the policy π_n^* will make from $\{X_{i+1}, X_{i+2}, \dots, X_n\}$ given the current value M_i of the running maximum. There is a useful notion of *conditional variance* that is perfectly analogous. Specifically, we set

$$\begin{aligned} w_{n-i}(M_i) &\stackrel{\text{def}}{=} \text{Var}[L_n(\pi_n^*) - L_i(\pi_n^*) \mid \mathcal{F}_i] \\ &= \mathbb{E}[\{L_n(\pi_n^*) - L_i(\pi_n^*) - v_{n-i}(M_i)\}^2 \mid \mathcal{F}_i]. \end{aligned} \quad (2.28)$$

Here, of course, if $i = 0$ we always have $M_0 = 0$ and

$$w_n(M_0) = \text{Var}[L_n(\pi_n^*)].$$

The martingale $\{Y_i, \mathcal{F}_i\}_{i=0}^n$ defined by (2.26) leads in a natural way to an informative representation for the conditional variance (2.28), and one starts with the difference sequence

$$d_j = Y_j - Y_{j-1}, \quad \text{where } 1 \leq j \leq n. \quad (2.29)$$

By (2.26) and telescoping of the sum we have

$$\sum_{j=i+1}^n d_j = L_n(\pi_n^*) - L_i(\pi_n^*) - v_{n-i}(M_i), \quad \text{for } 0 \leq i \leq n, \quad (2.30)$$

so by orthogonality of the martingale differences we get

$$w_{n-i}(M_i) = \text{Var}[L_n(\pi_n^*) - L_i(\pi_n^*) | \mathcal{F}_i] = \sum_{j=i+1}^n \mathbb{E}[d_j^2 | \mathcal{F}_i]. \quad (2.31)$$

This representation for the conditional variance w_{n-i} can be usefully reframed by taking a more structured view of the martingale differences (2.29). Specifically, we write

$$d_j = A_j + B_j, \quad (2.32)$$

where the variable

$$B_j \stackrel{\text{def}}{=} v_{n-j}(M_{j-1}) - v_{n-j+1}(M_{j-1}) \quad (2.33)$$

represents the change in the martingale Y_j when we do not select X_j , and where

$$A_j \stackrel{\text{def}}{=} (1 + v_{n-j}(X_j) - v_{n-j}(M_{j-1})) \mathbb{1}(X_j \in [M_{j-1}, h_{n-j+1}(M_{j-1})]) \quad (2.34)$$

is the *additional* contribution to the change in the martingale Y_j when we do select X_j . Since B_j is \mathcal{F}_{j-1} -measurable, we have

$$\mathbb{E}[d_j^2 | \mathcal{F}_{j-1}] = \mathbb{E}[A_j^2 | \mathcal{F}_{j-1}] + 2B_j \mathbb{E}[A_j | \mathcal{F}_{j-1}] + B_j^2,$$

and we also have $0 = \mathbb{E}[d_j | \mathcal{F}_{j-1}] = B_j + \mathbb{E}[A_j | \mathcal{F}_{j-1}]$, so

$$\mathbb{E}[d_j^2 | \mathcal{F}_{j-1}] = \mathbb{E}[A_j^2 | \mathcal{F}_{j-1}] - B_j^2. \quad (2.35)$$

Now, for $j = i + 1$ to n , we take the conditional expectation in (2.35) with respect to \mathcal{F}_i . When we sum these terms and recall (2.31) we get our final representation for conditional variances

$$w_{n-i}(M_i) = \text{Var}[L_n(\pi_n^*) - L_i(\pi_n^*) | \mathcal{F}_i] = \sum_{j=i+1}^n \{\mathbb{E}[A_j^2 | \mathcal{F}_i] - \mathbb{E}[B_j^2 | \mathcal{F}_i]\}. \quad (2.36)$$

The decomposition (2.36) was our main goal here, but before concluding the section we should make one further inference from (2.32). By the defining representation

(2.12) for h_k we have $0 \leq A_j \leq 1$, and by our bound (2.15) on the value function differences we have $-1 \leq B_j \leq 0$. Hence one has a uniform bound on the martingale differences

$$|d_j| = |A_j + B_j| \leq 1 \quad \text{for } 1 \leq j \leq n. \quad (2.37)$$

2.3.3 Inferences from the Uniform Model

We now consider the decompositions of Section 2.3 when F is the uniform distribution on $[0, 1]$, and we use superscripts to make this specialization explicit. In particular, we let $X_1^u, X_2^u, \dots, X_n^u$ be the underlying sequence of n independent uniformly distributed random variables, and we write M_i^u for the value of the last observation selected up to and including time $i \geq 1$ (and, as usual, we set $M_0^u = 0$). Lemma 2.4 tells us that the value function v_k^u is concave, and this is crucial to the proof of the lower bound for the conditional variance of $L_n^u(\pi_n^*)$.

Proposition 2.5 (Conditional Variance Lower Bound). *For $0 \leq i \leq n$ one has*

$$\frac{1}{3} v_{n-i}^u(M_i^u) - 2 \leq w_{n-i}^u(M_i^u).$$

2.3.4 A Cauchy-Schwarz Argument

If we take the total expectation in (A.13), then we have

$$\mathbb{E}[(h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u)^2] = 2\{\mathbb{E}[M_j^u] - \mathbb{E}[M_{j-1}^u]\}, \quad (2.38)$$

and, since $\mathbb{E}[(h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u)]$ is the unconditional probability that we accept the j 'th element of the sequence, one might hope to estimate $\mathbb{E}[L_n(\pi_n^*)]$ with help from (2.38) and a Cauchy-Schwarz argument.

In fact, by two applications of the Cauchy-Schwarz inequality, we get

$$\begin{aligned}
\mathbb{E}[L_n^u(\pi_n^*)] &= \sum_{j=1}^n \mathbb{E}[h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u] \\
&\leq n^{1/2} \left\{ \sum_{j=1}^n (\mathbb{E}[(h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u)^2]) \right\}^{1/2} \\
&\leq n^{1/2} \left\{ \sum_{j=1}^n \mathbb{E} \left[(h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u)^2 \right] \right\}^{1/2},
\end{aligned}$$

and, when we replace all of the summands using (2.38), we get a telescoping sum

$$\mathbb{E}[L_n^u(\pi_n^*)] \leq n^{1/2} \left\{ 2 \sum_{j=1}^n \{ \mathbb{E}[M_j^u] - \mathbb{E}[M_{j-1}^u] \} \right\}^{1/2} = (2n)^{1/2} \{ \mathbb{E}[M_n^u] \}^{1/2}.$$

We have $\mathbb{E}[M_n^u] < 1$ since the support of M_n^u equals $[0, 1]$, and, since the distribution of $L_n^u(\pi_n^*)$ does not depend on F , we find for all continuous F that

$$\mathbb{E}[L_n(\pi_n^*)] < (2n)^{1/2}. \tag{2.39}$$

This recaptures the mean upper bound (2.3) of Bruss and Robertson (1991) and Gnedin (1999) which was discussed in the introduction.

Here we should note that Bruss and Delbaen (2001) also used a Cauchy-Schwarz argument to show that for Poisson N_ν with mean ν , one has the analogous inequality

$$\mathbb{E}[L_{N_\nu}(\pi_{N_\nu}^*)] \leq (2\nu)^{1/2}. \tag{2.40}$$

We know by Proposition 2.2 that for $\nu = n$ we have $\mathbb{E}[L_{N_n}(\pi_{N_n}^*)] \leq \mathbb{E}[L_n(\pi_n^*)]$ but, even so, the bound (2.40) does not help directly with (2.39) — or vice versa. In addition to the usual issue that “policies do not de-Poissonize,” there is the a real *a priori* possibility that $\mathbb{E}[L_{N_n}(\pi_{N_n}^*)]$ might be much smaller than $\mathbb{E}[L_n(\pi_n^*)]$.

2.3.5 Inferences from the Exponential Model

Now we consider the exponential distribution $F(x) = 1 - e^{-x}$, for $x \geq 0$, and, as before, we use superscripts to make this specialization explicit. Thus $X_1^e, X_2^e, \dots, X_n^e$ denotes a sequence of n independent, mean one, exponential random variables, and M_i^e denotes the value of the last observation selected up to and including time $i \geq 1$ (and, again, we set $M_0^e = 0$). This time Lemma 2.6 provides the critical fact; it tells us that the value function v_k^e is convex, and this is at the heart of the argument.

Proposition 2.6 (Conditional Variance Upper Bound). *For each $0 \leq i \leq n$ one has*

$$w_{n-i}^e(M_i^e) \leq \frac{1}{3} v_{n-i}^e(M_i^e) + \frac{2}{3} \{1 + \log(n - i)\}. \quad (2.41)$$

The proof roughly parallels that of Proposition 2.5, but in this case some integrals are more troublesome to estimate. To keep the argument direct, we extract one calculation as a lemma.

Lemma 2.7. *For $0 \leq s < t < \infty$ one has*

$$\int_s^t \left(\frac{t-x}{t-s} \right)^2 e^{-x} dx \leq \frac{1}{3} (e^{-s} - e^{-t}) + \frac{2}{3} \{e^{-s} - e^{-t}(t-s+1)\}.$$

Proof of Proposition 2.6. Specialization of (2.36) to the exponential model and simplification give us

$$w_{n-i}^e(M_i^e) = \text{Var}[L_n^e(\pi_n^*) - L_i^e(\pi_n^*) | \mathcal{F}_i] \leq \sum_{j=i+1}^n \mathbb{E}[(A_j^e)^2 | \mathcal{F}_i], \quad (2.42)$$

where A_j^e is given by

$$A_j^e = (1 + v_{n-j}^e(X_j^e) - v_{n-j}^e(M_{j-1}^e)) \mathbb{1}(X_j^e \in [M_{j-1}^e, h_{n-j+1}^e(M_{j-1}^e)]).$$

Since M_{j-1}^e is \mathcal{F}_{j-1} -measurable, we have

$$\mathbb{E}[(A_j^e)^2 | \mathcal{F}_{j-1}] = \int_{M_{j-1}^e}^{h_{n-j+1}^e(M_{j-1}^e)} \{1 + v_{n-j}^e(x) - v_{n-j}^e(M_{j-1}^e)\}^2 e^{-x} dx, \quad (2.43)$$

and by Lemma 2.6 the map $x \mapsto 1 + v_{n-j}^e(x) - v_{n-j}^e(M_{j-1}^e)$ is convex in x and non-negative for all $x \in [M_{j-1}^e, h_{n-j+1}^e(M_{j-1}^e)]$.

If $h_{n-j+1}^e(M_{j-1}^e) < \infty$, the line through the left-end point $(M_{j-1}^e, 1)$ and the right-end point $(h_{n-j+1}^e(M_{j-1}^e), 0)$ provides us with an easy upper bound for the integrand (2.43). Specifically, for $x \in [M_{j-1}^e, h_{n-j+1}^e(M_{j-1}^e)]$, we have that

$$\{1 + v_{n-j}^e(x) - v_{n-j}^e(M_{j-1}^e)\}^2 e^{-x} \leq \left(\frac{h_{n-j+1}^e(M_{j-1}^e) - x}{h_{n-j+1}^e(M_{j-1}^e) - M_{j-1}^e} \right)^2 e^{-x}. \quad (2.44)$$

On the other hand, if $h_{n-j+1}^e(M_{j-1}^e) = \infty$, the right-side of (2.44) is replaced by e^{-x} , and (2.44) again holds since $0 \leq \{1 + v_{n-j}^e(x) - v_{n-j}^e(M_{j-1}^e)\} \leq 1$.

We now integrate (2.44) and use the bound of Lemma 2.7; the representation (2.43) then gives us

$$\begin{aligned} \mathbb{E}[(A_j^e)^2 \mid \mathcal{F}_{j-1}] &\leq \frac{1}{3}(e^{-M_{j-1}^e} - e^{-h_{n-j+1}^e(M_{j-1}^e)}) \\ &\quad + \frac{2}{3}\{e^{-M_{j-1}^e} - e^{-h_{n-j+1}^e(M_{j-1}^e)}(1 + h_{n-j+1}^e(M_{j-1}^e) - M_{j-1}^e)\}. \end{aligned} \quad (2.45)$$

Now we just need to interpret the two addends on the right-hand side of (2.45). The first addend is just the probability that observation X_j^e is selected when the value of the running maximum is M_{j-1}^e , that is,

$$\mathbb{E}[\mathbb{1}(X_j^e \in [M_{j-1}^e, h_{n-j+1}^e(M_{j-1}^e)]) \mid \mathcal{F}_{j-1}] = e^{-M_{j-1}^e} - e^{-h_{n-j+1}^e(M_{j-1}^e)}. \quad (2.46)$$

Similarly, the second addend of (2.45) is the one-period expected increment of the current running maximum M_{j-1}^e , or, to be explicit,

$$\begin{aligned} \mathbb{E}[M_j^e - M_{j-1}^e \mid \mathcal{F}_{j-1}] &= \int_{M_{j-1}^e}^{h_{n-j+1}^e(M_{j-1}^e)} (x - M_{j-1}^e) e^{-x} dx \\ &= e^{-M_{j-1}^e} - e^{-h_{n-j+1}^e(M_{j-1}^e)}(1 + h_{n-j+1}^e(M_{j-1}^e) - M_{j-1}^e). \end{aligned} \quad (2.47)$$

Given the two interpretations (2.46) and (2.47), our bound (2.45) now becomes

$$\begin{aligned} \mathbb{E}[(A_j^e)^2 \mid \mathcal{F}_{j-1}] &\leq \frac{1}{3} \mathbb{E}[\mathbb{1}(X_j^e \in [M_{j-1}^e, h_{n-j+1}^e(M_{j-1}^e)]) \mid \mathcal{F}_{j-1}] \\ &\quad + \frac{2}{3} \mathbb{E}[M_j^e - M_{j-1}^e \mid \mathcal{F}_{j-1}]. \end{aligned}$$

Next, we recall the variance upper bound (2.42), take conditional expectations with respect to \mathcal{F}_i , and sum over $i + 1 \leq j \leq n$, to obtain

$$\begin{aligned} w_{n-i}^e(M_i^e) &\leq \frac{1}{3} \mathbb{E}[L_n^e(\pi_n^*) - L_i^e(\pi_n^*) \mid \mathcal{F}_i] + \frac{2}{3} \mathbb{E}[M_n^e - M_i^e \mid \mathcal{F}_i] \\ &= \frac{1}{3} v_{n-i}^e(M_i^e) + \frac{2}{3} \mathbb{E}[M_n^e - M_i^e \mid \mathcal{F}_i], \end{aligned} \tag{2.48}$$

where in the last step we used the martingale identity (2.27).

To complete the proof, we only need to estimate the conditional expectation $\mathbb{E}[M_n^e - M_i^e \mid \mathcal{F}_i]$. We first set $\mathcal{M}_{[i+1, n]}^* = \max\{X_{i+1}^e, X_{i+2}^e, \dots, X_n^e\}$, and then we note that

$$M_n^e - M_i^e \leq \max\{\mathcal{M}_{[i+1, n]}^*, M_i^e\} - M_i^e \leq \mathcal{M}_{[i+1, n]}^*.$$

When we take the conditional expectations and use the independence of \mathcal{F}_i and $\{X_{i+1}^e, X_{i+2}^e, \dots, X_n^e\}$, we get

$$\mathbb{E}[M_n^e - M_i^e \mid \mathcal{F}_i] \leq \mathbb{E}[\mathcal{M}_{[i+1, n]}^* \mid \mathcal{F}_i] = \mathbb{E}[\mathcal{M}_{[1, n-i]}^*]. \tag{2.49}$$

The logarithmic bound for the last term is well-known, but, for completeness, we just note $\mathbb{P}(\mathcal{M}_{[1, n-i]}^* \leq t) = (1 - e^{-t})^{n-i}$, so

$$\mathbb{E}[\mathcal{M}_{[1, n-i]}^*] = \int_0^\infty 1 - (1 - e^{-t})^{n-i} dt = \sum_{j=1}^{n-i} j^{-1} \leq 1 + \log(n - i).$$

This last estimate then combines with the upper bounds (2.48) and (2.49) to complete the proof of (2.41). \square

2.3.6 Combined Inferences: Variance Bounds in General

The variance bounds obtained under the uniform and exponential models are almost immediately applicable to general continuous F . One only needs to make an appropriate translation.

Proposition 2.7. *For any continuous F and for all $0 \leq i \leq n$ one has the conditional variance bounds*

$$\frac{1}{3} v_{n-i}(M_i) - 2 \leq w_{n-i}(M_i) \leq \frac{1}{3} v_{n-i}(M_i) + \frac{2}{3} \{1 + \log(n-i)\}. \quad (2.50)$$

In particular, for $i = 0$ one has $M_0 = 0$ and

$$\frac{1}{3} \mathbb{E}[L_n(\pi_n^*)] - 2 \leq \text{Var}[L_n(\pi_n^*)] \leq \frac{1}{3} \mathbb{E}[L_n(\pi_n^*)] + \frac{2}{3} \{1 + \log n\}. \quad (2.51)$$

2.4 The Central Limit Theorem

Our proof of the central limit theorem for $L_n(\pi_n^*)$ depends on the most basic version of the martingale central limit theorem. Brown (1971), McLeish (1974), and Hall and Heyde (1980) all give variations containing this one.

Proposition 2.8 (Martingale Central Limit Theorem). *For each $n \geq 1$, we consider a martingale difference sequence $\{Z_{n,j} : 1 \leq j \leq n\}$ with respect to the sequence of increasing σ -fields $\{\mathcal{F}_{n,j} : 0 \leq j \leq n\}$. If*

$$\max_{1 \leq j \leq n} \|Z_{n,j}\|_\infty \rightarrow 0 \quad \text{as } n \rightarrow \infty \quad (2.52)$$

and

$$\sum_{j=1}^n \mathbb{E}[Z_{n,j}^2 | \mathcal{F}_{j-1}] \xrightarrow{\text{P}} 1 \quad \text{as } n \rightarrow \infty, \quad (2.53)$$

then we have the convergence in distribution

$$\sum_{j=1}^n Z_{n,j} \Longrightarrow N(0, 1) \quad \text{as } n \rightarrow \infty.$$

For each $n \geq 1$, we consider a driving sequence $X_{n,1}, X_{n,2}, \dots, X_{n,n}$ of independent random variables with the continuous distribution F . We then set

$$Z_{n,j} \stackrel{\text{def}}{=} \frac{3^{1/2}d_{n,j}}{(2n)^{1/4}}, \quad \text{for } 1 \leq j \leq n,$$

where the $d_{n,j}$'s are the differences defined by (2.29), although we now make explicit the dependence of the differences on n . This is a martingale difference sequence with respect to the increasing sequence of σ -fields $\mathcal{F}_{n,j} = \sigma\{X_{n,1}, X_{n,2}, \dots, X_{n,j}\}$, and when we take $i = 0$ in (2.30) we get the basic representation

$$\sum_{j=1}^n Z_{n,j} = \frac{3^{1/2}\{L_n(\pi_n^*) - \mathbb{E}[L_n(\pi_n^*)]\}}{(2n)^{1/4}}.$$

We know from (2.37) that we always have $|d_{n,j}| \leq 1$ so, by our normalization, the negligibility condition (2.52) is trivially valid. The heart of the matter is the proof of the weak law (2.53); more explicitly, we need to show that

$$\sum_{j=1}^n \frac{3 \mathbb{E}[d_{n,j}^2 \mid \mathcal{F}_{n,j-1}]}{(2n)^{1/2}} \xrightarrow{\text{p}} 1 \quad \text{as } n \rightarrow \infty. \quad (2.54)$$

The variance bounds (2.51) and the asymptotic relation (2.2) for the mean imply

$$\text{Var}[L_n(\pi_n^*)] \sim \frac{1}{3} \mathbb{E}[L_n(\pi_n^*)] \sim \frac{(2n)^{1/2}}{3} \quad \text{as } n \rightarrow \infty,$$

and telescoping and orthogonality of the differences $\{d_{n,j} : 1 \leq j \leq n\}$ give us

$$\text{Var}[L_n(\pi_n^*)] = \mathbb{E}\left[\sum_{j=1}^n \mathbb{E}[d_{n,j}^2]\right] = \mathbb{E}\left[\sum_{j=1}^n \mathbb{E}[d_{n,j}^2 \mid \mathcal{F}_{n,j-1}]\right],$$

so the weak law (2.54) will follow from Chebyshev's inequality if one proves that

$$\text{Var}\left[\sum_{j=1}^n \mathbb{E}[d_{n,j}^2 \mid \mathcal{F}_{n,j-1}]\right] = o(n) \quad \text{as } n \rightarrow \infty. \quad (2.55)$$

The proof of Theorem 1 is completed once one confirms the relation (2.55), and the next lemma gives us more than we need.

Lemma 2.8 (Conditional Variance Bound). *If F is continuous, then for $n \geq 1$, one has*

$$\text{Var} \left[\sum_{j=1}^n \mathbb{E}[d_{n,j}^2 \mid \mathcal{F}_{n,j-1}] \right] \leq \{18 + (\log n)^2\} (2n)^{1/2}.$$

2.5 Concluding Remarks

The idea of “spending symmetry” that was mentioned in Section 2.2.4 originates with an instructive essay of Tao (2009, Section 1.4). This notion can be cast in stunning generality, but here it turns out to be resolutely concrete and very useful.

The variance lower bound of (2.51) had been known to us for some years, but dogged analysis of the uniform model left us without an upper bound of comparable quality. A general Markov decision problem (MDP) bound in Arlotto et al. (2014) would give $\text{Var}[L_n(\pi_n^*)] \leq \mathbb{E}[L_n(\pi_n^*)]$, but here the MDP bound is too weak by a factor of three. It cannot serve even as good motivation for a central limit theorem.

With such a long tradition of immediate reduction to the uniform model, it was surprising to see how fruitful it could be to simultaneously use the exponential model — even though the distribution of $L_n(\pi_n^*)$ is the *same* under either model. Still, with different value functions come different qualitative features, and the convexity of the value functions under the exponential model leads in a natural way to the needed upper bound of the variance. This opened up the way to the rest of the analysis.

We mentioned one open problem earlier (see Remark 2.1), and there is a related problem that deserves some thought. In the offline selection problem, the distribution of the length of the longest increasing subsequence of a sequence of n independent uniformly distributed random variables is the same as the distribution of the length of the longest increasing subsequence of a random permutation of the

integers $\{1, 2, \dots, n\}$. This equivalence is lost in the online section problem, and it is unclear how much of Theorem 2.1 can be recaptured.

For example, if we write $L_n^{\text{perm}}(\pi_n^*)$ for the analog of $L_n(\pi_n^*)$ where now one chooses a random permutation of $\{1, 2, \dots, n\}$, then, by an argument of Burgess Davis given in Samuels and Steele (1981), one does have $\mathbb{E}[L_n^{\text{perm}}(\pi_n^*)] \sim (2n)^{1/2}$ as $n \rightarrow \infty$. Unfortunately, mean bounds like those of Theorem 2.1 cannot be achieved in this way, and variance bounds that would be good enough to support a central limit theorem are even more remote. Nevertheless, some analog of Theorem 2.1 is quite likely to be true.

Robust Bidding Policies

3.1 Introduction

3.1.1 Background and Overview

In the past few decades, there has been a tremendous growth in the use of auctions as transaction procedures for sales and purchases of high-value assets. For example, many governments around the world use spectrum auctions to sell telecommunication bandwidths and generate billions of dollars in revenue. Telecom companies who participate in spectrum auctions often have imperfect information about their rivals' business plans and consequently bids and bidding policies. Moreover, given high-stakes one-time nature of such auctions (in which prices auction winners have to pay could reach multiple billions of dollars), the bidder objective could be to maximize the worst-case payoff (rather than maximizing expected payoff). In this paper we study this kind of a bidder decision problem in detail. More precisely, we study optimal bidding policies for a bidder whose objective is to maximize the worst-case payoff.

This chapter is based on joint work with Saša Pekeč and Ozan Candogan.

Our focus on maximizing the worst-case payoff requires describing all possible scenarios, i.e., auction outcomes. Each scenario corresponds to some realization of rivals' bids and the set of all such bid realizations is the bidder's *uncertainty set*. The bidder's *robust bidding problem* is to determine a bidding policy that maximizes her worst-case payoff with respect to the uncertainty set. Solving the robust bidding problem does not require distributional assumptions on rivals' bids, and thus yields a distribution-free *robust bidding policy*.¹

The robust bidding problem is straightforward for the second-price auction and its Vickrey-Clarke-Groves (VCG) auction generalization. These auction formats satisfy the *incentive compatibility* condition, i.e., truthful-bidding (bidding one's valuation) is a weakly-dominant strategy for each bidder (see e.g., Krishna, 2009). As a result, truthful bidding is a straightforward solution to the robust bidding problem since bidders maximize their worst-case payoff by bidding truthfully. On the other hand, when incentive compatibility does not hold, bidding truthfully is not necessarily a robust bidding policy, and the robust bidding problem could be challenging. Moreover, in such settings, the robust bidding policies differ from the optimal policies of the expected payoff maximizing bidder.

We study the robust bidding problem for three auction formats: discriminatory, generalized second-price (GSP) and core-selecting auctions. These auction formats are widely used in practice. Discriminatory auction, a generalization of the first-price auction for a multi-item setting, is used, e.g., in electricity procurement and in sales of U.S. Treasury securities. GSP auction, a generalization of the second-price auction for a multi-item setting, is popular for selling advertisement slots next to Internet search results. Core-selecting auction has been recently adopted for sales of bundled items, such as spectrum licenses and airport take-off/landing rights.

¹ Our modeling approach fits the robust optimization framework, see, e.g., Ben-Tal and Nemirovski (2002).

The three auction formats differ in terms of allocation and payment rules, which are directly related to the manner in which (item) supply and (bidder) demand is treated. Specifically, items could be assumed to be homogenous or item heterogeneity could be handled. Similarly, bidders can be limited to unit demand, i.e., limiting each bidder to winning at most one-item, or could be allowed to have multiple demand and potentially win any number of items. Table 3.1.1 summarizes the settings for supply and demand in the three auction formats.² Analyzing the bidding problem in these auctions allows us to gain insights into the structure and performance of robust bidding policies under various auction settings.

Table 3.1: Summary of settings for studied auction formats.

	Discriminatory	GSP	Core-selecting
Demand	Unit	Unit	Multiple
Supply	Homogeneous	Heterogeneous	Heterogeneous

We first study unit-demand settings, namely the robust bidding problem in discriminatory and GSP auctions. In a perfect information setting, i.e., if there is no uncertainty about rivals' bids, the robust bidding policy is to bid the minimal amount to win one item (provided that it yields positive payoff). This policy readily extends to discriminatory auctions in the imperfect information setting: the bidder should bid minimal amount that guarantees winning one item, regardless of the realization of rivals' bids in the uncertainty set (provided positive payoff in the worst case). However such policy may not be optimal in GSP auctions under imperfect information. Due to the unit demand assumption, the bidder still wins at most one item. However, since the items are heterogeneous, the worst-case payoff from winning one item is different from the worst-case payoff from winning another item. Thus, establishing the robust bidding policy involves maximizing the minimum of item-specific

² Note that these three auctions could be considered as “natural” modifications of a VCG auction in each of the demand/supply settings.

worst-case payoff functions.

We then turn our attention to the core-selecting auction which allows for multiple bidder demand over bundles of heterogeneous items. Thus, bidder's valuation vector is multi-dimensional (exponential in the number of items) which makes the analysis of the robust bidding problem more challenging. We start our analysis by discussing conditions under which truthful bidding is optimal. In particular, the relationship between *VCG payoff* vector, i.e., the vector of payoffs that bidders would have obtained in the VCG auction, and the *core*, i.e., the set of all feasible payoff vectors in the core-selecting auction, is a determining factor for when truthful bidding is optimal. When the VCG payoff vector is in the core, the core-selecting auction is incentive compatible and hence truthful bidding is an optimal robust policy. However, when the VCG payoff vector is not in the core, the core-selecting auction is not incentive compatible, and the bidder has the incentive to bid non-truthfully. We then provide an optimal bidding policy under perfect information of rivals' bids for the general setting where the VCG payoff vector is not in the core. Under this bidding policy, the bidder underbids on the *truthful bundle*, i.e., the bundle of items that she would win if she bid truthfully, and bids on the *global bundle*, i.e., the bundle that consists of all items, so that she is guaranteed to win the truthful bundle.

For core-selecting auctions with imperfect information, we analyze the robust bidding problems of *single-minded* bidder (i.e., one with valuation for only one bundle), a *double-minded* bidder (i.e., one with distinct valuations for two inclusive bundles), and a *triple-minded* bidder (i.e., one with distinct valuations for a chain of three bundles). For the single-minded bidder, her optimal robust bidding policy is a generalization of the perfect-information optimal robust bidding policy: underbidding on her target bundle and bidding on the global bundle to guarantee winning the target bundle, regardless of the realization of rivals' bids in the uncertainty set. (We again assume that the bidder's valuation for her target bundle is high enough so that

by bidding truthfully she still wins her target bundle regardless of the realization of rivals' bids in the uncertainty set; if this condition is not true, then an optimal policy is simply bidding zero.)

For the double-minded bidder, if she wins either of her target bundles under truthful bidding and this allocation does not change for any realization of rivals' bids in the uncertainty set, then her optimal robust bidding policy is also a generalization of the perfect-information optimal bidding policy. In particular, similar to the single-minded bidder case, she underbids on the bundle that she would win under truthful bidding, and bids the highest amount on global bundle that still ensures winning the same bundle as in truthful bidding. If under truthful bidding the double-minded bidder receives different bundles for different realization of rivals' bids in the uncertainty set, then it is not clear which bundle yields larger worst-case payoff. In this case, the optimal robust policy corresponds to bids that maximize the minimum of bundle-specific worst-case payoff functions (this is similar to the case of GSP auctions with heterogeneous items).

It turns out that establishing a minimax type of an equality is critical in demonstrating and verifying the optimality of the presented robust bidding policies. However, we show the limit to this approach, as it does not extend to a triple-minded bidder in the core-selection auction. Hence, demonstrating optimality of candidate bidding policies would require a different approach and could become challenging.

Throughout our analysis, we evaluate the performance of robust bidding policies, by comparing them to a couple of benchmarks. In particular, we compare payoffs under robust bidding policies with payoffs under other bidding policies such as truthful bidding and expected-payoff maximizing policies, assuming some known distributions of rivals' bids over the uncertainty set. Our results show that for non-trivial³ robust bidding policies, the bidder's payoff is at least as large as her payoff under

³ i.e., positive bid for at least one item/bundle

truthful bidding, for all realizations of rivals' bids in the uncertainty set. In addition, for discriminatory and core-selecting auctions, robust bidding policies improve upon expected-payoff maximizing policies (assuming some known distributions of rivals' bids), as they reduce the risk of not winning the target items. In particular, robust policies yields a higher payoff compared to expected-payoff maximizing policies under adversarial realizations of rivals' bids. For GSP auctions, the comparison of the robust policy and the expected-payoff maximizing policy is more nuanced, which is due to the fact that the bidder's payment is less dependent on her own bids in this type of auction than in discriminatory and core-selecting auctions.

The paper proceeds as follows. Section 3.1.2 discusses related works in the literature. Section 3.2 follows with descriptions of the auction models and assumptions. In Section 3.3, we provide optimal bidding for unit demand bidders in discriminatory and GSP auctions. Section 3.4 discusses optimal bidding for multiple demand bidders in core-selecting auctions. Finally, in Section 3.5, we conclude and provide some possible directions for future research. Proofs are relegated to the appendix.

3.1.2 Related Literature

Our study is related to the classical topic of bidding in non-incentive-compatible (NIC) auctions (e.g., Stark and Rothkopf, 1979; Milgrom and Weber, 1982). Ausubel and Cramton (2002) studied incentive to reduce demand in discriminatory auctions with a divisible item. Edelman et al. (2005) examined GSP auctions and showed that truth-telling is not an equilibrium in such auctions. Day and Milgrom (2008) provided an optimal bid shading policy in core-selecting auctions under perfect information. Recently, Beck and Ott (2013) showed that bidder may have incentives to overbid in core-selecting auctions as well.

The modeling approach that we use lies under the framework of robust optimization. It has been recently developed as a decision tool to deal with decision making

when the input parameters are uncertain (e.g., Ben-Tal and Nemirovski, 2002; Bertsimas and Sim, 2004). Under this framework, the decision maker does not know the exact distribution of the uncertain parameters. Instead, it is assumed that these uncertain parameters belong to an uncertainty set that can be constructed based on historical data or expert’s belief. The decision maker wants to maximize her worst-case payoff with respect to this uncertainty set. There are close connections between maximizing worst-case objective over an uncertainty set and uncertainty aversion (e.g., Bertsimas and Brown, 2009). Recently, Bandi and Bertsimas (2014) adopt the robust optimization framework to study the optimal design problem for multi-item auctions, in the spirit of an earlier work by Myerson (1981). In our work, we focus on a particular bidder’s decision problem.

In some auction environments, bidders are uncertain about the distribution of rivals’ bids and are averse to this uncertainty. In such cases, worst-case maximization (or regret minimization) is a relevant decision criterion. Uncertainty aversion has been previously studied in the auction literature. Using both analytical and empirical approaches, Salo and Weber (1995) showed that overbidding behavior in first-price sealed-bid auctions can be explained by bidder’s aversion to distributional ambiguity. Lo (1998) employed the multiple priors model proposed by Gilboa and Schmeidler (1989) to model the bidder’s aversion to bid-distribution uncertainty and study the equilibrium bidding strategy in first-price sealed bid auction with private values. Bose et al. (2006) consider the optimal auction problem allowing for ambiguity about the distribution of valuation.

3.2 Model

We consider auctions that allocate m indivisible items from the set $M = \{1, 2, \dots, m\}$ among n bidders in $N = \{1, 2, \dots, n\}$. The monopolistic seller is indexed by 0. Items in M can be homogeneous (i.e., identical) or heterogeneous (i.e., distinct). A *bundle*

is a set of items. A bidder is said to have *unit demand* if she has positive valuation for only one item. Similarly, a bidder is said to have *multiple demand* if she has positive valuation for one or more bundles of items. For each bidder $j \in N$, we use $v_j(S)$ and $b_j(S)$ to denote her non-negative truthful and reported valuation for a bundle $S \subseteq M$, respectively. When bidder j has unit demand and the items are homogeneous, we abuse the notation and simply write v_j and b_j to denote her unique valuation/bid. The auctioneer uses an *allocation rule* to determine the set of items S_j to be allocated to bidder j (S_j is an empty set if bidder j does not obtain any item). Similarly, the auctioneer uses a *payment rule* to determine the amount p_j that bidder j has to pay. We assume that bidders have quasi-linear preferences, i.e., bidder j 's payoff is $\pi_j = v_j(S_j) - p_j$ for $j \in N$. For convenience, we write $b_{-j} = (b_1, b_2, \dots, b_{j-1}, b_{j+1}, \dots, b_n)$ to denote the profile of bidder j 's rivals' bids. When all b_j are scalar, the k^{th} highest bid in b_{-j} is denoted by $b_{-j}^{(k)}$. We choose to analyze the decision problem of bidder 1. It turns out that only $b_{-1}^{(k)}$ are relevant to our analysis. Thus, for simplicity, we omit the subscript and write $b^{(k)}$ to denote the k^{th} highest bid in b_{-1} . We use $\mathbb{1}_A$ to denote the indicator function of expression A , i.e., it has value one if A is true and zero otherwise. For any $x \in \mathbb{R}$, we denote $x^+ = \max(x, 0)$.

Next, we explain the allocation and payment rules in each particular auction setting and describe the robust optimization problem of bidder 1.

3.2.1 Discriminatory Auction

In discriminatory auctions, the items are assumed to be homogeneous. Each bidder has unit demand and submits a scalar bid b_j . We assume that $n \geq m$. The m highest bidders get allocated one item each and pay their bids b_j . When there is a tie, we assume that bidder 1 is favored, i.e., if $b_1 = b^{(m)}$ then bidder 1 wins an item. Thus, bidder 1 wins an item when her bid is no less than $b^{(m)}$, in which case she receives a

payoff of $v_1 - b_1$. Bidder 1's payoff function is given by:

$$\pi_1(b_1, b_{-1}) = (v_1 - b_1) \mathbb{1}_{b^{(m)} \leq b_1}. \quad (3.1)$$

3.2.2 Generalized Second-price Auction

In GSP auctions, the items are advertisement slots and are assumed to be heterogeneous. Each item $k \in M$ has an associated *click-through rate* α_k that quantifies the likelihood of the corresponding advertisement slot getting clicked. We assume that α_k are independent of bidder who obtain the corresponding items. Each bidder has unit demand and submits a scalar bid b_j . We assume that $n \geq m + 1$. The highest bidder wins item 1, the second highest bidder wins item 2, and so on. Each winning bidder pays exactly the amount of the next highest bid. Therefore, if bidder 1 wins item 1 then the payment is $p_1 = b^{(1)}$; if she wins item 2 then the payment is $p_1 = b^{(2)}$, and so on. We assume that bidder 1 is favored, so that if $b_1 = b^{(1)}$ then bidder 1 wins item 1, if $b_1 = b^{(2)}$ then bidder 1 wins item 2, and so on. If bidder 1 wins item k for $k \in M$, her (expected) payoff is the surplus $v_1 - b^{(k)}$ multiplied by the click-through rate α_k . Thus, bidder 1's profit function in this case is:

$$\pi_1(b_1, b_{-1}) = \begin{cases} \alpha_1(v_1 - b^{(1)}) & \text{if } b^{(1)} \leq b_1 \\ \alpha_2(v_1 - b^{(2)}) & \text{if } b^{(2)} \leq b_1 < b^{(1)} \\ \vdots & \\ \alpha_m(v_1 - b^{(m)}) & \text{if } b^{(m)} \leq b_1 < b^{(m-1)} \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

3.2.3 Core-selecting Auction

In our analysis of core-selecting auctions, we assume the items are heterogeneous. Each bidder has multiple demand and submits bids $b_j(S)$ for bundles $S \subseteq M$ of her interest. When $S = M$, we refer to it as the *global bundle*. The auctioneer decides the allocation outcome by solving the *winner determination problem*:

$$\begin{aligned}
w_b(N, M) &= \max_x \sum_{j \in N} \sum_{S \subseteq M} b_j(S) \cdot x_j(S) \\
\text{s.t.} \quad &\sum_{S \supseteq \{i\}} \sum_{j \in N} x_j(S) \leq 1, \quad \forall i \in M, \\
&\sum_{S \subseteq M} x_j(S) \leq 1, \quad \forall j \in N, \\
&x_j(S) \in \{0, 1\}, \quad \forall (S, j) \text{ s.t. bid } b_j(S) \text{ was submitted.}
\end{aligned} \tag{WD}$$

In (WD), the auctioneer maximizes the total reported value of bundles. The first constraint ensures that each item is only allocated once. The second constraint implies that the auctioneer only accepts at most one submitted bid from a bidder.⁴ Thus, each binary variable $x_j(S)$ equals to one if and only if bidder j is awarded bundle $S \subseteq M$. It is possible that the problem (WD) has multiple optimal solutions. In such cases, we assume a tie-breaking rule in favor of bidder 1 winning a pre-specified bundle. The function $w_b(C, S)$ is referred to as the *coalition value function* and is defined as the maximum surplus generated by allocating an arbitrary set of items $S \subseteq M$ among bidders in the set $C \subseteq N$ given the reported bids b . When $S = M$, we abuse the notation and write $w_b(C)$ instead of $w_b(C, M)$.

Next, we define *the core* to be the set of non-negative payoff vectors $\{\pi_j\}_{j \in N \cup 0}$ satisfying the *core constraints*:

$$\sum_{j \in C \cup 0} \pi_j \geq w_b(C), \quad \forall C \subseteq N. \tag{3.3}$$

The right hand side of (3.3) is the maximum surplus generated by allocating the items among members in the coalition C (hereafter referred to as a *blocking coalition*). Thus, the constraints (3.3) guarantee that no group of bidders can ensure better

⁴ This bidding rule is referred to as the ‘‘XOR’’ bidding language and is commonly used in practice for spectrum auctions.

payoff for themselves by excluding others from participation, i.e., there is no incentive to form a blocking coalition in the auction. We can re-write (3.3) to be constraints on payments as follows. First, recall that S_j is the allocated bundle for bidder j . By substituting $\pi_0 = \sum_{j \in N} p_j$ and $\pi_j = b_j(S_j) - p_j$, we get

$$\sum_{j \in W} p_j \geq w_b(C) - \sum_{j \in C} (b_j(S_j) - p_j), \quad \forall C \subseteq N, \quad (3.4)$$

where W is the set of bidders who receives nonempty bundles. After rearranging, the above constraints become:

$$\sum_{j \in W \setminus C} p_j \geq w_b(C) - \sum_{j \in C} b_j(S_j), \quad \forall C \subseteq N. \quad (3.5)$$

Let $\beta_C = w_b(C) - \sum_{j \in C} b_j(S_j)$ and $\beta \in \mathbb{R}^{2^n}$ be the vector of all β_C 's. Also, let A be a $n \times 2^n$ matrix comprising columns a_C that has the j th entry equals to zero if bidder j is in coalition C and one otherwise. Then the core constraints (3.5) can be written succinctly as

$$pA \geq \beta. \quad (3.6)$$

A *core-selecting payment rule* is a payment rule that selects a vector $\{p_j\}_{j \in N}$ satisfying the core constraint (3.6) and the individual rationality constraint $p \leq b$. For specificity, we use the quadratic core-selecting payment rule proposed by Day and Cramton (2012).⁵ Let p^0 be a reference payment vector. Under this rule, the payment vector p is the optimal solution of the following quadratic program:

⁵ Note that our results can be extended to other core-selecting payment rules as well. For other core-selecting payment rules, see e.g., Ausubel and Baranov (2010).

$$\begin{aligned}
& \min_p (p - p^0)(p - p^0)^T \\
& \text{s.t. } pA \geq \beta \\
& \quad p \leq b \\
& \quad p1 = \mu,
\end{aligned} \tag{QP}$$

where μ is defined as

$$\begin{aligned}
\mu &= \min_p p1 \\
& \text{s.t. } pA \geq \beta \\
& \quad p \leq b.
\end{aligned} \tag{3.7}$$

The payment vector p determined by (QP) minimizes the Euclidean distance from reference payment vector p^0 to the core. The quantity μ is the minimum value of total payment from bidders. Thus, the constraint $p1 = \mu$ guarantees that the payment rule is *bidder-optimal*, i.e., the total payment from bidders is minimized. This has the effect of minimizing the bidders' total incentive to deviate, as shown by Day and Milgrom (2008).

A related payment rule is the *VCG rule* (Vickrey, 1961; Clarke, 1971; Groves, 1973). Under this payment rule, each winner j pays the opportunity cost that she imposes on other bidders:

$$p_j^{VCG} = w_b(N \setminus j, M) - w_b(N \setminus j, M \setminus S_j). \tag{3.8}$$

The *VCG payoff* of bidder j is then simply $\pi_j^{VCG} = v_j(S_j) - p_j^{VCG}$. It is known that VCG payment rule satisfies incentive compatibility but may result in low seller's revenue and is vulnerable to collusive bidding Ausubel and Milgrom (2002). In contrast, core-selecting rules are robust to collusion but may have incentive compatibility issues, as we will see in later sections. When the quadratic rule (QP) is used with p^0

being the vector of VCG payments, we refer to such rule as the *nearest-VCG rule*.

3.2.4 Robust Optimization Formulation

In each of the aforementioned auction formats, we consider the decision problem of a particular bidder, who we choose to be bidder 1, without loss of generality. We assume that bidder 1 has a belief that her rivals' bids belong to an *uncertainty set* \mathcal{U}_{-1} . For discriminatory and GSP auctions, we have $\mathcal{U}_{-1} \subset \mathbb{R}_+^{n-1}$, while for core-selecting auctions, we have $\mathcal{U}_{-1} \subset \mathbb{R}_+^{(n-1)2^m}$. For simplicity, we assume that the uncertainty set \mathcal{U}_{-1} is a convex polytope, i.e., it can be specified by some linear constraints on b_{-1} :

$$\mathcal{U}_{-1} = \{b_{-1} \mid Pb \leq q\}, \quad (3.9)$$

where P and q are matrix and vector with appropriate dimensions. For example, \mathcal{U}_{-1} can be a box set (i.e., a hyperrectangle) that corresponds to the case where bids b_j for $j \neq 1$ are independent and belong to some known intervals $[\underline{b}_j, \bar{b}_j]$:

$$\mathcal{U}_{-1} = \{b_{-1} \mid \underline{b}_j \leq b_j \leq \bar{b}_j, \forall j \neq 1\}. \quad (3.10)$$

Despite the convex polytope assumption, many of our results can be extended to more general types of uncertainty set.⁶ Given an uncertainty set \mathcal{U}_{-1} , bidder 1's objective is to maximize her worst-case payoff with respect to this uncertainty set. In other words, she needs to solve the following robust optimization problem:

$$\pi_1^{MAXMIN} = \sup_{b_1 \in \mathcal{U}_1} \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1, b_{-1}). \quad (\text{RO})$$

In the above problem, we use \mathcal{U}_1 to refer to bidder 1's feasible policy space. We assume no restrictions, except for the non-negative condition, so $\mathcal{U}_1 = \mathbb{R}_+$ for dis-

⁶ Bandi and Bertsimas (2014) discussed several ways to construct the uncertainty set based on historical data.

criminatory and GSP auctions and $\mathcal{U}_1 = \mathbb{R}_+^{2^m}$ for core-selecting auctions. We call bidding policy b_1 a *robust policy* if it is an optimal solution to (RO).

Let π_1^{MINMAX} be the minimum value over \mathcal{U}_{-1} of bidder 1's maximum *ex post* payoff, i.e., payoff that she can achieve using an ex post optimal policy. Thus, π_1^{MINMAX} is given by:

$$\pi_1^{MINMAX} = \inf_{b_{-1} \in \mathcal{U}_{-1}} \sup_{b_1 \in \mathcal{U}_1} \pi_1(b_1, b_{-1}). \quad (3.11)$$

By minimax inequality (e.g., Boyd and Vandenberghe, 2004), one has that

$$\pi_1^{MAXMIN} \leq \pi_1^{MINMAX}. \quad (3.12)$$

Thus, π_1^{MINMAX} provides an upper bound for the optimal objective of the robust optimization problem (RO). If (3.12) holds with equality, i.e.,

$$\pi_1^{MAXMIN} = \pi_1^{MINMAX}, \quad (3.13)$$

then there are two implications. First, the robust policy can be viewed as an ex post optimal policy applied for a particular worst-case bid $b_{-1} \in \mathcal{U}_{-1}$. Second, if under a policy $b_1 \in \mathcal{U}_1$, bidder 1's worst-case payoff is the same with π_1^{MINMAX} , then such policy must be a robust policy. The later observation is particularly useful in proving the optimality of bidding policies in core-selecting auctions, as we will see in Section 3.4.

3.3 Robust Bidding with Unit Demand

In this section, we study the robust bidding problem (RO) in discriminatory and GSP auctions. We assume that bidders have unit demand in these auctions. For discriminatory auctions, the items are homogeneous, whereas for GSP auctions we consider both homogeneous and heterogeneous cases. We first introduce additional

notation. Recall that $b^{(k)}$ is the k^{th} highest bid among rivals of bidder 1. Let $l(k)$, $u(k)$ be the minimum and maximum values of $b^{(k)}$ over the uncertainty set \mathcal{U}_{-1} , i.e.,

$$\begin{aligned} l(k) &= \min_{b_{-1} \in \mathcal{U}_{-1}} b^{(k)} \\ \text{s.t. } & b_{-1} \in \mathcal{U}_{-1}, \end{aligned} \tag{3.14}$$

$$\begin{aligned} u(k) &= \max_{b_{-1} \in \mathcal{U}_{-1}} b^{(k)} \\ \text{s.t. } & b_{-1} \in \mathcal{U}_{-1}. \end{aligned} \tag{3.15}$$

In addition, let $u(k_1 | k_2, x)$ be the supremum of $b^{(k_1)}$ conditional on $b^{(k_2)} < x$:

$$\begin{aligned} u(k_1 | k_2, x) &= \sup_{\substack{b_{-1} \in \mathcal{U}_{-1} \\ x < b^{(k_2)}}} b^{(k_1)} \\ \text{s.t. } & b_{-1} \in \mathcal{U}_{-1} \\ & x < b^{(k_2)}. \end{aligned} \tag{3.16}$$

When $x = u(k_2)$ in (3.16), instead of writing $u(k_1 | k_2, u(k_2))$, we will use $u(k_1 | k_2)$ as a short notation. For any $x_1 < x_2$, any feasible solution to (3.16) when $x = x_2$ is also feasible to (3.16) when $x = x_1$. As a result, $u(k_1 | k_2, x)$ is non-increasing in x .

3.3.1 Discriminatory Auction

In a discriminatory auction, if bidder 1 knows her rivals' bids, she can best respond by bidding exactly $b^{(m)}$, the m^{th} highest bid among her rivals, if $b^{(m)} \leq v_1$. The following proposition characterizes bidder 1's robust policy under imperfect information.

Proposition 3.1. *In discriminatory auctions, a robust policy for bidder 1 is $b_1^{\text{RO}} = u(m)\mathbb{1}_{u(m) \leq v_1}$. The optimal payoff is $\pi_1^{\text{MAXMIN}} = (v_1 - u(m))^+$.*

Since truthful bidding always gives zero payoff, robust bidding results in a strictly better worst-case payoff when $u(m) < v_1$. Next, we provide a comparison between robust bidding policy and expected-payoff maximizing policy under uniform distribution assumption.

Example 3.1. Consider a discriminatory auction with $n = 2$ bidders and $m = 1$ item. In this case, the discriminatory auction is a first-price auction. Bidder 1 has a belief that $b_2 \in [c, d]$ for some non-negative constants c and d . According to Proposition 3.1, bidder 1's robust policy is $b_1^{RO} = d\mathbb{1}_{d \leq v_1}$. If bidder 1's belief is such that b_2 is uniformly distributed in $[c, d]$ and her objective is to maximize the expected payoff, the optimal bid is:⁷

$$b_1^{EM} = \begin{cases} d & \text{if } 2d - c < v_1 \\ \frac{1}{2}(v_1 + c) & \text{if } c \leq v_1 \leq 2d - c \\ 0 & \text{otherwise.} \end{cases}$$

Figure 3.1 shows a plot of b_1^{RO} and b_1^{EM} for the case where $c = 0$ and $d = 5$. There are two main differences between these bidding policies. First, when $v_1 < 5$, bidder 1 bids zero under b_1^{RO} , whereas under b_1^{EM} , she bids a positive amount. Second, when $v_1 \geq 5$, we have $b_1^{RO} = 5$ while $b_1^{EM} = \frac{1}{2}v_1$ is an increasing function in v_1 . When $v_1 \geq 10$, the two bidding policies are the same. In Figure 3.2, we show a comparison of bidder 1's payoffs (as functions of b_2) under b_1^{RO} and b_1^{EM} for different values of v_1 . For $v_1 = 3$, since $v_1 < 5$, according to Proposition 3.1, bidder 1's robust policy is $b_1^{RO} = 0$. As a result, her payoff is $\pi_1^{RO} = 0$ for all realization of $b_2 \in [0, 5]$. On the other hand, to maximize her expected payoff, bidder 1 bids $b_1^{EM} = 1.5$. Thus, she still gains a positive payoff of $\pi_1^{EM} = 1.5$ for $b_2 \in [0, 1.5]$. For $v_1 = 8$, we have $b_1^{RO} = 5$ and $b_1^{EM} = 4$. If $b_2 \in [0, 4)$ then bidder 1's payoffs under b_1^{RO} and b_1^{EM} are 3 and 4, respectively. However, when $b_2 \in [4, 5]$, bidder 1 does not win the item if she bids b_1^{EM} so her corresponding payoff is $\pi_1^{EM} = 0$, while her payoff when bidding b_1^{RO} is still $\pi_1^{RO} = 3$. \square

Example 3.1 illustrates the fact that under robust policy b_1^{RO} , bidder 1 bids the minimal amount that guarantees winning an item, if such bidding policy is profitable.

⁷ See Appendix B.1

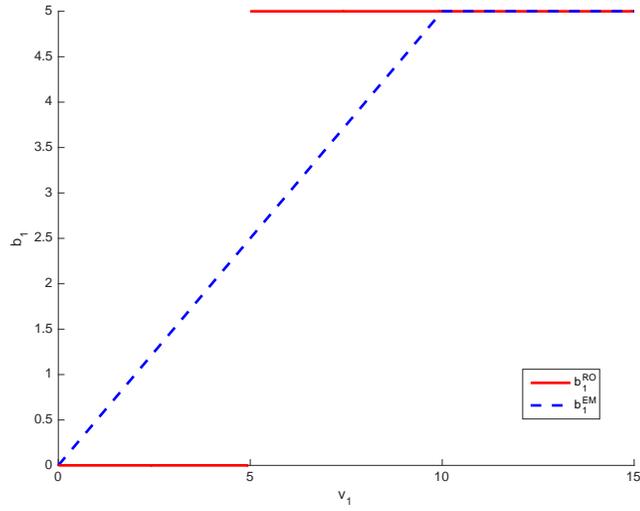


FIGURE 3.1: Illustration of Example 3.1 – Comparison of b_1^{RO} and b_1^{EM} .

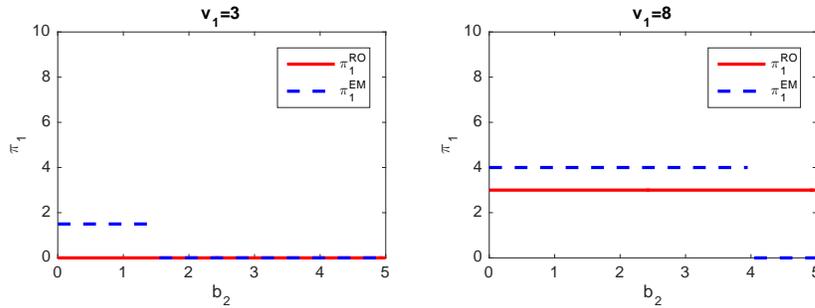


FIGURE 3.2: Illustration of Example 3.1 – Bidder 1's payoff function under b_1^{RO} and b_1^{EM} .

As we can see, this bidding policy is beneficial when rivals bid adversarially. Similar observations can be made for bidding policies in multiple demand settings, as we will see in Section 3.4.

3.3.2 Generalized Second-price Auction

In a GSP auction, if bidder 1 has perfect information about b_{-1} , her optimal policy is to bid the minimal amount to win the most profitable item. Specifically, let k^* be

the most profitable item:

$$k^* = \operatorname{argmax}_{k \in M} (\alpha_k (v_1 - b^{(k)})).$$

If $b^{(k^*)} \leq v_1$ then bidder 1's optimal policy is to bid $b_1 = b^{(k^*)}$. On the other hand, if $v_1 < b^{(k^*)}$ then winning any slot will result in negative payoff, so it is optimal for bidder 1 to bid $b_1 = 0$.

Now let us consider the case where bidder 1 has imperfect information about b_{-1} and has a belief that $b_{-1} \in \mathcal{U}_{-1}$. For each $k \in M$, let $f_k(x)$ be the worst-case payoff that bidder 1 gets when she wins item k by bidding x . Clearly, for a fixed k , the function $f_k(x)$ is only defined for $x \in [l(k), u(k-1))$. The following lemma characterizes the structure of $f_k(\cdot)$.

Lemma 3.1. *For $k \in \{1, 2, \dots, m\}$,*

$$f_k(x) = \begin{cases} \alpha_k (v_1 - x) & \text{if } l(k) \leq x < u(k) \\ \alpha_k (v_1 - u(k)) & \text{if } u(k) \leq x < u(k-1 | k) \\ \alpha_k (v_1 - u(k | k-1, x)) & \text{if } u(k-1 | k) \leq x < u(k-1) \end{cases}$$

From Lemma 3.1, we can see that $f_k(x)$ is decreasing on $[l(k), u(k))$, constant on $[u(k), u(k-1 | k))$ and weakly increasing on $[u(k-1 | k), u(k-1))$ (see Figure 3.3). Next, we look at bidder 1's robust policy in GSP auctions. We start with the case where the click-through rates are the same and are equal to α .

Proposition 3.2. *In GSP auctions with $\alpha_k = \alpha$ for all $k \in M$, a robust policy for bidder 1 is $b_1^{RO} = u(m) \mathbb{1}_{u(m) \leq v_1}$ and the worst-case payoff is $\pi_1^{MAXMIN} = \alpha(v_1 - u(m))^+$.*

Figure 3.4 shows the worst-case profit function of bidder 1 in a GSP auction with $n = 3$ bidders and $m = 2$ items with identical click-through rates. As we can see, it is optimal for bidder 1 to bid $b_1^{RO} = u(2) \mathbb{1}_{u(2) \leq v_1}$. Notice that if bidder 1 bids lower than $u(2)$, it is possible that she does not obtain any item, so her worst-case payoff

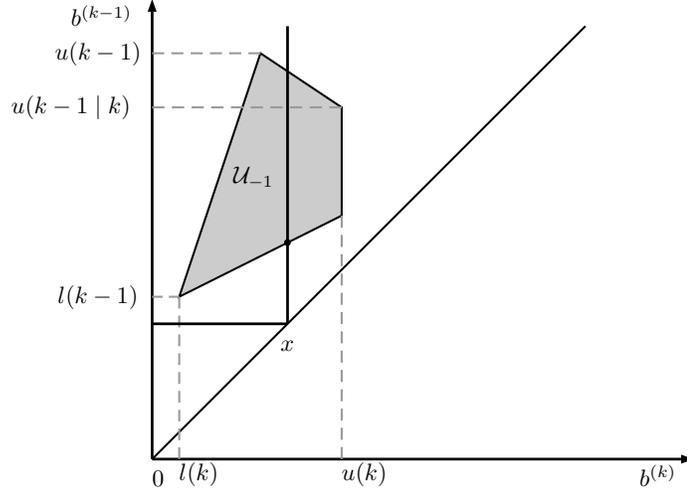


FIGURE 3.3: Illustration of Lemma 3.1 – The supremum of $b^{(k)}$ is increasing in x on $[l(k)u(k))$, constant on $[u(k), u(k-1 | k))$ and decreasing in x on $[u(k-1 | k), u(k-1))$.

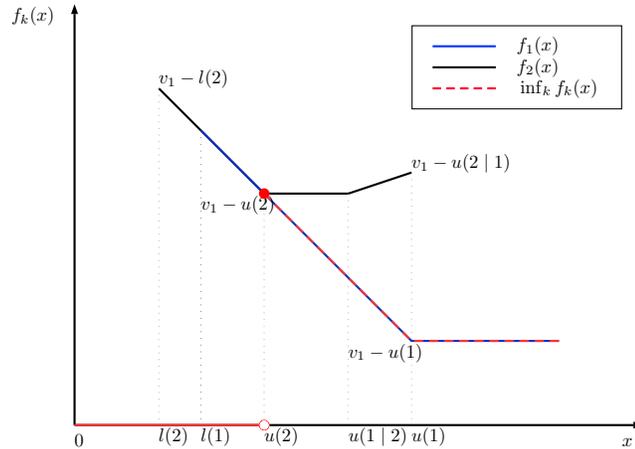


FIGURE 3.4: Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have identical click-through rates.

is zero. On the other hand, by bidding $b_1 > u(2)$, bidder 1's worst-case payoff is the same with $f_1(b_1)$, her payoff when winning item 1. Since this payoff is decreasing in b_1 (by Lemma 3.1), bidder 1's worst-case payoff is less than her worst-case payoff when bidding b_1^{RO} .

Next, we look at the case where click-through rates are not identical. We first

analyze case of $n = 3$ bidders and $m = 2$ items, and then discuss the extensions to more general settings.

Proposition 3.3. *Consider a GSP auction with $n = 3$ bidders and $m = 2$ items with click-through rates $\alpha_2 < \alpha_1$. A robust policy for bidder 1 is*

$$b_1^{RO} = \begin{cases} u(2) \mathbb{1}_{u(2) \leq v_1} & \text{if } \alpha_1(v_1 - u(1 | 2)) \leq \alpha_2(v_1 - u(2)) \\ u(1) \mathbb{1}_{u(1) \leq v_1} & \text{if } \alpha_2(v_1 - u(2 | 1)) \leq \alpha_1(v_1 - u(1)) \\ x^* \mathbb{1}_{x^* \leq v_1} & \text{otherwise} \end{cases} \quad (3.17)$$

where x^* is the solution of the following equation

$$\alpha_1(v_1 - x) = \alpha_2(v_1 - u(2 | 1, x)). \quad (3.18)$$

The worst-case payoff is given by

$$\pi_1^{MAXMIN} = \begin{cases} \alpha_2(v_1 - u(2))^+ & \text{if } \alpha_1(v_1 - u(1 | 2)) \leq \alpha_2(v_1 - u(2)) \\ \alpha_1(v_1 - u(1))^+ & \text{if } \alpha_2(v_1 - u(2 | 1)) \leq \alpha_1(v_1 - u(1)) \\ \alpha_1(v_1 - x^*)^+ & \text{otherwise} \end{cases} \quad (3.19)$$

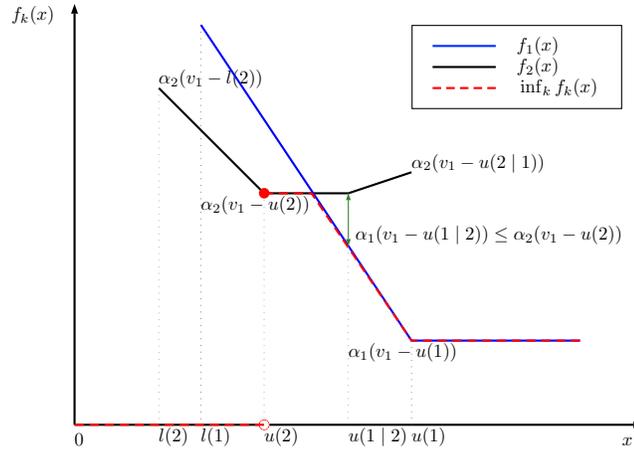


FIGURE 3.5: Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have different click-through rates satisfying $\alpha_1(v_1 - u(1 | 2)) \leq \alpha_2(v_1 - u(2))$.

Figure 3.5-3.7 show the worst-case profit function of bidder 1 under different conditions of the click-through rates α_1 and α_2 as described in (3.17). If $\alpha_1(v_1 - u(1 |$

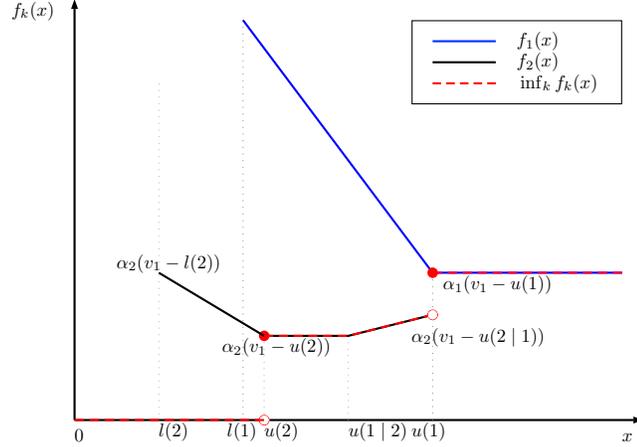


FIGURE 3.6: Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have different click-through rates satisfying $\alpha_2(v_1 - u(2 | 1)) \leq \alpha_1(v_1 - u(1))$.

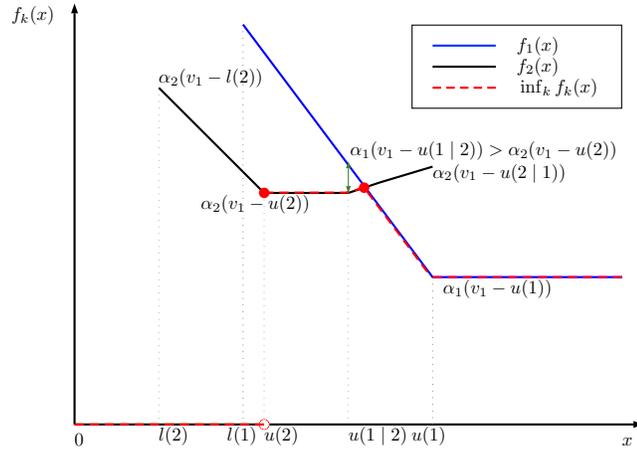


FIGURE 3.7: Worst-case profit function in a GSP auction with $n = 3$ bidders and $m = 2$ items which have different click-through rates satisfying $\alpha_2(v_1 - u(2)) < \alpha_1(v_1 - u(1 | 2))$ and $\alpha_1(v_1 - u(1)) < \alpha_2(v_1 - u(2 | 1))$.

$2)) \leq \alpha_2(v_1 - u(2))$ then from Figure 3.5 we see that, similar to the case of identical click-through rates, bidding $b_1^{RO} = u(2)\mathbb{1}_{u(2) \leq v_1}$ is optimal. If $\alpha_2(v_1 - u(2 | 1)) \leq \alpha_1(v_1 - u(1))$ then $b_1^{RO} = u(1)\mathbb{1}_{u(1) \leq v_1}$. Notice that if $v_1 > u(1)$ in this case then bidder 1 does not want to bid lower than $u(1)$, since she would risk winning item 2 or not winning any item, and receive a lower worst-case payoff. Finally, Figure 3.7

shows the worst-case payoff function when $\alpha_2(v_1 - u(2)) < \alpha_1(v_1 - u(1 | 2))$ and $\alpha_1(v_1 - u(1)) < \alpha_2(v_1 - u(2 | 1))$. In this case, the robust policy is to bid at the intersection of two worst-case payoff functions $f_1(x)$ and $f_2(x)$.

Next, in the following examples, we compare the robust policy with truthful bidding and an expected-payoff maximizing policy.

Example 3.2. Consider a GSP auction with $n = 3$ bidders and $m = 2$ items. Let $\alpha = (1, 0.7)$, $v_1 = 10$ and $\mathcal{U}_{-1} = \{(b^{(1)}, b^{(2)}) \mid 4 \leq b^{(2)} \leq 5, 6 \leq b^{(1)} \leq 7\}$. Since $\alpha_1(v_1 - u(1 | 2)) = 3 < 3.5 = \alpha_2(v_1 - u(2))$, according to Proposition 3.3, a robust policy for bidder 1 is $b_1^{RO} = 5$. The worst-case payoff under this policy is $\pi_1^{MAXMIN} = 3.5$, while the worst-case payoff under truthful bidding is $\pi_1^{TR} = 3 < 3.5 = \pi_1^{MAXMIN}$. One can also verify that in this case bidder 1's payoff under b_1^{RO} (as a function of b_{-1}) is larger than her payoff under truthful bidding for any realization of $b_{-1} \in \mathcal{U}_{-1}$. \square

Example 3.3. Consider a GSP auction with $n = 3$ bidders and $m = 2$ items. The settings for click-through rates and uncertainty set is similar to those in Example 3.2. However, in this case, we allow v_1 to vary and examine the performance of robust policy b_1^{RO} and expected-payoff maximizing policy b_1^{EM} corresponding to the distributional assumption that $b^{(1)}$ and $b^{(2)}$ are independent and uniformly distributed on $[6, 7]$ and $[4, 5]$, respectively. From Proposition 3.3, we have

$$b_1^{RO} = \begin{cases} 0 & \text{if } v_1 < 5 \\ 5 & \text{if } 5 \leq v_1 < \frac{35}{3} \\ 7 & \text{otherwise} \end{cases} \quad (3.20)$$

It can be shown that (see Appendix B.1) if $b^{(1)}$ and $b^{(2)}$ are independent and uniformly distributed, the expected-payoff maximizing policy is

$$b_1^{EM} = \begin{cases} 0 & \text{if } v_1 < 4 \\ v_1 & \text{if } 4 \leq v_1 < 5 \\ 5 & \text{if } 5 \leq v_1 < \frac{201}{18} \\ 7 & \text{otherwise} \end{cases} \quad (3.21)$$

Thus, b_1^{RO} and b_1^{EM} are the same except for when $v_1 \in [4, 5)$ and $v_1 \in [\frac{201}{18}, \frac{35}{3})$. Figures 3.8 and 3.9 show bidder 1's payoff functions under these policies for $v_1 = 4.5$ and $v_1 = 11.5$, respectively. When $v_1 = 4.5$, $b_1^{RO} = 0$ so bidder 1 receives $\pi_1^{RO} = 0$ for all realization of $b_{-1} \in \mathcal{U}_{-1}$. On the other hand, $b_1^{EM} = 4.5$ when $v_1 = 4.5$, so bidder 1 wins the second item and receives positive payoff when $b^{(2)} < 4.5$. For $v_1 = 11.5$, we have $b_1^{RO} = 5$ while $b_1^{EM} = 7$. From Figure 3.9, we see that under b_1^{RO} bidder 1's payoff is smallest when $b_1^{(2)}$ is highest, while under b_1^{EM} her payoff is smallest when $b_1^{(1)}$ is highest. b_1^{RO} gives slightly better worst-case payoff than b_1^{EM} , but it can yield smaller payoff at some other realizations of $b_{-1} \in \mathcal{U}_{-1}$. \square

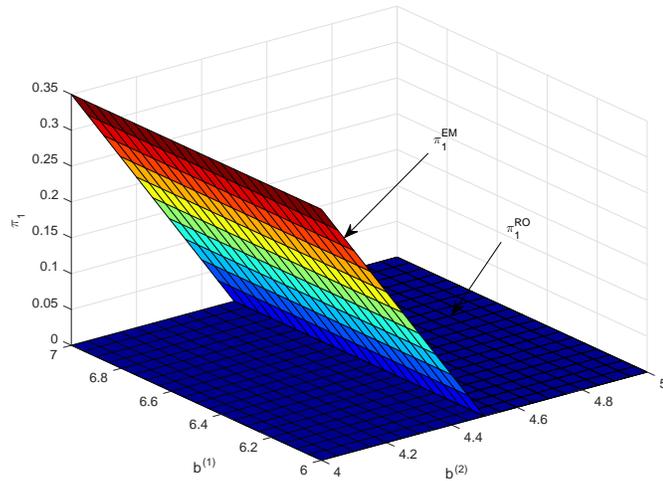


FIGURE 3.8: Illustration of Example 3.3 – Bidder 1's payoff function under b_1^{RO} and b_1^{EM} when $v_1 = 4.5$.

The result of Proposition 3.3 extends directly to the case of $n \geq 4$ and $m = 2$ since the robust bidding problem only depends on bidder 1's belief on $b^{(1)}$ and $b^{(2)}$.

Corollary 3.1. *Consider a GSP auction with $n \geq 4$ bidders and $m = 2$ items. A robust policy for bidder 1 is b_1^{RO} given in (3.17).*

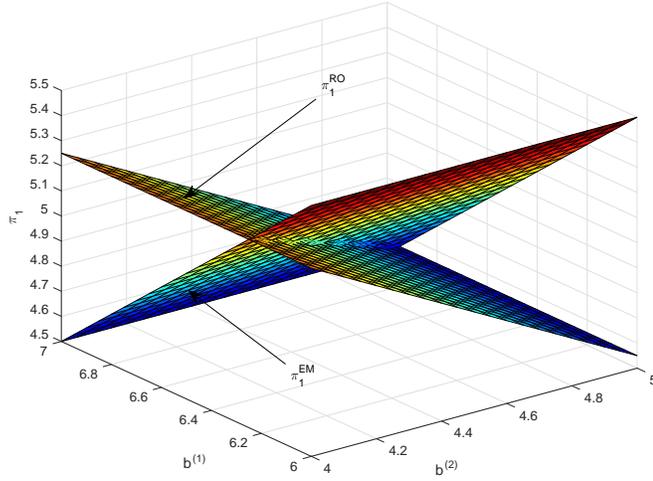


FIGURE 3.9: Illustration of Example 3.3 – Bidder 1’s payoff function under b_1^{RO} and b_1^{EM} when $v_1 = 11.5$.

Robust policies in more general settings with $n \geq 4$ bidders and $m \geq 3$ items with heterogeneous click-through rates is more involved as one needs to examine closely the relationship between $m \geq 3$ worst-case payoff functions. However, robust policies can be obtained in a similar fashion as in Proposition 3.3. Specifically, for all $k \in M$, we compute the k^{th} item’s worst-case payoff functions $f_k(x)$ and its corresponding domain $[l(k), u(k - 1))$. Next, we take the minimum of all these payoff functions to get the overall worst-case payoff. Finally, we decide the robust policy by choosing b_1^{RO} that maximizes this worst-case payoff function.

3.4 Robust Bidding with Multiple Demand

We now consider the robust bidding problem (RO) in core-selecting auctions. We assume that bidders in these auctions have multiple demand and report bids on bundles of items. We start our analysis with a discussion of the incentive to *misreport*, i.e., bidding different from true valuation, in core-selecting auctions. Bidder 1’s

optimal bidding policies under perfect information are discussed next. Finally, we study robust policies when bidder 1 has imperfect information of rivals' bids.

3.4.1 Relationship between VCG and the Core

The relationship between VCG payoffs and the core provides us useful information about whether bidders have incentive to misreport in core-selecting auctions. We say *VCG is in the core* if VCG payoffs satisfies core constraints (3.3) or, equivalently, the VCG payments satisfy core constraints (3.5). If VCG is in the core then it is the unique bidder-optimal point in the core and no bidder has the incentive to misreport.⁸ A sufficient condition for VCG being in the core is the *bidder-submodularity* property of the coalition value function $w_b(\cdot)$, defined as follows.

Definition 3.1. The coalition value function $w_b(\cdot)$ is *bidder-submodular* if for all $j \in N$ and all coalitions C and C' satisfying $0 \in C \subseteq C'$, one has that $w_b(C \cup \{j\}) - w_b(C) \geq w_b(C' \cup \{j\}) - w_b(C')$.

The bidder-submodularity property of coalition value function is closely related to the submodularity and supermodularity properties of bid functions $b_j(S)$. Thus, it is useful to review these definitions.

Definition 3.2. A set function $f(\cdot) : 2^M \rightarrow \mathbb{R}$ is said to be *submodular* if for every $S, S' \subseteq M$ with $S \subseteq S'$ and every $i \in M$ we have that

$$f(S \cup \{i\}) - f(S) \geq f(S' \cup \{i\}) - f(S'). \quad (3.22)$$

Definition 3.3. A set function $f(\cdot) : 2^M \rightarrow \mathbb{R}$ is said to be *supermodular* if for every $S, S' \subseteq M$ with $S \subseteq S'$ and every $i \in M$ we have that

$$f(S \cup \{i\}) - f(S) \leq f(S' \cup \{i\}) - f(S'). \quad (3.23)$$

⁸ See Ausubel and Milgrom (2002).

If $b_j(\cdot)$ is submodular for some $j \in N$ then bidder j 's reported valuation for having an extra item $i \in M$ decreases as S increases. Therefore, all goods are *substitutes* for bidder j (with respect to her reported valuation). Similarly, if $b_j(\cdot)$ is supermodular then all goods are *complements* for bidder j (with respect to her reported valuation). The following result establishes the relationship between the bidder-submodularity property of coalition valuation and the submodularity property of bid functions.

Proposition 3.4 (Ausubel and Milgrom (2006)). *If $b_j(\cdot)$ is a submodular set function for all $j \in N$, the corresponding coalition value function $w_b(\cdot)$ is bidder-submodular and VCG is in the core.*

Interestingly, if all bid functions are supermodular, the coalition value function is also bidder-submodular and thus VCG is also in the core. The following proposition establishes this result.

Proposition 3.5. *If $b_j(\cdot)$ is a supermodular set function for all $j \in N$, the corresponding coalition value function $w_b(\cdot)$ is bidder-submodular and VCG is in the core.*

When goods are substitutes, the marginal value of an extra bidder to a coalition is the difference between that new bidder's value for his assigned bundle and the opportunity cost of the coalition for the bundle. Since this opportunity cost increases with the coalition size under substitution, the marginal value of an extra bidder decreases as coalition size increases, which is the definition for bidder-submodular. When goods are complements, the marginal value of an extra bidder is the difference between his valuation for all goods and the coalition valuation for all goods, if such difference is positive. Thus, as the coalition size increases, the valuation of the coalition for all goods increases so the marginal value of the extra bidder decreases, which establishes the bidder-submodularity property of the coalition value function.

As we can see, bidders have incentive to misreport only when there are goods that are substitutes for some bidders and (possibly different) goods that are complements for (possibly different) bidders. These results extend to imperfect information as well.

Proposition 3.6. *If the uncertainty set \mathcal{U}_{-1} is such that for all $b_{-1} \in \mathcal{U}_{-1}$ the coalition value function $w_{v_1, b_{-1}}(\cdot)$ is bidder-submodular then truthful bidding is the optimal solution of (RO).*

Corollary 3.2. *If v_1 is submodular (supermodular) and $\{b_j\}_{j \neq 1}$ are submodular (supermodular) for all $b_{-1} \in \mathcal{U}_{-1}$, then truthful reporting is the optimal robust policy.*

3.4.2 Optimal Bidding Policies under Perfect Information

If bidder 1 has perfect information about rivals' bids b_{-1} , her VCG payoff is the maximum payoff that she can achieve.⁹ There are multiple policies that bidder 1 can use to achieve such payoff. In the next proposition, we show one such policy, which has direct generalization to imperfect information case.¹⁰ For convenience, we assume a tie-breaking rule in which bidder 1 winning bundle S_1 , the bundle that bidder 1 would win if she bids truthfully (hereafter referred to as *truthful bundle*), is favored.

Proposition 3.7. *The following policy is optimal for bidder 1:*

$$b_1^{PI}(S) = \begin{cases} 0 & \text{if } S \subsetneq S_1 \\ v_1(S_1) - \pi_1^{VCG} & \text{if } S_1 \subseteq S \subsetneq M \\ w_{b_{-1}}(N \setminus 1) & \text{if } S = M \end{cases} \quad (3.24)$$

⁹ See Day and Milgrom (2008).

¹⁰ For other optimal policies under perfect information, see Day and Milgrom (2008) and Beck and Ott (2013).

Remark 3.1. Policy (3.24) is also optimal for bidder 1 even if she uses multiple identities, also known as *shills*.¹¹

Under policy (3.24), bidder 1 shades (underbids) her valuation on truthful bundle S_1 and on any other bundles that contain it, except the global bundle M . In addition, she misreports her valuation for the global bundle and bids $w_v(N \setminus 1)$ for it. This bidding policy has two main effects. First, it guarantees bidder 1 still wins bundle S_1 under the assumed tie-breaking rule. Second, by bidding high on the global bundle, bidder 1 effectively inflates her demand for other bundles. As a result, other bidders are forced to pay for the high opportunity cost they impose on bidder 1. Since a winner's payment decreases if other winners' payments increase under a bidder-optimal payment rule, bidder 1 payment decreases. Notice that policy (3.24) can be modified to accommodate for different tie-breaking rules. Specifically, if bidder 1 winning bundle S_1 is not favored, bidder 1 can change her bid for S where $S_1 \subseteq S \subsetneq M$ to $b_1(S) = v_1(S_1) - \pi_1^{VCG} + \epsilon$ for some $\epsilon > 0$ arbitrarily small. In such case, bidder 1 still wins S_1 and achieves a payoff that is arbitrarily close to her VCG payoff.

We next show an existence result for perfect-information equilibria with non-truthful bidding strategies. First, we introduce a notation to represent an achievable payoff vector in core-selecting auctions.

Definition 3.4. For every valuation profile v , an *imputation* γ is a payoff vector that satisfies the rationality constraint, the core constraint and the efficiency constraint.

In other words, the imputation γ satisfies

$$\gamma \geq 0, \quad \sum_{j \in C} \gamma_j \geq w_v(C) \quad \forall C \subseteq N, \quad \sum_{j \in N} \gamma_j = w_v(N).$$

¹¹ For more discussion on shill bidding in combinatorial auctions, see e.g., Ausubel and Milgrom (2006)

The following proposition establishes a perfect-information equilibrium result in core-selecting auctions, assuming the tie-breaking rule is such that each bidder j winning her truthful bundle S_j is favored. Notice that we can accommodate for other tie-breaking rules and still obtain similar results by appropriately adding/subtracting $\epsilon > 0$ to the bidding functions, as we did in Proposition 3.7.

Proposition 3.8. *For a given valuation profile v , let γ^{EQ} be a corresponding imputation and let $\{S_j\}_{j \in N}$ be the allocation when everyone reports truthfully. The following bid profile*

$$b_j(S) = \begin{cases} 0 & \text{if } S \subsetneq S_1 \\ v_j(S_j) - \gamma_j^{EQ} & \text{if } S_1 \subseteq S \subsetneq M \\ \gamma_0^{EQ} & \text{if } S = M \end{cases} \quad (3.25)$$

is a full information equilibrium profile. Furthermore, $\{S_j\}_{j \in N}$ is an equilibrium allocation and the equilibrium payoff is γ^{EQ} .

As we can see in Proposition 3.8, one can vary the imputation γ^{EQ} to obtain different perfect-information equilibrium bidding profiles. The bidding policies in these profiles are in general non-truthful. In particular, they may include both underbidding and overbidding.

To conclude this section, we show that when deviating from truthful bidding is profitable, bidding truthfully is not even a locally optimal policy. In other words, there is a reporting policy that is arbitrarily close to bidder 1's true valuation that ensures bidder 1 a higher payoff than her truthful payoff.

Proposition 3.9. *Assume that the nearest-VCG payment rule is used. If bidder 1's truthful payoff is strictly less than her VCG payoff, then for any $\epsilon > 0$ there exists a bidding policy $b_1^\epsilon(\cdot)$ such that $\|v_1 - b_1^\epsilon\|_2 \leq \epsilon$ and such that by reporting $b_1^\epsilon(\cdot)$ bidder 1 gets a payoff that is larger than her truthful payoff.*

3.4.3 Single-minded Bidder

In this section, we consider the bidding problem of bidder 1 who has imperfect information about her rivals' bids. We assume that bidder 1 is *single-minded*, i.e., she has positive valuation if she wins a particular bundle S_1 and zero valuation otherwise:¹²

$$\begin{aligned} v_1(S) &= a > 0 & \text{if } S \supseteq S_1, \\ v_1(S) &= 0 & \text{if } S \not\supseteq S_1. \end{aligned} \tag{3.26}$$

Let ξ be the maximum value over the uncertainty set \mathcal{U}_{-1} of bidder 1's VCG payment if she wins S_1 , i.e.,

$$\xi = \max_{b_{-1} \in \mathcal{U}_{-1}} w_{b_{-1}}(N \setminus 1, M) - w_{b_{-1}}(N \setminus 1, M \setminus S_1). \tag{3.27}$$

If $v_1(S_1) \leq \xi$ then there exists $b_{-1}^* \in \mathcal{U}_{-1}$ such that bidder 1 does not win S_1 by bidding truthfully. Bidder 1's payoff is thus zero in that case. According to Proposition 3.7, there exists an optimal bidding policy that yields the same allocation outcome for bidder 1. Thus, given b_{-1}^* , bidder 1's payoff is still at most zero if she bids non-truthfully. Therefore, zero is an upperbound on bidder 1's worst-case payoff and bidding zero is a trivial robust policy. Consequently, without loss of generality, we assume throughout the remainder of our analysis in this section that $v_1(S_1) > \xi$, i.e., bidder 1's valuation for bundle S_1 is high enough so that she always wins bundle S_1 by bidding truthfully, regardless of the realization of her rivals' bids $b_{-1} \in \mathcal{U}_{-1}$. For convenience, we also assume a tie-breaking rule such that bidder 1 winning S_1 is favored. Notice that this assumption does not affect our results, by similar reason as in perfect-information case. The following proposition gives a robust policy for bidder 1.

¹² Here we assume free disposal, so bidder 1's valuation for winning $S \supseteq S_1$ is the same with her valuation for winning S_1

Proposition 3.10. *If bidder 1 is single-minded and $v_1(S_1) > \xi$ then a robust policy for bidder 1 is*

$$\begin{aligned} b_1^{RO}(S) &= \xi, \quad \text{if } S_1 \subseteq S \subsetneq M, \\ b_1^{RO}(M) &= \xi + \min_{b_{-1} \in \mathcal{U}_{-1}} w_{b_{-1}}(N \setminus 1, M \setminus S_1), \\ b_1^{RO}(S) &= 0, \quad \text{otherwise.} \end{aligned} \tag{3.28}$$

The optimal worst-case payoff is $\pi_1^{MAXMIN} = v_1(S_1) - \xi$.

Remark 3.2. Policy (3.28) is not uniquely optimal. For example, if $b_1(M)$ is any value in the interval $[0, \xi + \min_{b_{-1} \in \mathcal{U}_{-1}} w_{b_{-1}}(N \setminus 1, M \setminus S_1)]$, then the resulting policy is also optimal. However, bidding a higher value of $b_1(M)$ weakly increases the payoff of bidder 1 for any realization of $b_{-1} \in \mathcal{U}_{-1}$, since it would increase the payment of other bidders and thus reduce bidder 1's payment.

L \setminus L \setminus G Valuation Structure

We analyze the performance of robust policy given by Proposition 3.10 in an auction with $n = 3$ bidders and $m = 2$ identical items. In particular, we consider the Local-Local-Global (L \setminus L \setminus G) valuation structure in which bidder 1 is a *local* bidder who is interested in only one item while bidders 2 and 3 are *local* and *global* bidders who are interested in winning one and two items, respectively.¹³ Table 3.2 summarizes notation: a, b, c are problem parameters while x and y are decision variables. Note that in this particular section, we abuse the notation and use b to refer to bidder 2's bid on one and two items, rather than the bid profile of all bidders.

We consider the simple box-type uncertainty set:

$$\mathcal{U}_{-1} = \{(b_2, b_3) \mid \bar{b} - \epsilon_b \leq b \leq \bar{b} + \epsilon_b, \bar{c} - \epsilon_c \leq c \leq \bar{c} + \epsilon_c\},$$

¹³ Similar valuation structure has been used in equilibrium analysis of core-selecting auctions (e.g., Goeree and Lien, 2016)

Table 3.2: Bidder valuations under the L\L\G valuation structure.

# items	v_1	b_1	b_2	b_3
1	a	x	b	0
2	a	y	b	c

where $\epsilon_b < \bar{b}$ and $\epsilon_c < \bar{c}$. Notice that in this setting the coalition value function $w_{v_1, b_{-1}}(\cdot)$ need not to be bidder-submodular. Bidder 1's worst-case VCG payment defined in (3.27) in this case is simply $\xi = \max_{b_{-1} \in \mathcal{U}_{-1}} (c-b)^+$. If $\xi < a$ then according to Proposition 3.10, the robust bidding policy is

$$(x^*, y^*) = (\xi, \xi + \bar{b} - \epsilon_b) \quad (3.29)$$

Remark 3.3. Given the nearest-VCG payment rule, bidder 1's payoff under robust policy (3.29) is greater than her payoff under truthful bidding for any realization of b_{-1} in \mathcal{U}_{-1} .

Next, we provide concrete numerical examples to further illustrate the performance of robust bidding policy within the L\L\G structure.

Example 3.4. Consider a core-selecting auction with $n = 3$ bidders and $m = 2$ items. Bidders have L\L\G valuation structure (see Table 3.2). Let $v_1 = (10, 10)$, and $\mathcal{U}_{-1} = \{(b_2, b_3) \mid 7 \leq b \leq 13, 7 \leq c \leq 13\}$. According to Proposition 3.10, a robust bidding policy is $b_1^{RO,1} = (6, 13)$. However, as noted earlier in Remark 3.2, bidding any value of $y \in [6, 13]$ will not change the worst-case payoff, so policy such as $b_1^{RO,2} = (6, 6)$ also yields the same worst-case payoff. For a concrete example of rivals' bids defined by $b = 10$ and $c = 10$, perfect-information policy (3.24) is $b_1^{PI} = (0, 10)$.

Figure 3.10 shows a comparison of bidder 1's payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and truthful policy $b_1^{TR} = (10, 10)$. We can see that for any realization of $b_1 \in \mathcal{U}_{-1}$, bidder 1 receives a larger payoff by bidding according to robust policy $b_1^{RO,1}$

instead of reporting her true valuation. This agrees with our earlier observation in Remark 3.3.

Figure 3.11 shows a comparison of robust policies $b_1^{RO,1}$ and $b_1^{RO,2}$. Note that in this case the worst-case realization of the rivals' bids is $b = 7$ and $c = 13$, i.e., bidder 2 bids lowest and bidder 3 bids highest. Both policies $b_1^{RO,1}$ and $b_1^{RO,2}$ yield the same worst-case payoff. However, policy $b_1^{RO,1}$ gives larger payoff at other realizations of b_{-1} in the uncertainty set since bidder 1 bids higher on the global bundle in this policy (recall Remark 3.2).

In Figure 3.12, we show a comparison between robust policy $b_1^{RO,1}$ and perfect-information policy b_1^{PI} . Under b_1^{PI} , if bidder 2 and 3 bid such that $b \geq c$ and $b \geq 10$, then bidder 1 and 2 win an item each. Since bidder 1 bids zero for one item, bidder 2 has to incur the entire payment burden and bidder 1 is a free rider. Thus, bidder 1 extracts the entire surplus and obtains maximum payoff in this case. However, by bidding b_1^{PI} , bidder 1 also faces the risks of winning the unnecessary extra item or not winning any item. Specifically, if bidder 2 and 3 bid $b < c$ and $c \geq 10$, then bidder 3 wins both items and bidder 1 receives zero payoff. Similarly, if bidder 2 and 3 bid $b < 10$ and $c < 10$ then bidder 1 wins both item and pays a high payment $p_1 = \max(b, c)$. In both cases, bidder 1's payoff is reduced significantly. Robust policy $b_1^{RO,1}$ avoids these risks by ensuring that bidder 1 always wins an item, regardless of how much bidder 2 and 3 bid. \square

Remark 3.4. In Example 3.4, if bidder 1 uses a perfect-information optimal policy (3.24) with respect to the worst-case rivals' bids $b = 7$ and $c = 13$ then she will recover robust policy $b_1^{RO,1} = (6, 13)$. This is because in the single-minded setting, minimax equality (3.13) holds.

Example 3.5. Consider a core-selecting auction with similar settings as in Example 3.4. However, in this case, we allow bidder 1's valuation to vary and compare bidder

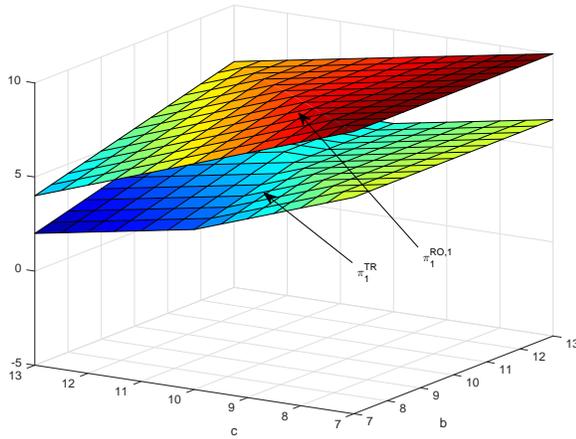


FIGURE 3.10: Illustration of Example 3.4 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and truthful policy $b_1^{TR} = (10, 10)$.

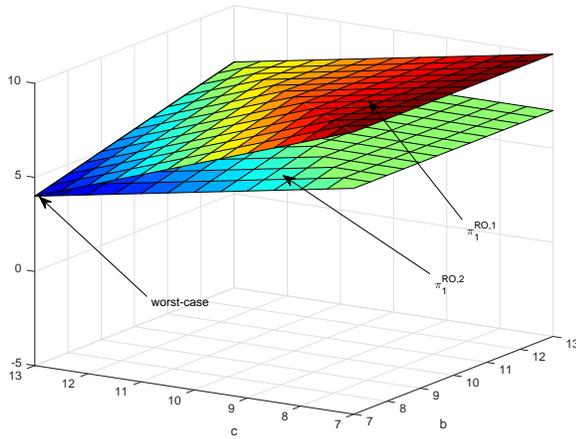


FIGURE 3.11: Illustration of Example 3.4 – Comparison of bidder 1’s payoffs under robust policies $b_1^{RO,1} = (6, 13)$ and $b_1^{RO,2} = (6, 6)$.

1’s payoff under robust policy (3.29) to her payoff under an expected-payoff maximizing policy. To compute the later, we consider the case where bidder 1 has a belief that her rivals’ bids are such that b and c are independent and the corresponding probability density functions are:

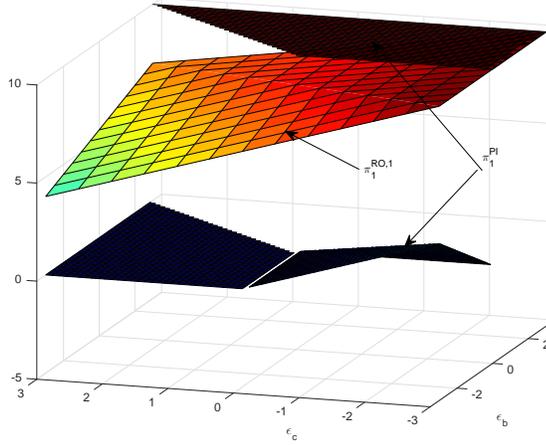


FIGURE 3.12: Illustration of Example 3.4 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and perfect-information policy $b_1^{PI} = (0, 10)$.

$$f_b(b) = \begin{cases} \frac{1}{18}(b - 7) & \text{if } 7 \leq b \leq 13 \\ 0 & \text{otherwise} \end{cases} \quad (3.30)$$

and

$$f_c(c) = \begin{cases} \frac{1}{18}(13 - c) & \text{if } 7 \leq c \leq 13 \\ 0 & \text{otherwise.} \end{cases} \quad (3.31)$$

Note that if b and c are uniformly distributed then it turns out that bidder 1’s expected-payoff maximizing policy is the same with her robust policy (3.29) and the comparison is trivial. Thus, for illustration purpose, we consider the distributions f_b and f_c instead.

When $v_1 = (4, 4)$, bidder 1’s robust policy and expected-payoff maximizing policies are $b_1^{RO} = b_1^{EM} = (0, 0)$. Thus, bidder 1’s payoff is the same under the two policies. When $v_1 = (10, 10)$, we have $b_1^{RO} = (6, 13)$ and $b_1^{EM} = (4, 12)$ (see Appendix B.2 for details). Notice that f_b is increasing in b for $b \in [7, 13]$ and f_c is decreasing in c for $c \in [7, 13]$, so under these distributional assumptions, bidder 2 are more likely to bid high and bidder 3 are more likely to bid low, relative to the

support ranges given by \mathcal{U}_{-1} . As a result, under b_1^{EM} , bidder 1 bids low for one item in anticipation that she would win one item with high probability and pay less as a result of her low bid. However, this policy exposes her to the risks of not winning any item or winning both items. Figure 3.13 shows a comparison of bidder 1's payoff under b_1^{RO} and b_1^{EM} when $v_1 = (10, 10)$. As we can see, bidder 1's payoff under b_1^{RO} is significantly greater than her payoff under b_1^{EM} at realizations of rivals' bids in \mathcal{U}_{-1} such that bidder 1 wins both items or does not win any items by bidding b_1^{EM} . Figure 3.14 shows a comparison of bidder 1's payoff under b_1^{RO} and b_1^{EM} when $v_1 = (20, 20)$. Similar observations can be made regarding the payoff functions under the two policies. Note that b_1^{RO} is the same as in the case of $v_1 = (10, 10)$ since the policy only depends on parameters of \mathcal{U}_{-1} as long as $a > \xi$. On the other hand, $b_1^{EM} = (3, 11)$ when $v_1 = (20, 20)$, i.e., bidder 1 shades more on her bids compared to the case of $v_1 = (10, 10)$. □

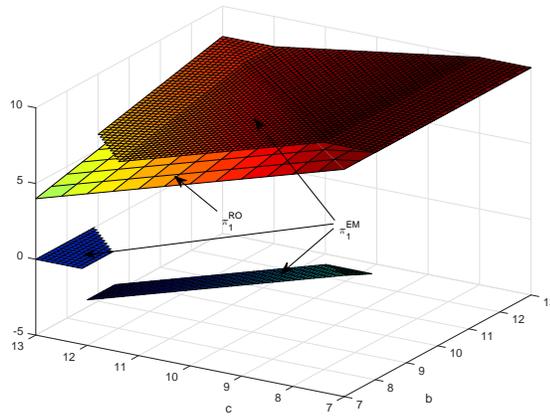


FIGURE 3.13: Illustration of Example 3.5 – Comparison of bidder 1's payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and expected-payoff maximizing policy $b_1^{EM} = (4, 12)$ when $v_1 = (10, 10)$.

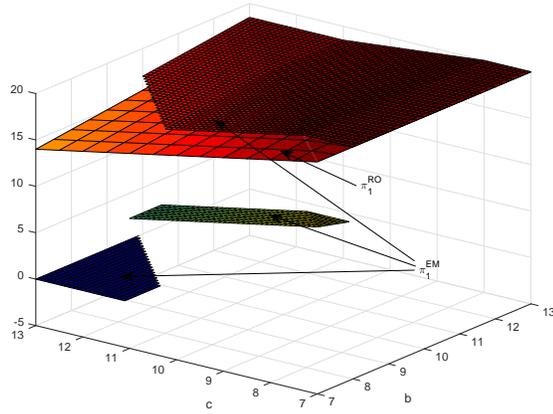


FIGURE 3.14: Illustration of Example 3.5 – Comparison of bidder 1’s payoffs under robust policy $b_1^{RO,1} = (6, 13)$ and expected-payoff maximizing policy $b_1^{EM} = (3, 11)$ when $v_1 = (20, 20)$.

3.4.4 Double-minded Bidder

We now extend our analysis to the case where bidder 1 is *double-minded*, i.e., she has positive valuation for two bundles. In particular, we assume that bidder 1 is interested in two bundles S_1 and S_2 where $\emptyset \subsetneq S_1 \subsetneq S_2 \subseteq M$:

$$\begin{aligned}
 v_1(S) &= a > 0 && \text{if } S_1 \subseteq S \subsetneq S_2, \\
 v_1(S) &= a' \geq a && \text{if } S_2 \subseteq S, \\
 v_1(S) &= 0 && \text{otherwise.}
 \end{aligned} \tag{3.32}$$

Let

$$\xi_1 = \max_{b_{-1} \in \mathcal{U}_{-1}} (w_{b_{-1}}(N \setminus 1, M) - w_{b_{-1}}(N \setminus 1, M \setminus S_1)) \tag{3.33}$$

and

$$\xi_2 = \max_{b_{-1} \in \mathcal{U}_{-1}} (w_{b_{-1}}(N \setminus 1, M) - w_{b_{-1}}(N \setminus 1, M \setminus S_2)) \tag{3.34}$$

be the maximum values of bidder 1’s VCG payment over the uncertainty set \mathcal{U}_{-1} when she wins bundle S_1 and S_2 , respectively. As in the single-minded bidder case,

if $v_1(S_1) < \xi_1$ then bidder 1's worst-case payoff is bounded above by zero, so bidding zero is a robust policy. Thus, without loss of generality, we assume that $v_1(S_1) \geq \xi_1$, i.e., bidder 1's valuation for S_1 is high enough so that she always wins either S_1 or S_2 under truthful reporting.

A bundle is said to be the *unique truthful allocation* for bidder 1 if by bidding truthfully bidder 1 always wins that bundle regardless of the realization of her rivals' bids in the uncertainty set. The following proposition shows robust policies for bidder 1 if either S_1 or S_2 is her unique truthful allocation.

Proposition 3.11. *Let bidder 1 be double-minded with $0 < \xi_1 < v_1(S_1) \leq v_1(S_2)$:*

(a) *If S_1 is bidder 1's unique truthful allocation then a robust policy is*

$$\begin{aligned} b_1^{RO}(S) &= \xi_1 \quad \text{if } S_1 \subseteq S \subsetneq M, \\ b_1^{RO}(M) &= \xi_1 + \min_{b_{-1} \in \mathcal{U}_{-1}} w_{b_{-1}}(N \setminus 1, M \setminus S_1), \\ b_1^{RO}(S) &= 0 \quad \text{otherwise.} \end{aligned} \tag{3.35}$$

The optimal worst-case payoff is $\pi_1^{MAXMIN} = v_1(S_1) - \xi_1$.

(b) *If S_2 is bidder 1's unique truthful allocation then a robust policy is*

$$\begin{aligned} b_1^{RO}(S) &= \xi_2 \quad \text{if } S_2 \subseteq S \subsetneq M, \\ b_1^{RO}(M) &= \xi_2 + \min_{b_{-1} \in \mathcal{U}_{-1}} w_{b_{-1}}(N \setminus 1, M \setminus S_2), \\ b_1^{RO}(S) &= 0 \quad \text{otherwise.} \end{aligned} \tag{3.36}$$

The optimal worst-case payoff is $\pi_1^{MAXMIN} = v_1(S_2) - \xi_2$.

Remark 3.5. Note that robust policies (3.35) and (3.36) are similar to robust policy (3.28) in the single-minded case. In fact, results from Proposition 3.11 can be extended to the case where bidder 1 has positive valuation for a collection of bundles

$\{S_k\}_{k=1}^K$ that satisfies $\emptyset \subsetneq S_1 \subsetneq S_2 \subsetneq \dots \subsetneq S_K$. In such case, if S_k is the unique truthful allocation for bidder 1 then a robust policy is

$$\begin{aligned} b_1^{RO}(S) &= \xi_k \quad \text{if } S_k \subseteq S \subsetneq M \\ b_1^{RO}(M) &= \xi_k + \min_{b_{-1} \in \mathcal{U}_{-1}} w_{b_{-1}}(N \setminus 1, M \setminus S_k) \\ b_1^{RO}(S) &= 0 \quad \text{otherwise.} \end{aligned}$$

We now examine the case where neither S_1 nor S_2 is the unique truthful allocation of bidder 1. For analytical tractability, we study this situation in an example auction setting similar to that of the single-minded case.

LG\L\G Valuation Structure

Consider a core-selecting auction with $n = 3$ bidders and $m = 2$ homogeneous items. We extend the setting of Section 3.4.3 by making bidder 1 double-minded, i.e., bidder 1 is now interested in winning either one or two items. We call this valuation structure the LG\L\G structure. Table 3.3 summarizes the notation for this case (notice that bidder 1's valuation for both items is now $a' \geq a$ instead of just a as in the L\L\G case).

Table 3.3: Bidder valuations under the LG\L\G valuation structure.

# items	v_1	b_1	b_2	b_3
1	a	x	b	0
2	a'	y	b	c

The quantities ξ_1 and ξ_2 defined in (3.33) and (3.34) specialized for this case are:

$$\begin{aligned} \xi_1 &= \max_{b_{-1} \in \mathcal{U}_{-1}} (c - b)^+, \\ \xi_2 &= \max_{b_{-1} \in \mathcal{U}_{-1}} \max(b, c). \end{aligned}$$

We assume that $\xi_1 < a$, so bidder 1 wins either one or two items under truthful bidding. According to Proposition 3.11, if $a' \leq a + \bar{b} - \epsilon_b$ then since winning one item is the unique truthful allocation for bidder 1, her robust bidding policy is $b_1^{RO} = (\xi_1, \xi_1 + \bar{b} - \epsilon_b)$. Similarly, if $a + \bar{b} + \epsilon_b < a'$ then winning both items is bidder 1's unique truthful allocation. As a result, her robust bidding policy is $b_1^{RO} = (0, \xi_2)$. The optimal worst-case payoff is thus $\pi_1^{MAXMIN} = a - \xi_1$ if $a' \leq a + \bar{b} - \epsilon_b$ and $\pi_1^{MAXMIN} = a' - \xi_2$ if $a + \bar{b} + \epsilon_b < a'$.

Now consider the case when $a' \in (a + \bar{b} - \epsilon_b, a + \bar{b} + \epsilon_b]$. In this case, bidder 1 can win either one item or both items under truthful bidding. To find the robust bidding policy for bidder 1, we first restrict our policy space to the set $\mathcal{U}'_1 = \{(x, y) \in \mathbb{R}_+^2 \mid x = \xi_1, y \geq x\}$. We will show that by choosing the optimal value for y in this reduced policy space \mathcal{U}'_1 , bidder 1 can obtain the highest possible worst-case payoff given by minimax inequality, so such policy is also optimal in the original policy space $\mathcal{U}_1 = \mathbb{R}_+^2$.

For any $(x, y) \in \mathcal{U}'_1$, let $\pi_1^{WO,1}(y)$ and $\pi_1^{WO,2}(y)$ be the worst-case payoff function when bidder 1 bids (x, y) and wins one and two items, respectively. We have

$$\inf_{b_{-1} \in \mathcal{U}'_{-1}} \pi_1(b_1, b_{-1}) = \min \left(\pi_1^{WO,1}(y), \pi_1^{WO,2}(y) \right)$$

The following lemma provides some properties of $\pi_1^{WO,1}(y)$ and $\pi_1^{WO,2}(y)$.

Lemma 3.2. *The worst-case payoff functions $\pi_1^{WO,1}(y)$ and $\pi_1^{WO,2}(y)$ satisfy:*

(a) $\pi_1^{WO,1}(y)$ is piece-wise linear and increasing in y . Furthermore,

$$\pi_1^{WO,1}(y) = a \quad \text{if } y \geq \xi_1 + \bar{c} + \epsilon_c.$$

(b) $\pi_1^{WO,2}(y)$ is piece-wise linear and decreasing in y . Furthermore,

$$\pi_1^{WO,2}(y) = a' - \bar{c} - \epsilon_c \quad \text{if } y \leq \xi_1 + \bar{c} + \epsilon_c.$$

In the next lemma, we show the relationship between these worst-case payoff functions and bidders' valuations.

Lemma 3.3. *We have:*

$$(a) \pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b) \leq \pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b) \text{ if and only if } a' \leq a + \bar{b} - \epsilon_b,$$

$$(b) \pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b) \leq \pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b) \text{ if and only if } a' \leq a + \bar{b} + \epsilon_b.$$

Figure 3.15 shows these worst-case payoff function under different valuation scenarios. When either winning one item or winning two items is the unique truthful allocation for bidder 1, the worst-case payoff function corresponding to that unique allocation dominates the other (Figures 3.15a and 3.15b). When this is not the case, the two worst-case payoff functions intersect (Figures 3.15c and 3.15d).

We now characterize of the robust bidding policy under the $LG \setminus L \setminus G$ setting when $a' \in (a + \bar{b} - \epsilon_b, a + \bar{b} + \epsilon_b]$.

Proposition 3.12. *Consider $LG \setminus L \setminus G$ setting. If $a' \in (a + \bar{b} - \epsilon_b, a + \bar{b} + \epsilon_b]$ then a robust bidding policy for bidder 1 is $b_1^{RO} = (x^*, y^*)$ with $x^* = \xi_1$ and y^* is a solution of the equation $\pi_1^{WO,1}(y) = \pi_1^{WO,2}(y)$. The worst-case payoff is $\pi_1^{MAXMIN} = \min(a, a' - \bar{c} - \epsilon_c)$.*

Remark 3.6. It turns out that in this $LG \setminus L \setminus G$ setting bidding truthfully is also a robust policy as long as there exists $b_{-1} \in \mathcal{U}_{-1}$ such that bidder 1 wins the global bundle under truthful bidding, i.e., $a + \bar{b} - \epsilon_b < a'$ (see Appendix B.2). This may not be true in general when S_2 is not the global bundle.

3.4.5 Beyond Double-minded Bidders: the Role of Minimax Equality

Demonstrating that the minimax equality (3.13) holds was central to our establishing of a robust bidding policy for the double-minded bidder (Proposition 3.12). Thus,

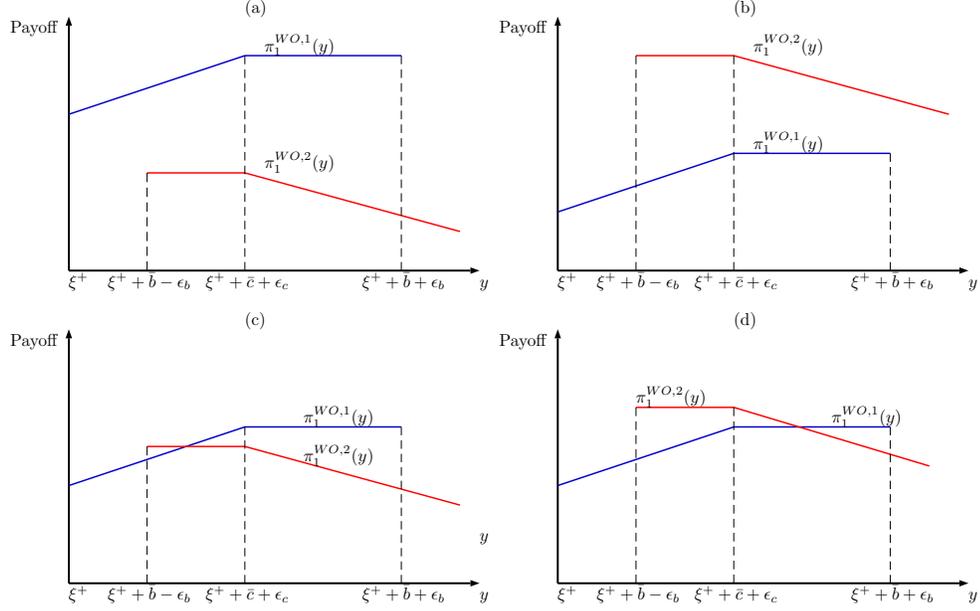


FIGURE 3.15: Worst-case payoff functions in different cases under LG\L\G setting. (a) $a' \leq a + \bar{b} - \epsilon_b$, winning one item always yields better worst-case payoff. (b) $a + \bar{b} + \epsilon_b < a'$, winning two items always yields better worst-case payoff. (c) and (d) $a + \bar{b} - \epsilon_b < a' \leq a + \bar{b} + \epsilon_b$, two worst-case payoff functions intersect.

identifying other settings in which minimax equality (3.13) holds would indicate a possibility for our approach to yield a robust bidding policy in such settings. For example, note that if there exists a unique truthful allocation for bidder 1, one can readily obtain minimax equality (3.13). Indeed, when there is a unique truthful allocation, bidder 1's ex post optimal policy is to bid the minimal amount to win this allocation. Since there is only one allocation outcome for the bidder, the resulting worst-case payoff is the same with that of robust bidding policies, so minimax equality (3.13) must hold. The following proposition summarizes this for the three auction formats we analyzed.

Proposition 3.13. *Minimax equality (3.13) holds in:*

- (a) *Discriminatory auctions,*
- (b) *GSP auctions with identical click-through-rates,*

(c) *Core-selecting auctions with single-minded bidders.*

In contrast, when bidder 1 receives different allocations for different realizations of rivals' bids in the uncertainty set under truthful bidding, minimax equality (3.13) may not hold due to the non-convexity of level sets associated with bidder 1's payoff functions. We next show that Sion's (1958) level-set based sufficient conditions for the minimax equality might not hold. Specifically, for any $\lambda \in \mathbb{R}$, the lower level set and upper level set of $f(x, y)$ are defined as follows:

$$\begin{aligned} LE(x, \lambda) &= \{y : y \in Y, f(x, y) \leq \lambda\}, \\ GE(\lambda, y) &= \{x : x \in X, f(x, y) \geq \lambda\}. \end{aligned}$$

Let X be a convex subset of a linear topological space, Y be a compact convex subset of a linear topological space, and $f : X \times Y \rightarrow \mathbb{R}$ be upper semicontinuous on X and lower semicontinuous on Y . Suppose that

$$GE(\lambda, y) \text{ is convex for all } y \in Y \text{ and } \lambda \in \mathbb{R}, \quad (3.37)$$

and,

$$LE(x, \lambda) \text{ is convex for all } x \in X \text{ and } \lambda \in \mathbb{R}. \quad (3.38)$$

Then we have Sion (1958):

$$\min_Y \sup_X f = \sup_X \min_Y f. \quad (3.39)$$

The following example shows that condition (3.37) on the convexity of upper level sets in b_1 of bidder 1's payoff function $\pi_1(b_1, b_{-1})$ is violated in a core-selecting auction with multiple demand settings.

Example 3.6. Consider a core-selecting auction with $n = 3$ bidders and $m = 3$ homogeneous items. Bidder 1's valuation vector is $v_1 = (5, 15, 16.5)$, bidder 2's bid vector is $b_2 = (5, 8.5, 14.5)$ and bidder 3's bid vector is $b_3 = (3, 6, 12)$. Figure 3.16

shows the upper level set in b_1 of bidder 1's payoff function $\pi_1(b_1, b_{-1})$ corresponding to level $\lambda = 1$. In this case, the upper level set is not connected, so condition (3.37) is violated. □

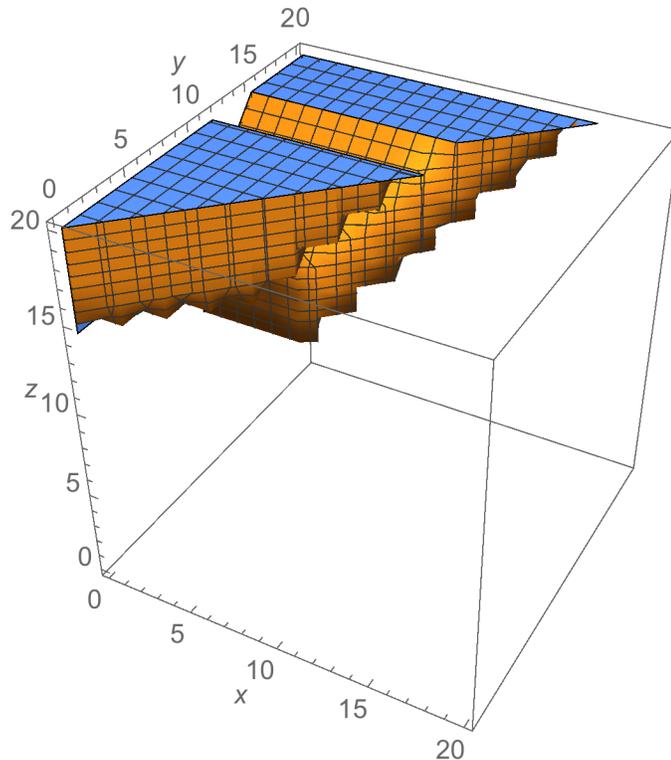


FIGURE 3.16: Illustration of Example 3.6: Non-convexity of the upper level set of in b_1 of $\pi_1(b_1, b_{-1})$ with $\lambda = 1$.

Note that the non-convexity of level sets of the bidder's payoff function does not exclude the possibility of minimax equality (3.13) holding. However, in the next example, we show that minimax equality (3.13) does not hold in a core-selecting auction where bidder 1 is *triple-minded*. (Recall that the bidder is triple minded if she has three distinct positive valuations $v(S_1) < v(S_2) < v(S_3)$ for some bundles $S_1 \subset S_2 \subset S_3$.)

Example 3.7. Consider a core-selecting auction with $n = 3$ bidders and $m = 3$ homogeneous items. Bidder 1 has valuation for one, two and three items. Bidder 2

has valuation for one and two items, but does not have an extra value for winning the third item. Bidder 3 is only interested in winning all items and has zero valuation for either one or two items. Table 3.4 summarizes the bidder valuation structure.

Table 3.4: Bidder valuations, Example 3.7.

# items	v_1	b_1	b_2	b_3
1	a	x	b	0
2	a'	y	c	0
3	a''	z	c	d

We assume that the uncertainty set \mathcal{U}_{-1} is of box-type and is given by:

$$\mathcal{U}_{-1} = \{(b_2, b_3) \mid \bar{b} - \epsilon_b \leq b \leq \bar{b} + \epsilon_b, \bar{c} - \epsilon_c \leq c \leq \bar{c} + \epsilon_c, \bar{d} - \epsilon_d \leq d \leq \bar{d} + \epsilon_d\}.$$

The numerical values we consider here are $v_1 = (a, a', a'') = (7, 13, 13.4)$, $\bar{b}_{-1} = (\bar{b}, \bar{c}, \bar{d}) = (4, 8, 10)'$, $\epsilon = (\epsilon_b, \epsilon_c, \epsilon_d) = 0.23 \bar{b}_{-1}$. It is straightforward to establish that, with these parameters, bidder 1 is guaranteed to win at least one item if she bids her true valuations (and consequently bidder 3 never wins). Bidder 1's worst-case profit functions corresponding to winning one, two and three items, respectively, are:

$$\begin{aligned} \pi_1^{WO,1}(x, y, z) &= \min \left(a, a - \frac{1}{2}x - \bar{d} - \epsilon_d + \frac{1}{2}c^*(x, y, z) + \frac{1}{2} \max(\bar{d} + \epsilon_d, z) \right), \\ \pi_1^{WO,2}(x, y, z) &= \min \left(a', a' - \bar{d} - \epsilon_d + \frac{1}{2}b^*(x, y, z) - \frac{1}{2}y + \frac{1}{2} \max(\bar{d} + \epsilon_d, z) \right), \\ \pi_1^{WO,3}(x, y, z) &= a'' - \bar{d} - \epsilon_d, \end{aligned}$$

where b^* and c^* are given by:

$$\begin{aligned} b^*(x, y, z) &= \max(\bar{b} - \epsilon_b, x + \bar{c} - \epsilon_c - y, z - y), \\ c^*(x, y, z) &= \max(\bar{c} - \epsilon_c, \max(y + \bar{b} - \epsilon_b - x, z - x)). \end{aligned}$$

The feasible region for $\pi_1^{WO,i}(x, y, z)$ is $\mathcal{U}_1^{(i)}$, for $i \in \{1, 2, 3\}$. These feasible regions

are given by:

$$\begin{aligned}
\mathcal{U}_1^{(1)} &= \{(x, y, z) \in \mathbb{R}_+^3 \mid x \geq \bar{d} - \bar{c} + \epsilon_c + \epsilon_d, \\
&\quad x + \bar{c} + \epsilon_c - \bar{b} + \epsilon_b \geq y, x + \bar{c} + \epsilon_c \geq z\}, \\
\mathcal{U}_1^{(2)} &= \{(x, y, z) \in \mathbb{R}_+^3 \mid x \geq \bar{d} - \bar{c} + \epsilon_c + \epsilon_d, \\
&\quad y \geq x + \bar{c} - \epsilon_c - \bar{b} - \epsilon_b, y \geq z - \bar{b} - \epsilon_b\}, \\
\mathcal{U}_1^{(3)} &= \{(x, y, z) \in \mathbb{R}_+^3 \mid x \geq \bar{d} - \bar{c} + \epsilon_c + \epsilon_d, \\
&\quad z \geq x + \bar{c} - \epsilon_c, z \geq y + \bar{b} - \epsilon_b\}.
\end{aligned}$$

Similar to the case of a double-minded bidder (Section 3.4.4), in order to find the robust policy, it is sufficient for bidder 1 to fix x at $x^* = \bar{d} - \bar{c} + \epsilon_c + \epsilon_d$ so that she wins at least one item regardless of the realization of rivals' bids $b_{-1} \in \mathcal{U}_{-1}$, and then optimize over y and z . Figure 3.17 shows $\pi_1^{WO,1}(x, y, z)$, $\pi_1^{WO,2}(x, y, z)$ and $\pi_1^{WO,3}(x, y, z)$ when $x = x^*$. We establish (by enumeration) that an optimal robust policy is $b_1^{RO} = (x^*, y^*, z^*) = (6.14, 12.5, 15.5)$, which is at the intersection of $\pi_1^{WO,1}(x, y, z)$ and $\pi_1^{WO,2}(x, y, z)$. The worst-case payoff corresponding to this robust policy is $\pi_1^{MAXMIN} = 3.765$. Note that $\pi_1^{WO,1}(x, y, z)$ and $\pi_1^{WO,3}(x, y, z)$ also intersect, but the resulting bidding policies are sub-optimal. We also have that $\pi_1^{MAXMIN} = 3.765 < 3.78 = \pi_1^{MINMAX}$, so minimax equality (3.13) does not hold. \square

Example 3.7 illustrates the limitations of the approach used to establish robust policies for double-minded bidders (Proposition 3.12). We have shown in the setting of double minded bidders that minimax equality (3.13) could be used to prove the optimality of robust bidding policies. Nevertheless, as established in Example 3.7, the equality may not hold in more general settings. Thus, finding robust bidding policies might not just be computationally challenging but would also require a new analytical approach that is different from the one presented here.

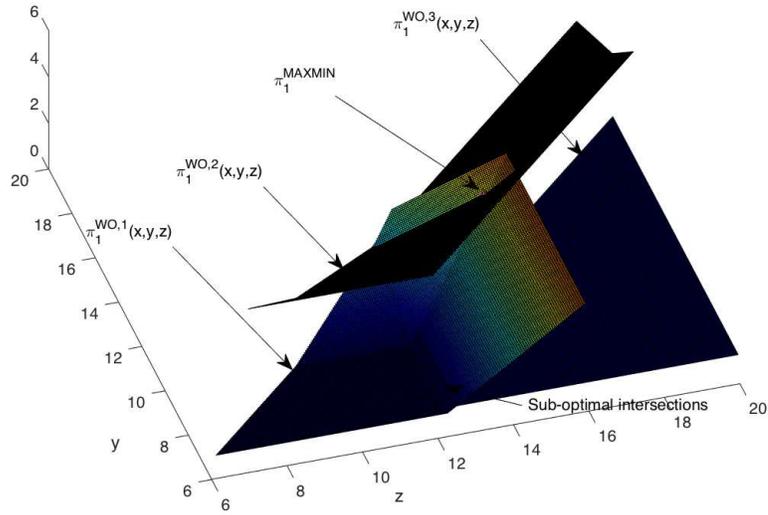


FIGURE 3.17: Illustration of Example 3.7: Worst-case payoff functions corresponding to winning one, two and three items.

3.5 Concluding Remarks

In this paper, we study the bidding problem of an auction bidder who has imperfect information of rivals' bids and wants to maximize her worst-case payoff. We model bidder's information about rivals' bids via an uncertainty set: every point of the uncertainty set is a possible realization of rivals' bids. The bidder's objective is to maximize her worst-case payoff with respect to this uncertainty set. The focus on the worst-case payoff objective requires a different approach from the expected-payoff maximizing bidder models that are prevalent in the auctions literature. Furthermore, unlike those models, our setting allows for distribution-free analysis.

One of the main challenges with our robust optimization approach to the bidding problem is the computational tractability. Our analysis indicates that solving a robust bidding problem involves maximizing a non-concave, discontinuous worst-case

profit function. When the profit function is multi-dimensional, finding the robust bidding policy could be a challenging task. Nevertheless, our analysis provides some insights that could apply when devising optimal or heuristic bidding policies to maximize worst-case payoff in other auction settings. For one simple example, if there is a unique favorable allocation outcome for the bidder, i.e., one that yields better payoff for the bidder than other allocation outcomes regardless of rivals' bids realization, our results readily apply to quite general settings: bidder's robust policy is to bid the minimal amount that ensures winning such allocation. If there is no such favorable allocation outcome, which may happen in settings with heterogeneous items and/or multiple demand, robust bidding policies correspond to the intersections of allocation-specific worst-case payoff functions. The search for optimal bids would then be restricted to bids at the intersections of those worst-case payoff functions. When dealing with multiple demand settings, to overcome the computational challenge of searching over high-dimensional policy spaces, one could employ variable reduction techniques and restrict the attention to only bids on bundles of interest – e.g., bundles with positive valuations, bundles demanded by competitors, or the global bundle.

Finally, note that in settings we considered, minimax equality (3.13) holding was important for establishing the existence of and providing descriptions of robust bidding policies. Hence, understanding of the settings in which minimax equality holds could be useful in finding optimal solutions to the robust bidding problem. Similarly, in settings in which minimax equality fails to hold, one needs to be aware of difficulties of finding and establishing such optimal robust policies, as demonstrated in the case of multiple demand core-selecting auctions.

Economic Order Quantity with Model Uncertainty and Business Interruption Insurance

4.1 Introduction

4.1.1 Background and Overview

Business interruption (BI) insurance offers substantial value to supply chain firms that face natural and man-made disaster risks. In this context, insurers typically offer a menu of contracts to supply chain customers to differentiate them according to their risk profiles. The judicious design of these insurance contracts significantly influences insurers' profits. For example, due to recent occurrences of catastrophic events such as the 2011 flood in Thailand, the 2011 earthquake in Japan and the 2013 explosion in Tianjin port, China, insurance companies suffered significant losses of billions of dollars in claims related to physical damage and business interruption (AGCS, 2015). As a result, insurers need to carefully design their products to remain competitive in the market.

Motivated by the above applications, we consider the dynamic pricing problem

This chapter is based on joint work with Bora Keskin.

of a monopolistic insurer who sells a business interruption insurance product over a planning horizon of N decision epochs. In each decision epoch, the insurer serves a customer who faces an Economic Order Quantity (EOQ) model with disruption risks. Each potential customer has a “type” that corresponds to the degree of financial loss the customer incurs due to disruption, and the customer types are unobservable to the insurer. Hence, the insurer employs vertical differentiation by providing a menu of contracts with different coverage levels. Regarding model uncertainty in this context, we consider two settings of learning: (a) the insurer is uncertain about the probability of disruption of a shipment; and (b) the insurer is uncertain about the distribution of customer’s type. We provide a Bayesian formulation of the learning problem for each of these settings. Using this formulation, we analyze the performance of a myopic Bayesian policy (MBP) and one of its variants in terms of regret, i.e., the profit loss relative to a clairvoyant who knows the underlying value of uncertain parameters.

Our paper provides several contributions to the literature on dynamic pricing with model uncertainty. In particular, we show that the MBP performs well when learning the probability of disruption. The insurer’s belief about this unknown probability of disruption converges to the true hypothesis at an exponential rate, regardless of the presence of information asymmetry. As a result, the regret under the MBP is bounded by a constant. With regard to learning the distribution of customer’s type, we show that under a variant of the MBP that guarantees separability of contracts in every decision epoch, the insurer’s belief also converges to the true hypothesis at an exponential rate and her regret is bounded by a constant.

The organization of this paper is as follows. Section 4.2 describes the basic model elements and assumptions. Section 4.3 follows with the analysis of the learning problems. In particular, Section 4.3.2 analyzes the performance of the MBP in the problem of learning about the probability of disruption, while Section 4.3.3 analyzes the performance of a variant of the MBP in the problem of learning about the

distribution of customer's type. Section 4.4 concludes with some managerial insights. All proofs are presented in the Appendix.

4.1.2 Related Literature

In the operations management literature, the trade-off between learning and earning has been extensively studied in the context of dynamic pricing with demand model uncertainty (Lobo and Boyd, 2003; Araman and Caldentey, 2009; Besbes and Zeevi, 2009; Farias and Van Roy, 2010; Harrison et al., 2012; Broder and Rusmevichientong, 2012; den Boer and Zwart, 2013; Keskin and Zeevi, 2014; Wang et al., 2014). In these studies, a customer's decision is simply whether to accept or reject an offered price. In our paper, the learning-and-earning trade-off is analyzed in the context of vertical product differentiation. Due to information asymmetry, the insurer wants to offer a menu of contracts to maximize her single-period expected profit. Keskin and Birge (2016) recently studied the learning-and-earning trade-off in a similar setting. They considered the case where the decision maker is uncertain about her production cost and learns about it by observing the customer's purchasing decisions and the realized costs. In our study, we consider the EOQ model with stochastic lead time and the application of BI insurance to this model.

Information asymmetry has been a focus of attention in the economic literature (Stiglitz, 1977; Mussa and Rosen, 1978; Myerson, 1982; Guesnerie and Laffont, 1984; Maskin and Riley, 1984; Matthews and Moore, 1987; Page, 1992; Jullien, 2000; Nöldeke and Samuelson, 2007; Winter, 2000; Hellwig, 2010; Chade and Schlee, 2012). In our paper, we study this subject in the specific context of EOQ model with BI insurance. Noting the empirical evidence of information asymmetry in the BI insurance market (Adachi et al., 2016), we model the insurer's decision problem in each decision epoch in the presence of information asymmetry as a mechanism design problem. While the study of the insurer's decision problem is in a dynamic setting,

we focus on the performance of a class of myopic policies under which the insurer solves a mechanism design problem in each decision epoch to maximize her profit in that epoch only. For analyses of optimal dynamic mechanisms, see Belloni et al. (2016) and the references therein.

Our paper also contributes to the growing literature at the interface between operations and finance. Studies in this literature analyze the use of financial instruments in hedging against operational risks, or the use of operational flexibility in hedging against financial risks (Ding et al., 2007; Gaur and Seshadri, 2005; Caldentey and Haugh, 2006; Oum et al., 2006; Caldentey and Haugh, 2009; Chod et al., 2010; Dong and Tomlin, 2012; Kouvelis et al., 2013; Dong et al., 2015; Serpa and Krishnan, 2016; Iancu et al., 2016). Similar to Dong and Tomlin (2012) and Dong et al. (2015), we study the interplay between inventory management and BI insurance in the presence of model uncertainty. As stated earlier, we adopt the economic order quantity (EOQ) framework to model inventory management (Liberatore, 1979; Sphicas, 1982; Sphicas and Nasri, 1984; Schmitt et al., 2010).

4.2 Problem Formulation and Preliminaries

4.2.1 Basic Model Elements

Consider a monopolistic insurer who offers a single BI product for sale to a collection of supply chain firms, hereafter called *the customers*, who sequentially arrive and decide on their insurance contracts and inventory policies over a planning horizon of N decision epochs. As will be explained in detail in Section 4.2.2, each customer manages her inventory via the EOQ framework. In this framework, an inventory policy is specified by a tuple (t, q) , where t is the time difference between order time and arrival time, and q is the length of demand interval satisfied by an order. Let ρ be the probability of a shipment being disrupted and X be a Bernoulli random variable with parameter ρ so that $X = 1$ if there is a disruption for a shipment and

$X = 0$ otherwise. There are I_n selling seasons corresponding to I_n shipments within each decision epoch n . Let X_{in} be a Bernoulli random variable with parameter ρ so that $X_{in} = 1$ if the shipment in selling season i of decision epoch n is disrupted. We use $TC(t, q, \beta, T, \gamma, X)$ to denote the total cost of a shipment, where t and q denote the parameters of the customer's inventory policy, β is the customer's *disruption penalty coefficient* (which will be explained in Section 4.2.2), (T, γ) is the customer's contract choice, and X indicates whether a disruption happens or not. We assume that I_n is sufficiently large for $n \in \{1, 2, \dots, N\}$ so that the customer's average cost in decision epoch n is approximated by:

$$\frac{1}{I_n} \sum_{i=1}^{I_n} TC(t, q, \beta, T, \gamma, X_{in}) \approx \mathbb{E}_X[TC(t, q, \beta, T, \gamma, X)]. \quad (4.1)$$

At the beginning of a decision epoch $n \in \{1, 2, \dots, N\}$, the insurer offers a menu of contracts Γ_n from which the arriving customer can choose an option (or not choose any at all). A collection $\mathbf{\Gamma} = (\Gamma_1, \Gamma_2, \dots, \Gamma_N)$ of menus of contracts is called a *policy*. There are two types of customers: low (L) type and high (H) type. Each customer type $s \in \{L, H\}$ is associated with a disruption penalty coefficient β_s , which represents the customer's financial risk level with respect to a disruption. We assume that $\beta_L < \beta_H$, so low type customers suffer smaller financial loss compared to high type customers. The customers' disruption penalty coefficient β can be different for different decision epochs, with $\phi \in [0, 1]$ being the probability that $\beta = \beta_H$. For each decision epoch n , the insurer wants to offer a contract for each type of customer, so that the menu of contracts Γ_n is a tuple $\{(T_{Ln}, \gamma_{Ln}), (T_{Hn}, \gamma_{Hn})\}$ with T_{Ln}, T_{Hn} being the *premiums* and γ_{Ln}, γ_{Hn} being the *coverage* parameters for low type customers and high type customers, respectively. We refer to a tuple (T, γ) as a γ -type contract. The BI insurance contract offered is in the form of co-insurance, so that the quality γ of a BI contract is represented by the *non-coverage percentage* specified under that

contract. In other words, $1 - \gamma$ is the fraction of *income loss* (defined in Section 4.2.2) reimbursed by the insurer when a disruption happens. A high value of γ corresponds to a low quality contract, and vice versa. The insurer selects contract quality γ in the range $[\underline{\gamma}, \bar{\gamma}]$, where $0 \leq \underline{\gamma} < \bar{\gamma} \leq 1$. If $T_{Ln} = T_{Hn}$ and $\gamma_{Ln} = \gamma_{Hn}$ then we say Γ_n is *pooling*, otherwise we say it is *separating*.

The insurer may not know the exact value of ρ but has an estimate $\hat{\rho}$ for it. Let \hat{X} be a Bernoulli random variable with parameter $\hat{\rho}$. Given a value of $\hat{\rho}$ and a menu of contracts $\Gamma = \{(T_L, \gamma_L), (T_H, \gamma_H)\}$, a customer's optimal inventory policy and contract choice can be determined by solving the problem $\mathcal{CP}(\hat{\rho}, \Gamma)$ given by:

$$\begin{aligned} \min_{s,t,q} \mathbb{E}_{\hat{X}}[TC(t, q, \beta, T_s, \gamma_s, \hat{X})] & \quad (4.2) \\ \text{s.t.} \quad s \in \{0, L, H\}, \quad (t, q) \in \mathcal{R}, & \end{aligned}$$

where $s = 0$ means the customer does not choose any insurance contract, with the convention that $T_0 = 0$ and $\gamma_0 = 1$, i.e., (T_0, γ_0) is a dummy contract with zero premium and insurance coverage. The set $\mathcal{R} \subset \mathbb{R}_+^2$ is a feasible region for (t, q) . Section 4.2.2 provides a detailed description of this feasible region. For a particular decision epoch, if a BI contract is selected, it is effective for all selling seasons within that decision epoch. Similarly, once chosen, a customer's optimal inventory policy remains fixed throughout the entire decision epoch.

We use $\mathcal{IP}(\hat{\rho}, \hat{\phi})$ to denote the insurer's single-period decision problem of designing a menu of contracts, where $\hat{\rho}$ and $\hat{\phi}$ are estimates of the actual probabilities ρ and ϕ , respectively. A detailed description of $\mathcal{IP}(\hat{\rho}, \hat{\phi})$ is given in Section 4.2.2. The sequence of events for each decision epoch n is as follows. Nature chooses the customer's type $\beta \in \{\beta_L, \beta_H\}$ with ϕ being the probability that $\beta = \beta_H$. The value of β is observed by the customer, but not the insurer. The insurer estimates ρ using $\hat{\rho}$ and ϕ using $\hat{\phi}$. The value of $\hat{\rho}$ is known to the customer. The insurer solves $\mathcal{IP}(\hat{\rho}, \hat{\phi})$ to find the optimal menu of contracts Γ_n that maximizes her expected profit for the

current decision epoch. The customer observes Γ_n and solves the problem $\mathcal{CP}(\hat{\rho}, \Gamma_n)$ for optimal contract choice and inventory policy. The chosen contract and inventory policy is then fixed for I_n selling seasons within decision epoch n . Nature chooses the sequence of Bernoulli variables X_{in} with parameter ρ . The customer's average cost and the insurer's profit are then realized based on the values of X_{in} .

4.2.2 Single-period Problems

Customers' Cost Model

This subsection describes a customer's decision problem $\mathcal{CP}(\rho, \Gamma)$ at a particular decision epoch. The customer operates a supply chain with continuous deterministic demand and the individual unit demands are non-interchangeable. We use λ to denote the customer's constant demand rate (units/time). Unfilled demand is backlogged, and each shipment is subject to disruption. With probability $1 - \rho$, there is no disruption and the lead time is a fixed amount τ_N . With probability ρ , a disruption occurs and results in a longer lead time $\tau_D > \tau_N$.¹ Recall that the customer's inventory policy is specified by t , the time difference between the order time and arrival time, and q , the duration over which an inventory order lasts. Since the demand is deterministic, q fully characterizes the order quantity; specifically, the size of an order is given by λq . For convenience, we will hereafter refer to q as the "order quantity". Let K be the fixed ordering cost.² Let h and b (dollars/unit/time) be the inventory holding and backlogging costs, respectively.

In the EOQ model, the operational cost $OC(t, q, \tau)$ corresponding to order time

¹ This lead time could correspond to that of an alternative supplier that the customer purchases from when disruption happens. We assume that the procurement cost from the alternative supplier is the same as the original supplier, but the increase in lead time results in an increase in operational cost due to shortage and an additional cost of capital due to the loss of income.

² Proportional ordering cost does not affect our analysis; hence, without loss of generality, we assume that this component is zero.

t , order quantity q and lead time τ is given by:

$$\begin{aligned}
OC(t, q, \tau) = & \mathbb{1}_{\tau \in [0, t)} \left\{ \frac{1}{2} h \lambda q + h \lambda (t - \tau) \right\} \\
& + \mathbb{1}_{\tau \in [t, t+q)} \left\{ \frac{1}{2q} h \lambda (t + q - \tau)^2 + \frac{1}{2q} b \lambda (\tau - t)^2 \right\} \\
& + \mathbb{1}_{\tau \in [t+q, \infty)} \left\{ \frac{1}{2} b \lambda q + b \lambda (\tau - t - q) \right\} + \frac{K}{q}.
\end{aligned} \tag{4.3}$$

The first term in (4.3) corresponds to the case when the ordered shipment arrives before the “demand interval”, i.e., the duration over which the inventory order is supposed to satisfy demand. In this case, there is no backlogging cost. The second term in (4.3) is when the shipment arrives in the middle of the demand interval. The customer incurs both backlogging cost before the shipment arrives and inventory holding cost after the shipment arrives. Finally, the third term in (4.3) is when the shipment arrives after the demand interval. In that case, all demand is backlogged and there is no inventory holding cost. In the subsequent analysis, we will focus on the case where only the first two cases can occur, while the realization of the third case can be treated similarly. Figure 4.1 illustrates the EOQ model in the first two cases.

We first look at the case without insurance. When there is no disruption, the customer’s total cost per unit time is given by:

$$TC(t, q, \beta, T_0, \gamma_0, 0) = OC(t, q, \tau_N). \tag{4.4}$$

On the other hand, when there is a disruption in the supply, the customer’s total cost per unit time is:

$$TC(t, q, \beta, T_0, \gamma_0, 1) = OC(t, q, \tau_D) + DP(t, q, \beta, \gamma_0), \tag{4.5}$$

where $DP(t, q, \beta, \gamma)$ is the disruption penalty corresponding to an inventory policy (t, q) , a customer’s disruption penalty coefficient β and an insurance non-coverage

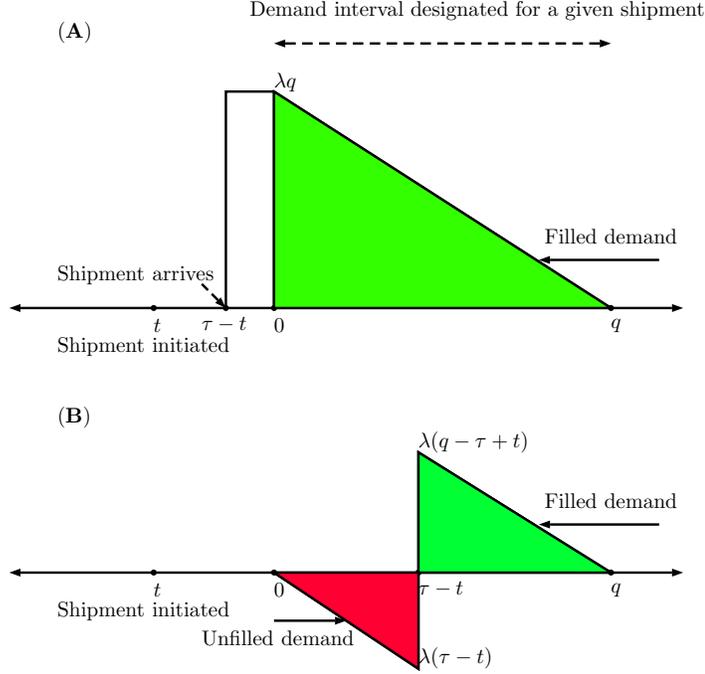


FIGURE 4.1: Illustration of the EOQ model for one selling season: (A) Shipment arrives before demand interval, (B) Shipment arrives during demand interval.

percentage γ . When there is no insurance, we model this disruption penalty as a linear function of the income loss $IL(t, q)$, with the slope being the disruption penalty coefficient β :

$$DP(t, q, \beta, \gamma_0) = \beta IL(t, q). \quad (4.6)$$

In general, if the customer has an insurance contract with non-coverage percentage γ then the *actual* income loss incurred by the customer is $\gamma IL(t, q)$, so the disruption penalty is:

$$DP(t, q, \beta, \gamma) = \beta \gamma IL(t, q). \quad (4.7)$$

The income loss $IL(t, q)$ is modeled as the increase in *shortage cost* due to disruption:

$$IL(t, q) = \mathbb{1}_{\tau_D \in [t, t+q]} \frac{1}{2q} b \lambda (\tau_D - t)^2 + \mathbb{1}_{\tau_D \in [t+q, \infty)} \left\{ \frac{1}{2} b \lambda q + b \lambda (\tau_D - t - q) \right\} \quad (4.8)$$

$$- \mathbb{1}_{\tau_N \in [t, t+q]} \frac{1}{2q} b \lambda (\tau_N - t)^2 - \mathbb{1}_{\tau_N \in [t+q, \infty)} \left\{ \frac{1}{2} b \lambda q + b \lambda (\tau_N - t - q) \right\}.$$

Hence, the expected total cost of a shipment in the case without insurance is:

$$\mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)] = (1 - \rho)TC(t, q, \beta, T_0, \gamma_0, 0) + \rho TC(t, q, \beta, T_0, \gamma_0, 1). \quad (4.9)$$

If the customer purchases a BI contract (T, γ) , her total cost under normal operation and under disruption are respectively:

$$\begin{aligned} TC(t, q, \beta, T, \gamma, 0) &= OC(t, q, \tau_N) + T, \\ TC(t, q, \beta, T, \gamma, 1) &= OC(t, q, \tau_D) + DP(t, q, \beta, \gamma) + T. \end{aligned}$$

The customer's total cost after purchasing insurance is thus:

$$\mathbb{E}_X[TC(t, q, \beta, T, \gamma)] = (1 - \rho)TC(t, q, \beta, T, \gamma, 0) + \rho TC(t, q, \beta, T, \gamma, 1). \quad (4.10)$$

Let $(t^*(\beta), q^*(\beta))$ be the customer's optimal *pre-insurance* inventory policy, i.e.,

$$(t^*(\beta), q^*(\beta)) = \underset{(t, q) \in \mathcal{R}}{\operatorname{argmin}} \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0), X]. \quad (4.11)$$

In problem (4.11), we assume that the feasible region for (t, q) takes the form $\mathcal{R} = \{(t, q) \in \mathbb{R}_+ \mid \tau_N \leq t \leq \tau_D \leq t + q\}$. That is, the customer selects the optimal inventory policies among those such that under normal operations, the shipment arrives before the demand interval starts, while under disruptive operations, the shipment arrives in the middle of the demand interval. We make the following assumption on the problem parameters to facilitate the derivation of closed-form expressions of $t^*(\beta)$ and $q^*(\beta)$ while guaranteeing that $(t^*(\beta), q^*(\beta))$ belongs to \mathcal{R} .

Assumption 4.1. The problem parameters satisfy the following conditions:

$$b > \frac{1 - \rho}{\rho} h, \quad (4.12)$$

$$\tau_D - \tau_N \geq \frac{(2Kh)^{1/2}}{(\lambda\rho)^{1/2}(h + b + b\beta_L\gamma)^{1/2}[\rho(h + b + b\beta_L\gamma) - h]^{1/2}}. \quad (4.13)$$

Conditions (4.12) and (4.13) present lower bounds on the shortage cost b and the delay due to disruption $\tau_D - \tau_N$. With Assumption 4.1 we limit our problem to cases where disruptions have substantial impact on the firm's cost due to high shortage cost and long disruption delay.

Insurer's Decision Problem

Consider the insurer's single-period decision problem $\mathcal{IP}(\rho, \phi)$. The insurer designs a menu of contracts $\Gamma = \{(T_L, \gamma_L), (T_H, \gamma_H)\}$ in each decision epoch. For $s \in \{L, H\}$, we use $R_s(\gamma)$ to denote the long-run average reimbursement for β_s -type customer choosing a γ -type contract:

$$R_s(\gamma) = \rho(1 - \gamma)IL(t^*(\beta_s), q^*(\beta_s)). \quad (4.14)$$

Note that for $(t, q) \in \mathcal{R}$, the income loss $IL(t, q)$ is simply

$$IL(t, q) = \frac{b\lambda}{2q}(\tau_D - t)^2 \mathbb{1}_{t \leq \tau_D}. \quad (4.15)$$

Let $P_s(T, \gamma)$ be the insurer's profit from β_s -type customer choosing a contract (T, γ) .

We have

$$P_s(T, \gamma) = T - R_s(\gamma). \quad (4.16)$$

Let $U(t, q, \beta, \gamma)$ be the customer's post-insurance expected cost excluding insurance premium:

$$U(t, q, \beta, \gamma) = \mathbb{E}_X[TC(t, q, \beta, T, \gamma, X)] - T. \quad (4.17)$$

Also, let $U(\beta, \gamma)$ to be the customer's *optimal* post-insurance expected cost excluding insurance premium:

$$U(\beta, \gamma) = U(t^*(\beta), q^*(\beta), \beta, \gamma). \quad (4.18)$$

We will hereafter refer to $U(\beta, \gamma)$ as the customer's *cost function*.

Consider the insurer's problem $\mathcal{IP}(\rho, \phi)$. Since the insurer does not observe β , she selects an optimal menu of contracts by solving a mechanism design problem:

$$\max_{T_L, T_H, \gamma_L, \gamma_H} (1 - \phi)P_L(T_L, \gamma_L) + \phi P_H(T_H, \gamma_H) \quad (4.19)$$

$$U(\beta_L, \gamma_L) + T_L \leq U(\beta_L, \gamma_H) + T_H, \quad (\text{IC-L})$$

$$U(\beta_H, \gamma_H) + T_H \leq U(\beta_H, \gamma_L) + T_L, \quad (\text{IC-H})$$

$$U(\beta_L, \gamma_L) + T_L \leq U(\beta_L, \gamma_0), \quad (\text{IR-L})$$

$$U(\beta_H, \gamma_H) + T_H \leq U(\beta_H, \gamma_0), \quad (\text{IR-H})$$

$$\underline{\gamma} \leq \gamma_L, \gamma_H \leq \bar{\gamma}.$$

Conditions (IC-L) and (IC-H) are *incentive compatibility* constraints. In particular, (IC-L) guarantees that a β_L -type customer prefers choosing (T_L, γ_L) to choosing (T_H, γ_H) . Similarly, (IC-H) guarantees that a β_H -type customer prefers choosing (T_H, γ_H) to choosing (T_L, γ_L) . Conditions (IR-L) and (IR-H) are, on the other hand, *individual rationality* constraints. They ensure that if the customer purchases an insurance contract, her post-insurance optimal expected cost is at least as low as her optimal expected cost if she does not purchase any insurance. The solution of problem (4.19) thus selects an optimal menu of contracts that maximizes the insurer's single-period expected profit, while satisfying the incentive compatibility and individual rationality constraints. If the customer is indifferent between purchasing an insurance contract and not purchasing one then we assume that she would choose the insurance contract. This assumption does not affect our subsequent analysis, since we can always modify (IR-L) and (IR-H) so that the customer's post-insurance optimal expected cost is guaranteed to be $\epsilon > 0$ smaller than her optimal expected cost without insurance.

4.2.3 Learning Problems

Consider the insurer's decision problem over a planning horizon of N decision epochs. There are two possible types of learning problems arising from this context. In the first problem, the insurer does not observe the exact value of ρ , the probability of a shipment being disrupted, at the beginning of the planning horizon. In the second problem, she does not observe the exact value of ϕ , the proportion of customers with high disruption penalty coefficients. The relationship between the insurer's decisions and the information that she collects to learn about the uncertain parameter is different in these learning problems. We will discuss these problems in turn.

The Problem of Learning ρ

Before the first decision epoch, nature chooses a value of $\rho \in \{\rho_0, \rho_1\}$ that is not observed by the insurer and this value of ρ remains fixed over the entire planning horizon. We denote by π_0 the *prior probability* assigned by the insurer to the event $\{\rho = \rho_1\}$. To exclude trivial cases, in which the insurer knows the parameter ρ with certainty, we assume that $0 < \pi_0 < 1$. We shall occasionally refer to the event $\{\rho = \rho_l\}$ as *hypothesis l* for $l \in \{0, 1\}$. Nature also chooses a value of ϕ but this is observed by the insurer.

Let π_{i_n} be the insurer's assigned probability for the event $\{\rho = \rho_1\}$ in selling season i of decision epoch n ; π_{i_n} will be occasionally referred to as the insurer's *belief* in selling season i of decision epoch n . We use the convention that $\pi_{I_{n-1}+1, n-1} = \pi_{1n}$. Let $\pi_n = \pi_{I_n, n}$ be the insurer's belief about ρ at the end of decision epoch n . We say the insurer's belief is equal to the *true hypothesis* if she assigns probability 1 to the event $\{\rho = \rho_l\}$ when the actual value of ρ is indeed ρ_l , for $l \in \{0, 1\}$. The insurer's dynamic pricing problem proceeds as follows. At the beginning of each decision epoch n , the insurer estimates ρ using $\hat{\rho}_n$ calculated from her most updated

belief:

$$\hat{\rho}_n = \pi_{n-1}\rho_1 + (1 - \pi_{n-1})\rho_0. \quad (4.20)$$

The insurer then selects a menu of contracts Γ_n by solving $\mathcal{IP}(\hat{\rho}_n, \phi)$; $\hat{\rho}_n$ is shared with the customer, so upon her arrival in decision epoch n , the customer selects the optimal contract and inventory policy by solving $\mathcal{CP}(\hat{\rho}_n, \Gamma_n)$. For each selling season $i = 1, 2, \dots, I_n$, the insurer observes the value of X_{in} and updates her belief about ρ using Bayes' rule:

$$\begin{aligned} \pi_{i,n} &= \begin{cases} \frac{\pi_{i-1,n}\rho_1}{\pi_{i-1,n}\rho_1 + (1 - \pi_{i-1,n})\rho_0} & \text{if } X_{i,n} = 1, \\ \frac{\pi_{i-1,n}(1 - \rho_1)}{\pi_{i-1,n}(1 - \rho_1) + (1 - \pi_{i-1,n})(1 - \rho_0)} & \text{otherwise} \end{cases} \\ &= \frac{\pi_{i-1,n}\rho_1^{X_{i,n}}(1 - \rho_1)^{1-X_{i,n}}}{\pi_{i-1,n}\rho_1^{X_{i,n}}(1 - \rho_1)^{1-X_{i,n}} + (1 - \pi_{i-1,n})\rho_0^{X_{i,n}}(1 - \rho_0)^{1-X_{i,n}}}. \end{aligned} \quad (4.21)$$

At the beginning of decision epoch $n + 1$, the insurer computes a new estimate $\hat{\rho}_{n+1}$ of ρ using the most updated belief, i.e., equation (4.20) with π_{n-1} replaced by π_n . The process is repeated in similar fashion.

The Problem of Learning ϕ

Before the first decision epoch, nature chooses a value of $\phi \in \{\phi_0, \phi_1\}$ that is not observed by the insurer and this value of ϕ remains fixed over the entire planning horizon. We denote by $\tilde{\pi}_0 \in (0, 1)$ the prior probability assigned by the insurer to the event $\{\phi = \phi_1\}$. Nature also chooses a value of ρ but it is observed by the insurer in this case.

Let $\tilde{\pi}_n$ be the insurer's assigned probability for the event $\{\phi = \phi_1\}$ in decision epoch n . At the beginning of each decision epoch n , the insurer estimates ϕ using

$\hat{\phi}_n$ calculated from her most updated belief:

$$\hat{\phi}_n = \tilde{\pi}_{n-1}\phi_1 + (1 - \tilde{\pi}_{n-1})\phi_0. \quad (4.22)$$

The insurer then selects a menu of contracts $\Gamma_n = \{(T_{Ln}, \gamma_{Ln}), (T_{Hn}, \gamma_{Hn})\}$ by solving $\mathcal{IP}(\rho, \hat{\phi}_n)$. Upon her arrival in decision epoch n , the customer selects the optimal contract and inventory policy by solving $\mathcal{CP}(\rho, \Gamma_n)$. Let Y_n be a random variable encoding the customer's contract choice in decision epoch n :

$$Y_n = \begin{cases} -1 & \text{if } \Gamma_n \text{ is pooling,} \\ 0 & \text{if } \Gamma_n \text{ is separating and customer chooses contract } (T_{Ln}, \gamma_{Ln}), \\ 1 & \text{if } \Gamma_n \text{ is separating and customer chooses contract } (T_{Hn}, \gamma_{Hn}). \end{cases} \quad (4.23)$$

The sequence $\{Y_0, Y_1, \dots\}$ is referred to as the *contract choice sequence*. We recall from Section 4.2.2 that by solving $\mathcal{IP}(\rho, \hat{\phi}_n)$, the insurer obtains a menu of contracts Γ_n that is *incentive compatible*, i.e., a low type customer selects contract (T_{Ln}, γ_{Ln}) and a high type customer selects contract (T_{Hn}, γ_{Hn}) . Thus, unless $Y_n = -1$, i.e., the offered menu of contracts Γ_n is pooling, the insurer completely observes the customer's type based on her choice of contract. In particular, if $Y_n = 1$ then the customer is of type β_H in decision epoch n and if $Y_n = 0$ then the customer is of type β_L in decision epoch n . Another observation is that the insurer updates her belief about the customer type only once in each decision epoch. For decision epoch n , we have $\tilde{\pi}_n = \tilde{\pi}_{n-1}$ if $Y_n = -1$ and

$$\tilde{\pi}_n = \begin{cases} \frac{\tilde{\pi}_{n-1}(1 - \phi_1)}{\tilde{\pi}_{n-1}(1 - \phi_1) + (1 - \tilde{\pi}_{n-1})(1 - \phi_0)} & \text{if } Y_n = 0, \\ \frac{\tilde{\pi}_{n-1}\phi_1}{\tilde{\pi}_{n-1}\phi_1 + (1 - \tilde{\pi}_{n-1})\phi_0} & \text{if } Y_n = 1, \end{cases} \\ = \frac{\tilde{\pi}_{n-1}\phi_1^{Y_n}(1 - \phi_1)^{1-Y_n}}{\tilde{\pi}_{n-1}\phi_1^{Y_n}(1 - \phi_1)^{1-Y_n} + (1 - \tilde{\pi}_{n-1})\phi_0^{Y_n}(1 - \phi_0)^{1-Y_n}}, \quad Y_n \in \{0, 1\}. \quad (4.24)$$

At the beginning of decision epoch $n + 1$, the insurer computes a new estimate $\hat{\phi}_{n+1}$ of ϕ using the most updated belief, i.e., equation (4.22) with $\tilde{\pi}_{n-1}$ replaced by $\tilde{\pi}_n$. The process is repeated in similar fashion.

4.2.4 Induced Probability Measures and Performance Metrics

In the problem of learning ρ , for a given policy $\Gamma = (\Gamma_1, \Gamma_2, \dots, \Gamma_N)$, there are two probability measures \mathbb{P}_0^Γ and \mathbb{P}_1^Γ on the outcome space of the claim sequence $\{X_{i,n}\}$ corresponding to two hypotheses $H_0 : \rho = \rho_0$ and $H_1 : \rho = \rho_1$, respectively. These probability measures are given by the following formula:

$$\mathbb{P}_l^\Gamma(X_{1,1} = x_{1,1}, \dots, X_{i,n} = x_{i,n}) = \rho_l^{x_{1,1}}(1 - \rho_l)^{1-x_{1,1}} \dots \rho_l^{x_{i,n}}(1 - \rho_l)^{1-x_{i,n}}, \quad (4.25)$$

$$l \in \{0, 1\}, \quad n \in \{1, 2, \dots, N\}.$$

Conditional on a hypothesis H_l , the insurer's expected profit (with the expectation taken over customer types) in decision epoch n from offering Γ_n is denoted by $v_l(\Gamma_n)$. Recall that $\Gamma_n = \{(T_{Ln}, \gamma_{Ln}), (T_{Hn}, \gamma_{Hn})\}$, so we have

$$v_l(\Gamma_n) = (1 - \phi)P_L(T_{Ln}, \gamma_{Ln}) + \phi P_H(T_{Hn}, \gamma_{Hn}). \quad (4.26)$$

Let $V_l^\Gamma(N)$ be the conditional expected total profit (with the expectation taken with respect to the induced probability measure P_l^Γ):

$$V_l^\Gamma(N) = \mathbb{E}_l^\Gamma \left\{ \sum_{n=1}^N \tilde{v}_l(\Gamma_n) \right\}. \quad (4.27)$$

Let Γ_l^* be the optimal single-period menu of contracts if the insurer has perfect information about ρ given that $\rho = \rho_l$. We are interested in the following performance metric:

$$\Delta_l^\Gamma(N) = \frac{1}{v_l(\Gamma_l^*)} [Nv_l(\Gamma_l^*) - V_l^\Gamma(N)], \quad l \in \{0, 1\}. \quad (4.28)$$

The first term in the bracket on the right hand side of the above equation corresponds to the expected total profit for a clairvoyant who knows the exact value of ρ before

the start of decision epoch 1. The second term in the bracket is the expected total profit of the insurer who is uncertain about the value of ρ and follows a policy $\mathbf{\Gamma}$. Thus, $\Delta_l^\Gamma(N)$ is the *regret* of the insurer over a planning horizon of N decision epochs, i.e., it is the number of decision epochs of ideal expected profit that the insurer loses over N decision epochs due to the uncertainty in ρ when the actual value of ρ is ρ_l . We also consider the average regret over the two hypotheses of ρ :

$$\Delta^\Gamma(N) = \frac{1}{2} [\Delta_1^\Gamma(N) + \Delta_2^\Gamma(N)]. \quad (4.29)$$

Similarly, in the problem of learning ϕ , for any policy $\mathbf{\Gamma} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_n\}$, there are two induced probability measures $\tilde{\mathbb{P}}_0^\Gamma$ and $\tilde{\mathbb{P}}_1^\Gamma$ corresponding to hypotheses $\tilde{H}_0 : \phi = \phi_0$ and $\tilde{H}_1 : \phi = \phi_1$. Let $\mathcal{P}(\mathbf{\Gamma})$ be the set of indices of decision epoch n in which Γ_n is a pooling contract. For $l \in \{0, 1\}$ and $n \in \{1, 2, \dots, N\}$, we have

$$\tilde{\mathbb{P}}_l^\Gamma(Y_1 = y_1, Y_2 = y_2, \dots, Y_n = y_n) = \prod_{m \in \{1, 2, \dots, n\} \setminus \mathcal{P}(\mathbf{\Gamma})} \phi_l^{y_m} (1 - \phi_l)^{1 - y_m}. \quad (4.30)$$

As in the problem of learning ρ , we quantify the performance of a policy by the regret of the insurer when compared with a clairvoyant who knows the exact value of ϕ . In particular, let $\tilde{v}_l(\Gamma_n)$ be the insurer's expected profit (with expectation taken over customers' types) from offering $\Gamma_n = \{(T_{Ln}, \gamma_{Ln}), (T_{Hn}, \gamma_{Hn})\}$ in decision epoch n , conditioned on the hypothesis $\tilde{H}_l : \phi = \phi_l$:

$$\tilde{v}_l(\Gamma_n) = (1 - \phi_l)P_L(T_{Ln}, \gamma_{Ln}) + \phi_l P_H(T_{Hn}, \gamma_{Hn}). \quad (4.31)$$

Let $\tilde{V}_l^\Gamma(N)$ be the insurer's conditional expected total profit (with expectation taken with respect to the induced probability measure $\tilde{\mathbb{P}}_l^\Gamma$) over the first N decision epoch:

$$\tilde{V}_l^\Gamma(N) = \tilde{\mathbb{E}}_l^\Gamma \left\{ \sum_{n=1}^N v_l(\Gamma_n) \right\}. \quad (4.32)$$

Let $\tilde{\Gamma}_l^*$ be the optimal single-period menu of contracts if the insurer has perfect information about ϕ given that $\phi = \phi_l$. The insurer's regret metrics are defined as follow:

$$\tilde{\Delta}_l^\Gamma(N) = \frac{1}{\tilde{v}_l(\tilde{\Gamma}_l^*)} \left[N\tilde{v}_l(\tilde{\Gamma}_l^*) - \tilde{V}_l(N) \right], \quad l \in \{0, 1\}, \quad (4.33)$$

$$\tilde{\Delta}^\Gamma(N) = \frac{1}{2} \left[\tilde{\Delta}_0^\Gamma(N) + \tilde{\Delta}_1^\Gamma(N) \right]. \quad (4.34)$$

4.2.5 The Myopic Bayesian Policy

In our Bayesian formulation of the insurer's dynamic pricing problem, we focus on the case where the insurer selects the menu of contracts Γ_n in each decision epoch n that maximizes her expected profit in that decision epoch only, using the most updated belief of the uncertain parameter, i.e., either ρ or ϕ . The pricing policy obtained in this way is called the myopic Bayesian policy (MBP). For clarity, we will use $\Psi = (\Psi_1, \Psi_2, \dots, \Psi_N)$ to denote the MBP. We are interested in the performance of Ψ in terms of the performance metrics defined in Section 4.2.4, in the problems of learning ρ and ϕ . In establishing the results, it is useful to note that if there exists an m such that Ψ_m is pooling in decision epoch m , then $\Psi_{m'}$ must be pooling for all $m' > m$. In other words, for some $n \geq 1$,

$$\tilde{\mathbb{P}}_l^\Gamma(Y_1 = y_1, Y_2 = y_2, \dots, Y_n = y_n) = 0, \quad (4.35)$$

if $\exists m \in \{1, 2, \dots, n-1\}$ such that $y_m = -1$ and $\exists m' \in \{m+1, \dots, n\}$ such that $y_{m'} \geq 0$. We next define several classes of policies for which the performance is relevant to that of Ψ in the problem of learning ϕ .

Definition 4.1. A policy $\Gamma = \{\Gamma_1, \Gamma_2, \dots, \Gamma_N\}$ is *separating* if for each decision epoch $n \in \{1, 2, \dots, N\}$, the menu of contracts Γ_n is separating.

Definition 4.2. A policy $\Gamma = \{\Gamma_1, \Gamma_2, \dots, \Gamma_N\}$ is *self-selecting* if it is separating and for each decision epoch $n \in \{1, 2, \dots, N\}$, the menu of contract $\Gamma_n =$

$\{(T_{Ln}, \gamma_{Ln}), (T_{Hn}, \gamma_{Hn})\}$ is such that a low type customer chooses (T_{Ln}, γ_{Ln}) and a high type customer chooses (T_{Hn}, γ_{Hn}) .

Note that Ψ is not necessarily a self-selecting policy, since it is possible that the menu of contracts under that policy is pooling in some decision epochs. On the other hand, a self-selecting policy is not necessarily the MBP, since a menu of contracts in a decision epoch might not be optimal for the insurer's decision problem in that decision epoch.

4.3 Analysis

4.3.1 Preliminary Results

To facilitate subsequent analysis, we provide some preliminary results of the single-period decision problems. The following proposition establishes the customer's optimal order quantity and time.

Proposition 4.1. *For each $\beta > 0$, we have*

$$q^*(\beta) = \frac{(2K\rho)^{1/2}(h+b+b\beta)^{1/2}}{(\lambda h)^{1/2}[\rho(h+b+b\beta)-h]^{1/2}}, \quad (4.36)$$

$$t^*(\beta) = \tau_D - q^*(\beta) \frac{h}{\rho(h+b+b\beta)}. \quad (4.37)$$

Corollary 4.1. *$q^*(\beta)$ is decreasing and $t^*(\beta)$ is increasing in β .*

As we can see, customers with high disruption penalty coefficient β place orders early and at small quantities. The following proposition provides a characterization of the customer's cost function $U(\beta, \gamma)$.

Proposition 4.2. *The customer's cost function $U(\beta, \gamma)$ is given by*

$$U(\beta, \gamma) = f_0(\beta) + f(\beta)\gamma, \quad (4.38)$$

where $f(\beta)$ is defined by:

$$f(\beta) = \rho \frac{\beta b \lambda}{2q^*(\beta)} (\tau_D - t^*(\beta))^2, \quad (4.39)$$

and $f_0(\beta)$ is a function of β .

Note that the customer's cost function $U(\beta, \gamma)$ is linear in the contract quality γ , with $f(\beta)$ being the customer's marginal cost. From (4.39), we can see that this marginal cost is equal to the expected disruption penalty due to income loss. The next results are useful for characterizing the problem $\mathcal{IP}(\rho, \phi)$. First, observe that $\mathbb{E}_X[TC(t, q, \beta, T, \gamma)]$ is increasing in β and γ for all $t > 0$ and $q > 0$. Thus, we have the following lemma.

Lemma 4.1. $U(\beta, \gamma)$ is monotonically increasing in β and γ for $\beta > 0$, $\gamma \in [\underline{\gamma}, \bar{\gamma}]$.

To characterize $\mathcal{IP}(\rho, \phi)$, we need another property of $U(\beta, \gamma)$, namely the *single-crossing property*. This is guaranteed by the following technical assumption on the problem parameters.

Assumption 4.2. The problem parameters satisfy $\beta_H < \bar{\beta}_H$, where

$$\bar{\beta}_H \triangleq \frac{(h^2 - 16\rho b h - 16\rho h^2 + 16\rho^2 b^2 + 32\rho^2 b h + 16\rho^2 h^2)^{1/2}}{4\rho b} + \frac{h}{4\rho b}. \quad (4.40)$$

Lemma 4.2. The customer's cost function $U(\beta, \gamma)$ satisfies the single-crossing property, i.e.,

$$U_{\beta\gamma}(\beta, \gamma) > 0, \quad \forall \beta \in [\beta_L, \beta_H], \forall \gamma \in [0, 1]. \quad (4.41)$$

The following lemma characterizes the mechanism design problem (4.19) based on the monotonicity and single-crossing property of $U(\beta, \gamma)$.

Lemma 4.3. The problem $\mathcal{IP}(\rho, \phi)$ satisfies the following:

(i) for any feasible solution $\Gamma = \{(T_L, \gamma_L), (T_H, \gamma_H)\}$, the non-coverage levels γ_L, γ_H satisfy

$$\underline{\gamma} \leq \gamma_H \leq \gamma_L \leq \bar{\gamma}, \quad (4.42)$$

(ii) at the optimal solution, constraint (IR-L) binds,

(iii) at the optimal solution, constraint (IC-H) binds.

Lemma 4.3 shows that the insurer offers small coverage for low type customers and high coverage for high type customers. Furthermore, since (IR-L) binds at the optimal solution, the insurer extracts the entire surplus of low type customers by offering a contract that makes them indifferent between purchasing insurance and not purchasing any insurance.

Proposition 4.3. *The optimal menu of contracts $\Gamma^* = \{(T_L^*, \gamma_L^*), (T_H^*, \gamma_H^*)\}$ to the problem $\mathcal{IP}(\rho, \phi)$ satisfy*

$$T_L^* = U(\beta_L, \gamma_0) - U(\beta_L, \gamma_L^*), \quad (4.43)$$

$$T_H^* = U(\beta_L, \gamma_0) - U(\beta_L, \gamma_L^*) + U(\beta_H, \gamma_L^*) - U(\beta_H, \gamma_H^*). \quad (4.44)$$

Observe that (4.44) can be rewritten as

$$T_H^* = [U(\beta_H, \gamma_0) - U(\beta_H, \gamma_H^*)] - [U(\beta_H, \gamma_0) - U(\beta_L, \gamma_0) + U(\beta_L, \gamma_L^*) - U(\beta_H, \gamma_L^*)]. \quad (4.45)$$

The first term on the right hand side of (4.45) is the premium level at which high type customers are indifferent between purchasing and not purchasing insurance. The second term on the right hand side of (4.45) is positive (due to the single-crossing property) and is the *information rent* that must be paid to high type customers so that they have an incentive to select contract (T_H^*, γ_H^*) .

Proposition 4.4. *The optimal menu of contracts $\Gamma^* = \{(T_L^*, \gamma_L^*), (T_H^*, \gamma_H^*)\}$ to the problem $\mathcal{IP}(\rho, \phi)$ satisfy*

(i) if $\beta_H < \min(1, \bar{\beta}_H)$ then $\gamma_L^* = \gamma_H^* = \bar{\gamma}$ (pooling optimum),

(ii) if $1 \leq \beta_H < \bar{\beta}_H$, $\frac{f(\beta_L)}{\beta_L} < f(\beta_H)$ and

$$\phi < \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)} \quad (4.46)$$

then $\gamma_L^* = \gamma_H^* = \underline{\gamma}$ (pooling optimum),

(iii) if $1 \leq \beta_H < \bar{\beta}_H$ and either

$$\frac{f(\beta_L)}{\beta_L} \geq f(\beta_H) \quad (4.47)$$

or

$$\frac{f(\beta_L)}{\beta_L} < f(\beta_H), \quad \phi \geq \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)} \quad (4.48)$$

then $\gamma_L^* = \bar{\gamma}$ and $\gamma_H^* = \underline{\gamma}$ (separating optimum).

We can see that separating menus of contracts never occur when the disruption penalty coefficient for high type customer is low, i.e., $\beta_H < 1$. If such a condition does not hold, then the insurer selects a separating menu of contracts when the marginal cost of low type customers is large compared to that of high type customers, i.e., condition (4.47), or if the proportion of high type customers is above a certain threshold, i.e., condition (4.48).

4.3.2 Belief Convergence and Regret in Learning ρ

We show that in the problem of learning ρ , the insurer's belief converges to the true hypothesis at an exponential rate. Thus, the insurer learns about the actual value of ρ after the first decision epoch, given that I_1 is sufficiently large.

Proposition 4.5. *For any policy Γ , there exist $\mu_0, \mu_1 > 0$ such that for all $i \geq 1$, we have*

$$\begin{aligned}\mathbb{E}_0^\Gamma[\pi_{i,1}] &\leq \mu_0 \exp(-\mu_1 i), \\ \mathbb{E}_1^\Gamma[1 - \pi_{i,1}] &\leq \mu_0 \exp(-\mu_1 i).\end{aligned}$$

The main idea behind the proof of Proposition 4.5 is that the sequence of cumulative sums of differences between random variables $X_{i,1}$ and their expected value is a martingale. By utilizing the probability measure (4.25), one can express a belief $\pi_{i,1}$ as a function of that martingale. The Azuma-Hoeffding inequality is then used to bound any large deviations of the described martingale, establishing the exponential convergence results.

Corollary 4.2. *For each policy Γ , the posterior probabilities $\{\pi_{i,1}\}$ converge almost surely as $i \rightarrow \infty$ to the true hypothesis under both \mathbb{P}_0^Γ and \mathbb{P}_1^Γ .*

Note that Proposition 4.5 and Corollary 4.2 are true for any policy Γ and are not necessarily limited to the MBP. They are direct consequences of the fact that the realization of the claim sequence $\{X_{in}\}$ is independent of the pricing policy Γ . These results imply that under the MBP Ψ , the insurer's regret is bounded by a constant. The following proposition formalizes this observation.

Proposition 4.6. *There exists a constant $C > 0$ such that $\Delta^\Psi(N) \leq C$ for all $N \geq 1$.*

4.3.3 Belief Convergence and Regret in Learning ϕ

Unlike in the problem of learning ρ , in the problem of learning ϕ , the realization of observed data, i.e., the contract choice sequence, may depend on a particular policy Γ . This makes it less obvious whether the insurer's belief would eventually converge to the true hypothesis. The following proposition shows that the insurer's

belief about ϕ converges to *some* value, even though that might not be the true hypothesis.

Proposition 4.7. *For each policy Γ , the posterior probabilities $\{\tilde{\pi}_n\}$ converges almost surely as $n \rightarrow \infty$ to a limit probability $\tilde{\pi}_\infty$ under both $\tilde{\mathbb{P}}_0^\Gamma$ and $\tilde{\mathbb{P}}_1^\Gamma$.*

In the absence of pooling contracts, the problem of learning ϕ is similar to the problem of learning ρ in that the realization of observed data is independent of the insurer's action. Thus, under a self-selecting policy, the insurer's belief converges to the true hypothesis at an exponential rate.

Proposition 4.8. *If a policy Γ is self-selecting then there exists $\mu_0, \mu_1 > 0$ such that for all $n \geq 1$, we have*

$$\begin{aligned}\tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n] &\leq \mu_0 \exp(-\mu_1 n), \\ \tilde{\mathbb{E}}_1^\Gamma[1 - \tilde{\pi}_n] &\leq \mu_0 \exp(-\mu_1 n).\end{aligned}$$

Under the condition of information asymmetry, the insurer's selection of an optimal menu of contracts depends on her estimation of the actual value of ϕ . Due to the possibility of having pooling contracts (Proposition 4.4), the insurer's belief under the MBP Ψ is not guaranteed to converge to the true hypothesis. However, we show that by employing a variant of the MBP with flexibly chosen constraints on the difference of quality levels to enforce separability in each decision epoch, the insurer gets a regret that is also bounded by a constant.

For $\epsilon > 0$ and $\alpha > 0$, let $\epsilon_n(\alpha) = \epsilon \exp(-\alpha n)$. Consider the policy $\Psi_{\epsilon, \alpha}$ obtained in a similar fashion as the MBP Ψ , except that in each decision epoch n , the optimal menu of contracts is obtained by solving the following modified version of the problem

$\mathcal{IP}(\rho, \hat{\phi}_n)$, which we will refer to as the problem $\mathcal{IP}_{\epsilon, \alpha}(\rho, \hat{\phi}_n)$:

$$\max_{T_L, T_H, \gamma_L, \gamma_H} (1 - \hat{\phi}_n)P_L(T_L, \gamma_L) + \hat{\phi}_n P_H(T_H, \gamma_H) \quad (4.49)$$

$$U(\beta_L, \gamma_L) + T_L \leq U(\beta_L, \gamma_H) + T_H, \quad (\text{IC-L})$$

$$U(\beta_H, \gamma_H) + T_H \leq U(\beta_H, \gamma_L) + T_L, \quad (\text{IC-H})$$

$$U(\beta_L, \gamma_L) + T_L \leq U(\beta_L, \gamma_0), \quad (\text{IR-L})$$

$$U(\beta_H, \gamma_H) + T_H \leq U(\beta_H, \gamma_0), \quad (\text{IR-H})$$

$$|\gamma_L - \gamma_H| \geq \epsilon_n(\alpha), \quad (4.50)$$

$$\underline{\gamma} \leq \gamma_L, \gamma_H \leq \bar{\gamma}.$$

The policy $\Psi_{\epsilon, \alpha}$ is referred to as the constrained variant of the myopic Bayesian policy with parameters (ϵ, α) , or $\text{CMBP}(\epsilon, \alpha)$. Note that in the problem $\mathcal{IP}_{\epsilon, \alpha}(\rho, \hat{\phi}_n)$, we enforce the condition (4.50) so that the optimal menu of contracts is guaranteed to be separating. Any solution Γ to the problem $\mathcal{IP}_{\epsilon, \alpha}(\rho, \hat{\phi}_n)$ is feasible to the problem $\mathcal{IP}(\rho, \hat{\phi}_n)$. However, due to the addition of constraint (4.50), the insurer's single-period profit under Γ is no more than that of the optimal solution of $\mathcal{IP}(\rho, \hat{\phi}_n)$. Hence, under $\text{CMBP}(\epsilon, \alpha)$, the insurer sacrifices a certain degree of optimality in exchange for separability. The benefit of this policy is that, in terms of learning, the insurer's belief is guaranteed to converge to the true hypothesis. Consequently, the regret is bounded by a constant, as we have in the case of learning ρ .

Proposition 4.9. *For any $\epsilon > 0$ and $\alpha > 0$, under the policy $\Psi_{\epsilon, \alpha}$, there exist $\mu_0, \mu_1 > 0$ such that for all $n \geq 1$, we have*

$$\tilde{\mathbb{E}}_0^{\Psi_{\epsilon, \alpha}}[\tilde{\pi}_n] \leq \mu_0 \exp(-\mu_1 n),$$

$$\tilde{\mathbb{E}}_1^{\Psi_{\epsilon, \alpha}}[1 - \tilde{\pi}_n] \leq \mu_0 \exp(-\mu_1 n).$$

Proposition 4.10. *For $\epsilon > 0$ sufficiently small and $\alpha > 0$, there exists a constant $C > 0$ such that $\tilde{\Delta}^{\Psi_{\epsilon, \alpha}}(N) \leq C$ for all $N \geq 1$.*

Several useful observations help establish the results of Proposition 4.10. First, if the problem primitives are such that the insurer's optimal single-period decision is the same regardless of the value of ϕ then the insurer's regret is obviously zero. If the clairvoyant selects a separating menu of contracts, then for ϵ sufficiently small, a sufficiently small value of $\tilde{\pi}_{n-1}$ implies that the menu of contracts $\Psi_{\epsilon,\alpha}^n$ is the same with that of the clairvoyant. On the other hand, due to the belief convergence results in Proposition 4.8, the probability of $\tilde{\pi}_{n-1}$ exceeding a constant threshold is decaying at an exponential rate. Together, these two facts establish a constant upper bound on the regret.

In the case where the clairvoyant selects a pooling contract, we can show that when $\tilde{\pi}_{n-1}$ is sufficiently small, the separating constraint (4.50) in the problem $\mathcal{IP}_{\epsilon,\alpha}(\rho, \hat{\phi}_n)$ must bind. Thus, the insurer offers a separating menu of contracts in which the quality level for high type customers is ideal, i.e., it is the same with that of the clairvoyant's selection, while the quality level for low type customers differs only by an $\epsilon_n(\alpha)$ amount to its high type counterpart. The single-period regret in this case is a linear function of $\epsilon_n(\alpha)$, which is decaying at an exponential rate. In addition, the probability of $\tilde{\pi}_{n-1}$ exceeding a constant threshold is also decaying at an exponential rate. Hence, the regret must be bounded above by a constant.

4.4 Concluding Remarks

In this paper, we formulate and study an Economic Order Quantity model in the presence of business interruption insurance and model uncertainty. In this setting, we analyze the performance of the myopic Bayesian policy and one of its variants in the dynamic pricing problem of a monopolistic insurer. Unlike most studies in the literature of dynamic pricing with model uncertainty, our paper considers the setting in which the insurer wants to perform price discrimination by offering a menu of contracts instead of a simple take-it-or-leave-it offer. In such a setting, information

asymmetry is an important factor that affects the insurer's optimal decisions. We show that in the problem of learning about the probability of disruption, the realization of observed data is independent of the pricing policies, so the insurer can learn about the uncertain parameter at an exponential rate and her regret is bounded by a constant. In the problem of learning about the distribution of customer's type, this is no longer the case. However, by utilizing a variant of the MBP to enforce separability of contracts in every decision epoch, the insurer can achieve similar performance as in the case of learning about the disruption probability, i.e., belief converging to the true hypothesis at an exponential rate and regret being bounded by a constant.

Appendix A

Proofs in Chapter 2

A.1 Proofs in Section 2.2

Lemma 2.1

Proof. The first assertion is trivial, so we focus on the second. To organize our induction we denote by \mathcal{H}_k the assertion

$$v_k(s + \epsilon) < v_k(s) \quad \text{for all } s \in \mathcal{I} \text{ and all } \epsilon > 0.$$

When $k = 1$, we have $v_1(s) = 1 - F(s)$, and admissibility of F implies v_1 is strictly decreasing on \mathcal{I} . This establishes the base case \mathcal{H}_1 .

For $k > 1$ we assume that \mathcal{H}_{k-1} holds, and we note by the Bellman recursion

(2.11) and the characterizing properties of admissible distributions that

$$\begin{aligned}
v_k(s + \epsilon) - v_k(s) &= F(s + \epsilon)v_{k-1}(s + \epsilon) + \int_{s+\epsilon}^{\infty} \max\{v_{k-1}(s + \epsilon), 1 + v_{k-1}(x)\}f(x) dx \\
&\quad - F(s)v_{k-1}(s) - \int_s^{\infty} \max\{v_{k-1}(s), 1 + v_{k-1}(x)\}f(x) dx \\
&\leq F(s + \epsilon)v_{k-1}(s + \epsilon) + \int_{s+\epsilon}^{\infty} \max\{v_{k-1}(s), 1 + v_{k-1}(x)\}f(x) dx \\
&\quad - F(s + \epsilon)v_{k-1}(s) - \int_{s+\epsilon}^{\infty} \max\{v_{k-1}(s), 1 + v_{k-1}(x)\}f(x) dx \\
&= F(s + \epsilon) \{v_{k-1}(s + \epsilon) - v_{k-1}(s)\},
\end{aligned}$$

where we first used $v_{k-1}(s + \epsilon) < v_{k-1}(s)$ and then used the trivial estimate

$$\{F(s + \epsilon) - F(s)\}v_{k-1}(s) \leq \int_s^{s+\epsilon} \max\{v_{k-1}(s), 1 + v_{k-1}(x)\}f(x) dx.$$

For $s \in \mathcal{I}$ one has strict positivity of $F(s + \epsilon)$, so by the induction hypothesis \mathcal{H}_{k-1} we have $v_k(s + \epsilon) - v_k(s) \leq F(s + \epsilon) \{v_{k-1}(s + \epsilon) - v_{k-1}(s)\} < 0$. \square

Lemma 2.2

Proof. First, note that after subtracting $v_{k-1}(s)$ from both sides of equation (2.13), we have

$$v_k(s) - v_{k-1}(s) = \int_s^{h_k(s)} \{1 + v_{k-1}(x) - v_{k-1}(s)\} dF(x).$$

The map $x \mapsto v_{k-1}(x)$ is monotone decreasing, so the factor $\{1 + v_{k-1}(x) - v_{k-1}(s)\}$ is bounded above by one. This gives us our upper bound in (2.15). The representation (2.12) for h_k tells us the integrand is non-negative on $[s, h_k(s)]$, and this gives the lower bound in (2.15). \square

Proposition 2.1

Proof. We first derive a recursion for the difference $v_k(s) - v_k(t)$. For $0 \leq s \leq t < \infty$, we have from (2.11) that

$$\begin{aligned} v_k(s) - v_k(t) &= F(s)\{v_{k-1}(s) - v_{k-1}(t)\} \\ &+ \int_s^t \max\{v_{k-1}(s) - v_{k-1}(t), 1 + v_{k-1}(x) - v_{k-1}(t)\} dF(x) \\ &+ \int_t^\infty [\max\{v_{k-1}(s), 1 + v_{k-1}(x)\} - \max\{v_{k-1}(t), 1 + v_{k-1}(x)\}] dF(x). \end{aligned} \tag{A.1}$$

Next, we let

$$\begin{aligned} a_{k-1}(x) &\stackrel{\text{def}}{=} \min\{v_{k-1}(s) - v_{k-1}(t), v_{k-1}(s) - v_{k-1}(x) - 1\}, \\ b_{k-1}(x) &\stackrel{\text{def}}{=} \min\{1 + v_{k-1}(x) - v_{k-1}(t), 0\}, \end{aligned}$$

and we note that the difference

$$\max\{v_{k-1}(s), 1 + v_{k-1}(x)\} - \max\{v_{k-1}(t), 1 + v_{k-1}(x)\}$$

which appears in the last integrand of (A.1) can be written as

$$\max\{v_{k-1}(s), 1 + v_{k-1}(x)\} - \max\{v_{k-1}(t), 1 + v_{k-1}(x)\} = \max\{a_{k-1}(x), b_{k-1}(x)\}.$$

Here $s \leq t$, so when $b_{k-1}(x) \leq 0$ the monotonicity of the value functions in Lemma 2.1 implies that $0 \leq v_{k-1}(s) - v_{k-1}(x) - 1$. It then follows that $0 \leq a_{k-1}(x)$ and $\max\{a_{k-1}(x), b_{k-1}(x)\} = \max\{a_{k-1}(x), 0\}$. In general, we then have the equivalence

$$\max\{v_{k-1}(s), 1 + v_{k-1}(x)\} - \max\{v_{k-1}(t), 1 + v_{k-1}(x)\} = \max\{a_{k-1}(x), 0\},$$

and we can substitute this representation and the explicit expression for $a_{k-1}(x)$ in (A.1) to obtain the simplified *difference recursion*

$$\begin{aligned} v_k(s) - v_k(t) &= F(s)\{v_{k-1}(s) - v_{k-1}(t)\} \\ &+ \int_s^t \max\{v_{k-1}(s) - v_{k-1}(t), 1 + v_{k-1}(x) - v_{k-1}(t)\} dF(x) \\ &+ \int_t^\infty \max\{\min\{v_{k-1}(s) - v_{k-1}(t), v_{k-1}(s) - v_{k-1}(x) - 1\}, 0\} dF(x). \end{aligned} \quad (\text{A.2})$$

We now let \mathcal{H}_k be the assertion that

$$v_{k-1}(s) - v_{k-1}(t) \leq v_k(s) - v_k(t) \quad \text{for all } 0 \leq s \leq t < \infty,$$

and we prove by induction that \mathcal{H}_k holds for all $k \geq 1$. We first note that for $k = 1$ we have $v_0(s) = 0$ for all $s \in [0, \infty)$. By the difference recursion (A.2) we obtain $v_1(s) - v_1(t) = F(t) - F(s) \geq 0 = v_0(s) - v_0(t)$, so the base case \mathcal{H}_1 holds.

Next, we suppose that \mathcal{H}_{k-1} holds, and we apply \mathcal{H}_{k-1} to *all* of the terms on the right-hand side of (A.2). We then obtain that

$$\begin{aligned} v_k(s) - v_k(t) &\leq F(s)\{v_k(s) - v_k(t)\} \\ &+ \int_s^t \max\{v_k(s) - v_k(t), 1 + v_k(x) - v_k(t)\} dF(x) \\ &+ \int_t^\infty \max\{\min\{v_k(s) - v_k(t), v_k(s) - v_k(x) - 1\}, 0\} dF(x). \end{aligned}$$

We can now apply the difference recursion (A.2) a second time after we replace k by $k + 1$. This tells us that the right-hand side above is equal to the difference $v_{k+1}(s) - v_{k+1}(t)$, thus completing the proof of \mathcal{H}_k and of the proposition. \square

Corollary 2.1

Proof. Here we only have to note that

$$\begin{aligned} h_{k+1}(s) &= \sup\{x \in [s, \infty) : F(x) < 1 \text{ and } v_k(s) - v_k(x) \leq 1\} \\ &\leq \sup\{x \in [s, \infty) : F(x) < 1 \text{ and } v_{k-1}(s) - v_{k-1}(x) \leq 1\} \\ &= h_k(s), \end{aligned}$$

where the one inequality comes directly from the submodularity (2.16) and the two equalities come from (2.12). \square

Corollary 2.2

Proof. By the monotonicity (2.17) of the optimal threshold functions, the recursion (2.13) gives us the difference identity

$$\begin{aligned} v_{k+1}(s) - v_k(s) &= v_k(s) - v_{k-1}(s) \\ &+ \int_s^{h_{k+1}(s)} \{v_{k-1}(s) - v_{k-1}(x) - v_k(s) + v_k(x)\} dF(x) \\ &+ \int_{h_{k+1}(s)}^{h_k(s)} \{v_{k-1}(s) - 1 - v_{k-1}(x)\} dF(x), \end{aligned}$$

and it suffices to check that the two integrands on the right-hand side are non-positive. Non-positivity of the first integrand follows from the submodularity (2.16), and non-positivity of the second integrand follows from the characterization of $h_k(s)$ in (2.12). \square

Lemma 2.3

Proof. Set $Q(x, y) = 1 + v_{k-1}(x) - v_{k-1}(y)$. By our hypothesis, the partial derivatives of Q exist, and Lemma 2.1 implies that Q_y is strictly positive. Now, if (x_0, y_0) satisfies $Q(x_0, y_0) = 0$, then by the implicit function theorem there is a neighborhood \mathcal{N}_0 of x_0 where one can solve $Q(x, y) = 0$ uniquely for y , and the solution y is a differentiable

function of x for all $x \in \mathcal{N}_0$. Moreover, if $x < s_{k-1}^*$ then (2.14) tells us $Q(x, y) = 0$ if and only if $y = h_k(x)$, so h_k is differentiable as claimed. Given the differentiability of h_k at s , the formula (2.21) follows directly from $v_{k-1}(x) = v_{k-1}(h_k(x)) + 1$ by differentiation and the chain rule. The non-negativity of $h'_k(s)$ then follows because the value function v_{k-1} is strictly decreasing. \square

Proposition 2.3

Proof. We argue by induction on k , and we first note for $k = 1$ that $v_1(s) = 1 - F(s)$, so $v'_1(s) = -f(s)$ and (2.22) holds since $v_0(s) \equiv 0$. Next, we assume by induction that v_{k-1} is continuously differentiable on \mathcal{I} . If $s < s_{k-1}^*$ then the induction assumption and Lemma 2.3 imply that h_k is differentiable at s . We then differentiate (2.13) to find

$$\begin{aligned} v'_k(s) &= -f(s) + \{1 - F(h_k(s)) + F(s)\}v'_{k-1}(s) \\ &\quad + f(h_k(s))\{1 - v_{k-1}(s) + v_{k-1}(h_k(s))\}h'_k(s) \\ &= -f(s) + \{1 - F(h_k(s)) + F(s)\}v'_{k-1}(s), \end{aligned}$$

where the last step used the characterization (2.14) of h_k . Alternatively, if $s > s_{k-1}^*$ we have $F(h_k(s)) = 1$ and (2.13) says simply that

$$v_k(s) = F(s)v_{k-1}(s) + \int_s^\infty \{1 + v_{k-1}(x)\}f(x) dx.$$

Differentiation of this integral then gives us (2.22). Thus, one has that (2.22) holds on all of $\mathcal{I}_k = \mathcal{I} \setminus \{s_{k-1}^*\}$. Moreover, taking left and right limits in (2.22) gives us

$$\lim_{s \nearrow s_{k-1}^*} v'_k(s) = -f(s_{k-1}^*) + F(s_{k-1}^*)v'_{k-1}(s_{k-1}^*) = \lim_{s \searrow s_{k-1}^*} v'_k(s).$$

It is almost obvious that these relations imply the continuous differentiability v_k , but to make it crystal clear let γ be the common value of the limits above and define a

continuous function $\bar{v} : \mathcal{I} \rightarrow \mathbb{R}$ by setting

$$\bar{v}(s) = \begin{cases} v'_k(s) & \text{if } s < s_{k-1}^* \\ \gamma & \text{if } s = s_{k-1}^* \\ v'_k(s) & \text{if } s > s_{k-1}^*. \end{cases}$$

Next, we obtain by piecewise integration that

$$v_k(s) = v_k(0) + \int_0^s \bar{v}(u) du \quad \text{for all } s \in \mathcal{I},$$

implying, as expected, that v_k is continuously differentiable on \mathcal{I} . \square

Lemma 2.4

Proof. Proposition 2.3 tells us v_k^u is continuously differentiable on $(0, 1)$, and we prove concavity by showing that $s \mapsto (v_k^u)'(s)$ non-increasing on $(0, 1)$. We let \mathcal{H}_k be the assertion

$$(v_k^u)'(s + \epsilon) \leq (v_k^u)'(s) \quad \text{for all } s \in (0, 1) \text{ and } 0 < \epsilon < 1 - s,$$

and we argue by induction. For $k = 1$ we have $v_1^u(s) = 1 - s$, so $(v_1^u)'(s) = -1$ and \mathcal{H}_1 holds trivially.

Now, if we specialize the derivative recursion (2.22) to the uniform model we have

$$(v_k^u)'(s) = -1 + \{1 - h_k^u(s) + s\}(v_{k-1}^u)'(s), \quad \text{for } s \in (0, 1),$$

so if we assume that \mathcal{H}_{k-1} holds then we have

$$\begin{aligned} (v_k^u)'(s + \epsilon) - (v_k^u)'(s) &= \{1 - h_k^u(s) + s\} \{(v_{k-1}^u)'(s + \epsilon) - (v_{k-1}^u)'(s)\} \\ &\quad + \{h_k^u(s) - h_k^u(s + \epsilon) + \epsilon\} (v_{k-1}^u)'(s + \epsilon). \end{aligned} \quad (\text{A.3})$$

Since $0 \leq s \leq h_k^u(s) \leq 1$, we see from \mathcal{H}_{k-1} that the first summand on the right-hand side of (A.3) is non-positive. Monotonicity of v_k^u also tells us $(v_{k-1}^u)'(s + \epsilon) \leq 0$, so to complete the induction step we just need to check that

$$g(s, \epsilon) \stackrel{\text{def}}{=} h_k^u(s + \epsilon) - h_k^u(s) \leq \epsilon. \quad (\text{A.4})$$

From the definition of the critical value s_{k-1}^* we have

$$g(s, \epsilon) = \begin{cases} h_k^u(s + \epsilon) - h_k^u(s) & \text{if } s < s + \epsilon < s_{k-1}^* \\ 1 - h_k^u(s) & \text{if } s < s_{k-1}^* \leq s + \epsilon \\ 0 & \text{if } s_{k-1}^* \leq s < s + \epsilon, \end{cases}$$

so we only need to check (A.4) in the first two cases.

For $s < s + \epsilon < s_{k-1}^*$, we know by Lemma 2.3 that $h_k^u(s)$ is differentiable at s , so by the induction assumption \mathcal{H}_{k-1} and the negativity of $(v_{k-1}^u)'(s)$ we have

$$0 \leq (h_k^u)'(s) \leq 1 \quad \text{for all } s \in (0, s_{k-1}^*).$$

Thus, h_k^u is Lipschitz-1 continuous on $(0, s_{k-1}^*)$, and we have

$$g(s, \epsilon) = h_k^u(s + \epsilon) - h_k^u(s) \leq \epsilon \quad \text{for all } s < s + \epsilon < s_{k-1}^*.$$

For the second case where $s < s_{k-1}^* \leq s + \epsilon$, we first note that $h_k^u(s_{k-1}^*) = 1$ and that h_k^u is continuous, so we have

$$g(s, \epsilon) = \lim_{u \nearrow s_{k-1}^*} \{h_k^u(u) - h_k^u(s)\} \leq \lim_{u \nearrow s_{k-1}^*} \{u - s\} = s_{k-1}^* - s < \epsilon,$$

where the inequality follows from the Lipschitz-1 property of h_k^u . This completes the second check and the proof of the induction step. \square

Lemma 2.5

Proof. By Proposition 2.3 we know that v_k^e is continuously differentiable and by (2.22) we have

$$(v_k^e)'(s) = (1 - e^{-s} + e^{-h_k^e(s)})(v_{k-1}^e)'(s) - e^{-s}. \quad (\text{A.5})$$

We now let \mathcal{H}_k be the assertion that

$$-\{1 - e^{-h_{k+1}^e(s)+s}\}^{-1} \leq (v_k^e)'(s), \quad \text{for all } s \in (0, \infty),$$

and we argue induction. For $k = 1$ we have $v_1^e(s) = e^{-s} < 1$, so by (2.12) we have $h_2^e(s) = \infty$. In turn this gives us

$$-\{1 - e^{-h_2^e(s)+s}\}^{-1} = -1 \leq -e^{-s} = (v_1^e)'(s),$$

which verifies \mathcal{H}_1 .

Next, if we assume that \mathcal{H}_{k-1} holds and we substitute the lower bound from \mathcal{H}_{k-1} into (A.5), then rearrangement gives us

$$-\{1 - e^{-h_k^e(s)+s}\}^{-1}[1 - e^{-s}\{1 - e^{-h_k^e(s)+s}\}] - e^{-s} \leq (v_k^e)'(s).$$

From (2.17) we have $h_{k+1}^e(s) \leq h_k^e(s)$, so we now have

$$-\{1 - e^{-h_{k+1}^e(s)+s}\}^{-1} \leq -\{1 - e^{-h_k^e(s)+s}\}^{-1} \leq (v_k^e)'(s),$$

and this is just what one needs to complete the induction step. \square

Lemma 2.6

Proof. By Proposition 2.3 we know that v_k^e is continuously differentiable, and we again argue by induction. This time we take \mathcal{H}_k to be the assertion

$$(v_k^e)'(s) \leq (v_k^e)'(s + \epsilon) \quad \text{for all } s \in (0, \infty) \text{ and } \epsilon > 0.$$

For $k = 1$, we have $v_1^e(s) = e^{-s}$ and $(v_1^e)'(s) = -e^{-s}$ so the base case \mathcal{H}_1 of the induction is valid.

Now, by (A.5) applied twice we have

$$\begin{aligned} (v_k^e)'(s) - (v_k^e)'(s + \epsilon) &= [1 - e^{-s-\epsilon} + e^{-h_k^e(s+\epsilon)}]\{(v_{k-1}^e)'(s) - (v_{k-1}^e)'(s + \epsilon)\} \\ &\quad + \{e^{-h_k^e(s)}[1 - e^{-h_k^e(s+\epsilon)+h_k^e(s)}] - e^{-s}[1 - e^{-\epsilon}]\}(v_{k-1}^e)'(s) \\ &\quad - e^{-s}[1 - e^{-\epsilon}]. \end{aligned} \tag{A.6}$$

The induction hypothesis \mathcal{H}_{k-1} , tells us that $s \mapsto v_{k-1}^e(s)$ is convex, so by (2.21) we have $(h_k^e)'(s) \geq 1$ for $s \in (0, \infty)$, and this gives us the bound

$$-h_k^e(s + \epsilon) + h_k^e(s) \leq -\epsilon.$$

We always have $s \leq h_k^e(s)$ and $(v_{k-1}^e)'(s) \leq 0$, so (A.6) implies the simpler bound

$$\begin{aligned} (v_k^e)'(s) - (v_k^e)'(s + \epsilon) &\leq [1 - e^{-s-\epsilon} + e^{-h_k^e(s+\epsilon)}] \{ (v_{k-1}^e)'(s) - (v_{k-1}^e)'(s + \epsilon) \} \\ &\quad - e^{-s} [1 - e^{-\epsilon}] [1 - e^{-h_k^e(s)+s}] (v_{k-1}^e)'(s) \\ &\quad - e^{-s} [1 - e^{-\epsilon}]. \end{aligned}$$

We only need to check that this bound is non-positive. By the induction hypothesis \mathcal{H}_{k-1} and $s + \epsilon \leq h_k^e(s + \epsilon)$, we see the first term is non-positive. The bound (2.23) tells us

$$-[1 - e^{-h_k^e(s)+s}] (v_{k-1}^e)'(s) \leq 1,$$

so, when we replace $-[1 - e^{-h_k^e(s)+s}] (v_{k-1}^e)'(s)$ with its upper bound, we also see that the second and the third terms sum to zero. This completes the proof of the induction step and of the lemma. \square

A.2 Proofs in Section 2.3

Proposition 2.5

Proof. Specialization of the representation (2.36) gives us

$$w_{n-i}^u(M_i^u) = \sum_{j=i+1}^n \mathbb{E}[(A_j^u)^2 | \mathcal{F}_i] - \sum_{j=i+1}^n \mathbb{E}[(B_j^u)^2 | \mathcal{F}_i], \quad (\text{A.7})$$

where the definitions (2.33) and (2.34) now become

$$B_j^u = v_{n-j}^u(M_{j-1}^u) - v_{n-j+1}^u(M_{j-1}^u)$$

and

$$A_j^u = (1 + v_{n-j}^u(X_j) - v_{n-j}^u(M_{j-1}^u)) \mathbb{1}(X_j^u \in [M_{j-1}^u, h_{n-j+1}^u(M_{j-1}^u)]). \quad (\text{A.8})$$

First, we work toward a lower bound for the leading sum in (A.7). If we square both sides of (A.8) and take conditional expectations, then we have

$$\mathbb{E}[(A_j^u)^2 | \mathcal{F}_{j-1}] = \int_{M_{j-1}^u}^{h_{n-j+1}^u(M_{j-1}^u)} \{1 + v_{n-j}^u(x) - v_{n-j}^u(M_{j-1}^u)\}^2 dx. \quad (\text{A.9})$$

By Lemma 2.4 the map $x \mapsto v_{n-j}^u(x)$ is concave in x , so the line through the points $(M_{j-1}^u, 1)$ and $(h_{n-j+1}^u(M_{j-1}^u), 0)$ provides a lower bound on the integrand in (A.9). Integration of this linear lower bound then gives

$$\frac{1}{3} (h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u) \leq \int_{M_{j-1}^u}^{h_{n-j+1}^u(M_{j-1}^u)} \{1 + v_{n-j}^u(x) - v_{n-j}^u(M_{j-1}^u)\}^2 dx. \quad (\text{A.10})$$

From the definition of $L_i^u(\pi_n^*)$ we have the identity

$$\sum_{j=i+1}^n \mathbb{E}[h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u \mid \mathcal{F}_i] = \mathbb{E}[L_n^u(\pi_n^*) - L_i^u(\pi_n^*) \mid \mathcal{F}_i] = v_{n-i}^u(M_i^u),$$

so (A.9) and (A.10) give us

$$\frac{1}{3} v_{n-i}^u(M_i^u) \leq \sum_{j=i+1}^n \mathbb{E}[(A_j^u)^2 \mid \mathcal{F}_i]. \quad (\text{A.11})$$

Now, to work toward an upper bound on $\mathbb{E}[(B_j^u)^2 \mid \mathcal{F}_i]$, we first note by the crude Lemma 2.2 that

$$(B_j^u)^2 = (v_{n-j}^u(M_{j-1}^u) - v_{n-j+1}^u(M_{j-1}^u))^2 \leq (h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u)^2. \quad (\text{A.12})$$

The definition (2.24) of the running maximum M_j^u , the uniform distribution of X_j^u , and calculus give us the identity

$$\begin{aligned} \mathbb{E}[M_j^u - M_{j-1}^u \mid \mathcal{F}_{j-1}] &= \int_{M_{j-1}^u}^{h_{n-j+1}^u(M_{j-1}^u)} (x - M_{j-1}^u) dx \\ &= \frac{1}{2} (h_{n-j+1}^u(M_{j-1}^u) - M_{j-1}^u)^2, \end{aligned} \quad (\text{A.13})$$

so (A.12) gives us the succinct bound

$$(B_j^u)^2 \leq 2 \mathbb{E}[M_j^u - M_{j-1}^u \mid \mathcal{F}_{j-1}].$$

Now we take the conditional expectation with respect to \mathcal{F}_i and sum over $i < j \leq n$. Telescoping then gives us

$$\sum_{j=i+1}^n \mathbb{E}[(B_j^u)^2 \mid \mathcal{F}_i] \leq 2 \mathbb{E}[M_n^u - M_i^u \mid \mathcal{F}_i] \leq 2, \quad (\text{A.14})$$

where, in the last step, we used $0 \leq M_i^u \leq M_n^u \leq 1$. The representation (A.7) and the bounds (A.11) and (A.14) complete the proof of the lemma. \square

Lemma 2.7

Proof. If we set

$$g(y) \stackrel{\text{def}}{=} y^{-2} \{-6y + e^{-y}(2y^3 + 3y^2 - 6 + 6e^y)\} \quad \text{for } y \geq 0,$$

then by integration and simplification one has for $0 \leq s < t < \infty$ that

$$\frac{e^{-s}}{3} g(t-s) = \int_s^t \left(\frac{t-x}{t-s} \right)^2 e^{-x} dx - \frac{1}{3}(e^{-s} - e^{-t}) - \frac{2}{3}\{e^{-s} - e^{-t}(t-s+1)\},$$

and the lemma follows if we verify that $g(y) \leq 0$ for all $y \geq 0$. By the integral representation

$$g(y) = y^{-2} \left(\int_0^y (-6) dx + \int_0^y e^{-x}(6 + 6x + 3x^2 - 2x^3) dx \right),$$

we see that it suffices to show that

$$6 + 6x + 3x^2 - 2x^3 \leq 6e^x \quad \text{for all } x \in [0, \infty),$$

and the last inequality is obvious from the power series of e^x . \square

Proposition 2.7

Proof. If X_1, X_2, \dots, X_n is a sequence of independent random variables with the continuous distribution F , then the familiar transformations

$$X_i^u \stackrel{\text{def}}{=} F(X_i) \quad \text{and} \quad X_i^e \stackrel{\text{def}}{=} -\log\{1 - F(X_i)\}$$

define sequences that have the uniform and exponential distribution, respectively. These transformations give us a dictionary that we can use to translate results between our models; specifically we have:

$$\begin{aligned} v_k(s) &= v_k^u(F(s)) & \text{and} & & v_k(s) &= v_k^e(-\log\{1 - F(s)\}), \\ M_i^u &= F(M_i) & \text{and} & & M_i^e &= -\log\{1 - F(M_i)\}, \\ w_{n-i}(M_i) &= w_{n-i}^u(M_i^u) & \text{and} & & w_{n-i}(M_i) &= w_{n-i}^e(M_i^e). \end{aligned}$$

Proposition 2.5 tells us that

$$\frac{1}{3} v_{n-i}^u(M_i^u) - 2 \leq w_{n-i}^u(M_i^u),$$

so the first column of the dictionary gives us the first inequality of (2.50). Similarly, Proposition 2.6 tells us

$$w_{n-i}^e(M_i^e) \leq \frac{1}{3} v_{n-i}^e(M_i^e) + \frac{2}{3} \{1 + \log(n - i)\},$$

and the second column of the dictionary gives us the second inequality of (2.50). \square

Lemma 2.8

Proof. We fix $n \geq 1$ and simplify the notation by dropping the subscript n on the martingale difference sequence and the filtration. We then let

$$V \stackrel{\text{def}}{=} \sum_{j=1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}]$$

and consider the martingale $\{V_i : 0 \leq i \leq n\}$ defined by setting

$$V_i \stackrel{\text{def}}{=} \mathbb{E}[V \mid \mathcal{F}_i] \quad \text{for } 0 \leq i \leq n.$$

One has the initial and terminal values

$$V_0 = \sum_{j=1}^n \mathbb{E}[d_j^2] = \text{Var}[L_n(\pi_n^*)] \quad \text{and} \quad V_n = V,$$

and if we introduce the new martingale differences $\Delta_i = V_i - V_{i-1}$, $1 \leq i \leq n$, then telescoping and orthogonality give us

$$V_n - V_0 = \sum_{i=1}^n \Delta_i \quad \text{and} \quad \text{Var}[V_n] = \text{Var}\left[\sum_{j=1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}]\right] = \sum_{i=1}^n \mathbb{E}[\Delta_i^2]. \quad (\text{A.15})$$

For $1 \leq j \leq i+1$ all of the summands $\mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}]$ are \mathcal{F}_i -measurable, so we have

$$\begin{aligned} \Delta_i &= \sum_{j=1}^i \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}] + \mathbb{E}\left[\sum_{j=i+1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}] \mid \mathcal{F}_i\right] \\ &\quad - \sum_{j=1}^i \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}] - \mathbb{E}\left[\sum_{j=i+1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}] \mid \mathcal{F}_{i-1}\right]. \end{aligned}$$

The first and the third sum cancel, and we obtain

$$\begin{aligned} \Delta_i &= \sum_{j=i+1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_i] - \mathbb{E}\left[\sum_{j=i+1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_i] \mid \mathcal{F}_{i-1}\right] \\ &= w_{n-i}(M_i) - \mathbb{E}[w_{n-i}(M_i) \mid \mathcal{F}_{i-1}], \end{aligned} \quad (\text{A.16})$$

where in the last line we twice used the formula (2.31) for the conditional variance.

Next, we set

$$G_i \stackrel{\text{def}}{=} \{\omega : X_i(\omega) \in [M_{i-1}(\omega), h_{n-i+1}(M_{i-1}(\omega))]\},$$

so, in words, G_i is the set of all ω for which the observation $X_i(\omega)$ is selected at time i under the optimal policy π_n^* . By the recursive definition (2.24) of the running maximum M_i , we then have the decomposition

$$w_{n-i}(M_i) = w_{n-i}(M_{i-1}) + \{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i). \quad (\text{A.17})$$

In fact, one can replace w_{n-i} with any function here, and it will be useful to also note that

$$v_{n-i}(M_i) = v_{n-i}(M_{i-1}) + \{v_{n-i}(X_i) - v_{n-i}(M_{i-1})\} \mathbb{1}(G_i). \quad (\text{A.18})$$

The first summand on the right-hand side of (A.17) is \mathcal{F}_{i-1} -measurable, so, if we rewrite (A.16) using (A.17) we obtain

$$\Delta_i = \{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i) - \mathbb{E}[\{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i) \mid \mathcal{F}_{i-1}].$$

When we square this identity and take the conditional expectation we find

$$\mathbb{E}[\Delta_i^2 \mid \mathcal{F}_{i-1}] \leq \mathbb{E}[\{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\}^2 \mathbb{1}(G_i) \mid \mathcal{F}_{i-1}], \quad (\text{A.19})$$

and all that remains is to estimate the difference $\{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i)$.

Now consider the *upper bound* in (2.50) and replace $w_{n-i}(M_i)$ and $v_{n-i}(M_i)$ with their decompositions (A.17) and (A.18). When we move the term $w_{n-i}(M_{i-1})$ to the right side, we have

$$\begin{aligned} \{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i) &\leq \frac{1}{3} \{v_{n-i}(X_i) - v_{n-i}(M_{i-1})\} \mathbb{1}(G_i) & (\text{A.20}) \\ &+ \frac{1}{3} v_{n-i}(M_{i-1}) - w_{n-i}(M_{i-1}) \\ &+ \frac{2}{3} (1 + \log n). \end{aligned}$$

By the *lower bound* in (2.50) the second summand is bounded by two, and, to estimate the first summand, we note that the characterization (2.12) for the optimal threshold function and the monotonicity of the value function give us

$$\frac{1}{3} |v_{n-i}(X_i) - v_{n-i}(M_{i-1})| \mathbb{1}(G_i) \leq \frac{1}{3} \mathbb{1}(G_i). \quad (\text{A.21})$$

The left-hand side of (A.20) is zero off of the set G_i , so using (A.21) we see that (A.20) gives us

$$\{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i) \leq \{3 + \frac{2}{3} \log n\} \mathbb{1}(G_i). \quad (\text{A.22})$$

A parallel argument gives us the complementary inequality,

$$-\{3 + \frac{2}{3} \log n\} \mathbb{1}(G_i) \leq \{w_{n-i}(X_i) - w_{n-i}(M_{i-1})\} \mathbb{1}(G_i). \quad (\text{A.23})$$

Specifically, one now begins with the *lower bound* in (2.50) and replaces $w_{n-i}(M_i)$ and $v_{n-i}(M_i)$ with their decompositions.

Taken together (A.22) and (A.23) imply

$$|w_{n-i}(X_i) - w_{n-i}(M_{i-1})| \mathbb{1}(G_i) \leq \left\{3 + \frac{2}{3} \log n\right\} \mathbb{1}(G_i),$$

so we can square both sides and take conditional expectations with respect to \mathcal{F}_{i-1} .

By (A.19) and the definition of G_i , we then have

$$\mathbb{E}[\Delta_i^2 \mid \mathcal{F}_{i-1}] \leq \left\{18 + \frac{8}{9}(\log n)^2\right\} \mathbb{E}[\mathbb{1}(X_i \in [M_{i-1}, h_{n-i+1}(M_{i-1})]) \mid \mathcal{F}_{i-1}],$$

so if we drop the factor 8/9, take total expectations, and sum we get

$$\sum_{i=1}^n \mathbb{E}[\Delta_i^2] \leq \{18 + (\log n)^2\} \sum_{i=1}^n \mathbb{E}[\mathbb{1}(X_i \in [M_{i-1}, h_{n-i+1}(M_{i-1})])]. \quad (\text{A.24})$$

By (A.15) the sum on the left is variance of $V_n = \sum_{j=1}^n \mathbb{E}[d_j^2 \mid \mathcal{F}_{j-1}]$, and by (2.25) the sum of the expected values on the right is equal to $\mathbb{E}[L_n(\pi_n^*)]$. Finally, we know that $\mathbb{E}[L_n(\pi_n^*)] < (2n)^{1/2}$ from (2.3) and the argument of Section 2.3.3, so (A.24) completes the proof of the lemma. \square

Appendix B

Proofs in Chapter 3

B.1 Proof in Section 3.3

Example 3.1

Proof. Bidder 1's payoff function is:

$$\pi_1(b_1, b_2) = \begin{cases} v_1 - b_1 & \text{if } b_1 \geq b_2 \\ 0 & \text{otherwise.} \end{cases}$$

Since b_2 is uniform on $[c, d]$, bidder 1's expected payoff as a function of b_1 is:

$$\mathbb{E}[\pi_1(b_1)] = \begin{cases} v_1 - b_1 & \text{if } d < b_1 \\ \frac{(v_1 - b_1)(b_1 - c)}{d - c} & \text{if } c \leq b_1 \leq d \\ 0 & \text{otherwise.} \end{cases}$$

By solving first-order condition, the maximizer of $\mathbb{E}[\pi_1(b_1)]$ is:

$$b_1^{EM} = \begin{cases} d & \text{if } 2d - c < v_1 \\ \frac{1}{2}(v_1 + c) & \text{if } c \leq v_1 \leq 2d - c \\ 0 & \text{otherwise.} \end{cases}$$

□

Example 3.3

Proof. Using (3.2) and the fact that b_2 and b_3 are independent and uniformly distributed, we have:

$$\mathbb{E}[\pi_1(b_1)] = \begin{cases} 0 & \text{if } b_1 < 4 \\ 0.7(8 - \frac{1}{2}b_1^2 - 4v_1 + b_1v_1) & \text{if } 4 \leq b_1 < 5 \\ 0.7(v_1 - \frac{9}{2}) & \text{if } 5 \leq b_1 < 6 \\ 18 - \frac{1}{2}b_1^2 - 6v_1 + b_1v_1 & \text{if } 6 < b_1 \leq 7 \\ v_1 - \frac{13}{2} & \text{if } 7 < b_1. \end{cases} \quad (\text{B.1})$$

One can verify that the maximizer of $\mathbb{E}[\pi_1(b_1)]$ is given by (3.21). \square

Proposition 3.1

Proof. If $b_1 < u(m)$, there exists $b_{-1} \in \mathcal{U}_{-1}$ such that $b_1 < b^{(m)}$ so that bidder 1 loses the auction and receives zero payoff. Thus, bidder 1's worst-case payoff is no more than zero if she bids $b_1 < u(m)$. On the other hand, if $b_1 \geq u(m)$ then bidder 1 always wins an item and receives a payoff $\pi_1(b_1, b_{-1}) = v_1 - b_1$. Hence, her optimal policy is to bid exactly $u(m)$ if $u(m) \leq v_1$ and bid zero if $v_1 < u(m)$. In other words, a solution to (RO) is $b_1^{RO} = u(m)\mathbb{1}_{u(m) \leq v_1}$. The optimal worst-case payoff is thus $\pi_1^{MAXMIN} = (v_1 - u(m))^+$. \square

Lemma 3.1

Proof. By definition, we have

$$\begin{aligned} f_k(x) &= \inf_{b_{-1}} \alpha_k(v_1 - b^{(k)}) \\ \text{s.t. } & b^{(k)} \leq x < b^{(k-1)} \\ & b_{-1} \in \mathcal{U}_{-1}. \end{aligned} \quad (\text{B.2})$$

When $l(k) \leq x < u(k)$, the constraint $b^{(k)} \leq x$ is always binding and thus we have $f_k(x) = \alpha_k(v_1 - x)$. If $u(k) \leq x < u(k-1 | k)$ then bidder 1 wins item k for sure

by bidding x and it is possible for rivals to bid in such a way that bidder 1 pays the maximum possible amount $u(k)$. Thus, $f_k(x) = \alpha_k(v_1 - u(k))$ in this case. Finally, if $u(k-1 | k) \leq x < u(k-1)$ then the constraint $b^{(k)} \leq x$ is never binding. Therefore, $f_k(x) = \alpha_k(v_1 - u(k | k-1, x))$. \square

Proposition 3.2

Proof. We first prove the proposition for the case of $m = 2$. Without loss of generality, we can assume that $\alpha = 1$. If bidder 1 bids $b_1 < u(2)$ then there is always a possibility that she loses the auction, so her worst-case payoff in this case is zero. On the other hand, if $b_1 \geq u(2)$ then by Lemma 3.1, the worst-case payoff for bidder 1 if she wins the first slot is at most $v_1 - u(2)$, whereas her worst-case payoff should she win the second slot is at least $v_1 - u(2)$ (see Figure 3.4). In other words, $f_1(b_1) \leq v_1 - u(2) \leq f_2(b_1), \forall b_1 \geq u(2)$. Thus, it is optimal for bidder 1 to bid $b_1^{RO} = u(2)\mathbb{1}_{u(2) \leq v_1}$. The worst-case payoff is therefore $\pi_1^{MAXMIN} = (v_1 - u(2))^+$. The same argument extends directly for $m > 2$. \square

Proposition 3.3

Proof. When $\alpha_1(v_1 - u(1 | 2)) \leq \alpha_2(v_1 - u(2))$, according to Lemma 3.1, bidder 1's worst-case payoff is at most $\alpha_2(v_1 - u(2))$ (see Figure 3.5). Thus, a robust policy is $b_1^{RO} = u(2)\mathbb{1}_{u(2) \leq v_1}$ and the optimal worst-case payoff is $\pi_1^{MAXMIN} = \alpha_2(v_1 - u(2))^+$. On the other hand, if $\alpha_2(v_1 - u(2 | 1)) \leq \alpha_1(v_1 - u(1))$ then from Figure 3.6 we can see that bidder 1 is always better off by winning item 1 than by winning item 2 (in terms of worst-case payoff). Thus, $b_1^{RO} = u(1)\mathbb{1}_{u(1) \leq v_1}$ and $\pi_1^{MAXMIN} = \alpha_1(v_1 - u(1))^+$. Finally, when $\alpha_2(v_1 - u(2)) < \alpha_1(v_1 - u(1 | 2))$ and $\alpha_1(v_1 - u(1)) < \alpha_2(v_1 - u(2 | 1))$, bidder 1's worst-case payoff function on $x \in [u(2), \infty)$ is maximized at the intersection of $y = \alpha_1(v_1 - x)$ and $y = \alpha_2(v_1 - u(2 | 1, x))$ (Figure 3.7). Thus, $b_1^{RO} = x^*\mathbb{1}_{x^* \leq v_1}$. The optimal worst-case payoff in this case is therefore $\pi_1^{MAXMIN} = \alpha_1(v_1 - x^*)^+$. \square

B.2 Proofs in Section 3.4

Proposition 3.5

Proof. Since VCG is in the core if the coalition value function is bidder-submodular, it suffices to show this bidder-submodularity property. For any coalition S , let $b_S(M) = \max_{j \in S} b_j(M)$. Then due to supermodularity, we have $w_b(S) = b_S(M)$ and $w_b(S \cup l) = \max(b_S(M), b_l(M))$. Therefore, $w_b(S \cup l) - w_b(S) = \max(b_S(M), b_l(M)) - b_S(M) = \max(0, b_l(M) - b_S(M))$. As a result, if $0 \in S \subset S'$ then $\max(0, b_l(M) - b_S(M)) \geq \max(0, b_l(M) - b_{S'}(M))$ so that $w_b(S \cup l) - w_b(S) \geq w_b(S' \cup l) - w_b(S')$. Thus, $w_b(\cdot)$ is bidder-submodular. \square

Proposition 3.6

Proof. For any $b_1 \in \mathcal{U}_1$ and $b_{-1} \in \mathcal{U}_{-1}$, since $w_{v_1, b_{-1}}$ is bidder-submodular, we have that $\pi_1(b_1, b_{-1}) \leq \pi_1(v_1, b_{-1})$ (Ausubel and Milgrom, 2006). Thus,

$$\inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1, b_{-1}) \leq \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(v_1, b_{-1}),$$

which implies

$$\sup_{b_1 \in \mathcal{U}_1} \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1, b_{-1}) \leq \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(v_1, b_{-1}).$$

As a result, truthful reporting is the optimal solution to (RO). \square

Proposition 3.7

Proof. Let S_1 be the bundle that bidder 1 wins when she bids truthfully. By definition, we have

$$w_{v_1, b_{-1}}(N) = v_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1).$$

We observe that by using policy (3.24), bidder 1 also wins S_1 . In fact, the maximum reported valuation generated by allocating S_1 to bidder 1 and $M \setminus S_1$ to the remaining

bidders is

$$\begin{aligned} v_1(S_1) - \pi_1^{VCG} + w_{b_{-1}}(N \setminus 1, M \setminus S_1) &= w_{v_1, b_{-1}}(N) - (w_{v_1, b_{-1}}(N) - w_{b_{-1}}(N \setminus 1)) \\ &= w_{b_{-1}}(N \setminus 1), \end{aligned}$$

which is equal to the maximum reported valuation generated by allocating items in M to the remaining bidders. Under the assumed tie-breaking rule, S_1 is her allocation outcome. The VCG payoffs with respect to the reported valuation profile (b_1, b_{-1}) is

$$\pi_j^{VCG} = w_b(N) - w_b(N \setminus j) = w_{b_{-1}}(N \setminus 1) - w_{b_{-1}}(N \setminus 1) = 0,$$

for $j \in N$, and

$$\pi_0^{VCG} = w_b(N) - \sum_{j \in N} \pi_j^{VCG} = w_{b_{-1}}(N \setminus 1).$$

This VCG profile is in the core with respect to the reported bids b so bidder 1 is charged exactly her VCG payment p_1^{VCG} , which is a function of b_{-1} only. As a result, bidder 1 gets

$$\pi_1 = v_1(S_1) - p_1^{VCG} = v_1(S_1) - (v_1(S_1) - \pi_1^{VCG}) = \pi_1^{VCG},$$

which is her VCG payoff with respect to the reported valuation (v_1, b_{-1}) . Since π_1^{VCG} is the maximum payoff that bidder 1 can get, policy (3.24) is optimal. \square

Remark 3.1

Proof. Let $C \subseteq N$ be the set of bidders corresponding to bidder 1's skills. Also, given a bid profile b , let S_j be the bundle allocated to bidder j . We have

$$\begin{aligned} \pi_0 + \sum_{j \in N \setminus C} \pi_j &= \sum_{j \in N} p_j + \sum_{j \in N \setminus C} (b_j(S_j) - p_j) \\ &= \sum_{j \in C} p_j + \sum_{j \in N \setminus C} b_j(S_j) \\ &\leq \sum_{j \in C} p_j + w_b(N \setminus C, M \setminus \cup_{j \in C} S_j), \end{aligned}$$

where the last inequality follows from the definition of $w_b(\cdot)$. From the core constraints, we have $w_b(N \setminus S) \leq \pi_0 + \sum_{j \in N \setminus C} \pi_j$. Hence,

$$p_1^{VCG} = w_b(N \setminus C) - w_b(N \setminus C, M \setminus \cup_{j \in C} S_j) \leq \sum_{j \in C} p_j.$$

Thus, bidder 1 pays at least p_1^{VCG} even when she uses shills, which implies that her payoff is at most π_1^{VCG} . Since bidding policy (3.24) gives bidder 1 this VCG payoff, it is also optimal in the extended policy space in which bidder 1 uses shills. \square

Proposition 3.8

Proof. First observe that $\sum_{j \in N} (v_j(S_j) - \gamma_j^{EQ}) = w_v(N) - \sum_{j \in N} \gamma_j^{EQ} = \gamma_0^{EQ}$, so when everyone bids according to (3.25), the maximum total reported valuation is γ_0^{EQ} and the efficient allocation is $\{S_j\}_{j \in N}$ (under the tie-breaking rule favoring such allocation). The VCG payoff for each bidder j with respect to this bidding profile is $\pi_j^{VCG} = w_b(N) - w_b(N \setminus j) = \gamma_0^{EQ} - \gamma_0^{EQ} = 0$. This VCG payoff profile is in the core, so each bidder is charged exactly $p_j^{VCG} = v_j(S_j) - \gamma_j^{EQ}$. As a result, the actual payoff that bidder j receives is $\gamma_j = v_j(S_j) - p_j^{VCG} = \gamma_j^{EQ}$.

Next, we show that each bidder j earns no more than γ_j^{EQ} when her rivals bid according to (3.25). Under bidding profile (3.25), no optimal allocation yields a total reported valuation more than γ_0^{EQ} . In fact, assume the opposite and let C be the smallest coalition such that the total reported value on assigned bundles exceed γ_0^{EQ} . We have that

$$\gamma_0^{EQ} \leq \max_{x: \cup_{j \in C} x_j = M} \sum_{j \in C} b_j(x_j) \leq \sum_{j \in C} (v_j(S_j) - \gamma_j^{EQ}) = w_v(C) - \sum_{j \in C} \gamma_j^{EQ},$$

so the set C can form a blocking coalition with the seller, which is a contradiction to the fact that γ^{EQ} satisfies the core constraints. As a result, for every bidder j , the coalition value among the rivals $w_b(N \setminus j)$ cannot exceed γ_0^{EQ} . If bidder j deviates then

either she wins all items, in which case the resulting total valuation is still no more than γ_0^{EQ} , or she wins a subset of M and contributes at most an additional γ_j^{EQ} to the total valuation on assigned bundles. Thus, we have $w_{b'_j, b_{-j}}(N) \leq \gamma_0^{EQ} + \gamma_j^{EQ}$ for any alternative reporting policy b'_j . But this implies that by reporting b'_j , bidder j receives a payoff that can be no more than $w_{b'_j, b_{-j}}(N) - w_{b'_j, b_{-j}}(N \setminus j) \leq (\gamma_0^{EQ} + \gamma_j^{EQ}) - \gamma_0^{EQ} = \gamma_j^{EQ}$. Thus, b_j is indeed bidder j 's best reply and the equilibrium results follow directly. \square

Proposition 3.9

Proof. Let b_1 be the optimal bidding policy defined by (3.24). For any $\alpha \in [0, 1]$, we define b_1^α to be a linear combination of the true valuation and the optimal policy b_1 :

$$b_1^\alpha(S) = (1 - \alpha)v_1(S) + \alpha b_1(S), \quad \forall S \subseteq M.$$

We observe that when bidder 1 bids according to b_1^α and others bid truthfully, the optimal allocation is the same with that of truthful bidding. In fact, if S_1 is bidder 1's truthful bundle, then for any $S \subseteq M$, we have

$$v_1(S) + w_{b_{-1}}(N \setminus 1, M \setminus S) \leq v_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1). \quad (\text{B.3})$$

Bidder 1 also wins S_1 under her optimal bidding policy b_1 :

$$b_1(S) + w_{b_{-1}}(N \setminus 1, M \setminus S) \leq b_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1). \quad (\text{B.4})$$

Multiplying (B.3) by $(1 - \alpha)$ and (B.4) by α , and summing up yields

$$b_1^\alpha(S) + w_{b_{-1}}(N \setminus 1, M \setminus S) \leq b_1^\alpha(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1). \quad (\text{B.5})$$

As a result, the allocation outcome for bidder 1 does not change under the bidding profile (b_1^α, v_{-1}) . Furthermore, by bidding b_1^α , bidder 1 reduces her reported valuation for the winning bundle by $\alpha \pi_1^{VCG}$. Consequently, all the winners together have to pay

an extra amount of $\alpha\pi_1^{VCG}$ to overcome blocking coalitions. Under the nearest-VCG rule, this amount is split equally among the winners, and bidder 1 only has to pay a fraction of the quantity $\alpha\pi_1^{VCG}$. Therefore, her payoff under b_1^α is strictly larger than that of truthful bidding. This holds for any $\alpha > 0$ so we can choose $b_1^\epsilon = b_1^\alpha$ for a sufficiently small α so that $\|v_1 - b_1^\epsilon\|_2 \leq \epsilon$. \square

Proposition 3.10

Proof. We first show that bidder 1's worst-case payoff under policy (3.28) is at least $v_1(S_1) - \xi$. By construction, we have that

$$b_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1) \geq b_1(M), \quad \forall b_{-1} \in \mathcal{U}_{-1}. \quad (\text{B.6})$$

Furthermore, from the definition of ξ , we have

$$b_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1) \geq w_{b_{-1}}(N \setminus 1, M), \quad \forall b_{-1} \in \mathcal{U}_{-1}. \quad (\text{B.7})$$

Since $b(S) = 0$ for all $S \not\supseteq S_1$ and $w_{b_{-1}}(N \setminus 1, M \setminus S) \leq w_{b_{-1}}(N \setminus 1, M)$, the above inequality also implies that

$$b_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1) > b_1(S) + w_{b_{-1}}(N \setminus 1, M \setminus S), \quad (\text{B.8})$$

for all $S \not\supseteq S_1$ and $b_{-1} \in \mathcal{U}_{-1}$. The inequalities (B.6), (B.7) and (B.8) jointly show that it is always optimal for the seller to allocate bundle S_1 to bidder 1 and the rest of the items to other bidders. Since a bidder's payment can never exceed her bid, bidder 1's payoff under policy (3.28) is at least $v_1(S_1) - \xi$. Hence, the optimal worst-case payoff under robust bidding satisfies $\pi_1^{MAXMIN} \geq v_1(S_1) - \xi$. On the other hand, recall that we have the minimax inequality (3.12), so $\pi_1^{MAXMIN} \leq \pi_1^{MINMAX}$. For any realization of $b_{-1} \in \mathcal{U}_{-1}$, bidder 1 can response optimally by bidding according to a perfect-information optimal policy, e.g., policy (3.24). Under such policy, bidder

1 receives her VCG payoff, so we have

$$\begin{aligned}
\pi_1^{MINMAX} &= \min_{b_{-1} \in \mathcal{U}_{-1}} (w_b(N, M) - w_b(N \setminus 1, M)) \\
&= \min_{b_{-1} \in \mathcal{U}_{-1}} (v_1(S_1) + w_b(N \setminus 1, M \setminus S_1) - w_b(N \setminus 1, M)) \\
&= v_1(S_1) - \max_{b_{-1} \in \mathcal{U}_{-1}} (w_b(N \setminus 1, M) - w_b(N \setminus 1, M \setminus S_1)) \\
&= v_1(S_1) - \xi.
\end{aligned}$$

Since $v_1(S_1) - \xi \leq \pi_1^{MAXMIN} \leq \pi_1^{MINMAX} = v_1(S_1) - \xi$, we must have that $\pi_1^{MAXMIN} = \pi_1^{MINMAX} = v_1(S_1) - \xi$ and the bidding policy (3.28) is optimal. \square

Remark 3.3

Proof. When bidders 1 and 2 win exactly one item each, their VCG payments are $p_1^{VCG} = \max(0, c - b)$ and $p_2^{VCG} = \max(y - x, c - x)$ and the core constraints are $p_1 \leq x$, $p_2 \leq b$ and $p_1 + p_2 \geq c$. The projection of VCG on the core gives us bidder 1's payment:

$$p_1 = \left(x - \frac{1}{2} \min(x, x + b - c) + \frac{1}{2} \min(c - y, 0) \right)^+. \quad (\text{B.9})$$

The right hand side of (B.9) corresponds to projecting the VCG payment onto the blocking constraint $p_1 + p_2 \geq c$ created by bidder 3's bid on the global bundle. This term can also be written as

$$\max(0, c - b) + \frac{1}{2}(c - \max(0, c - b) - \max(y - x, c - x)),$$

which is bidder 1's VCG payment plus an extra amount that is half of the total surcharge that bidder 1 and 2 together have to pay to overcome the blocking global bidder.¹ However, such projection does not always yield a non-negative payment, so

¹ See e.g., Goeree and Lien (2016).

it is truncated at zero as in (B.9). By substituting $(x, y) = (\xi, \xi + \bar{b} - \epsilon_b)$ into (B.9), we get bidder 1's payment under the robust policy:

$$p_1^{RO} = \left(\xi - \frac{1}{2} \min(\xi, \xi + b - c) + \frac{1}{2} \min(c - \xi - \bar{b} + \epsilon_b, 0) \right)^+ \quad (\text{B.10})$$

Bidder 1's payoff under the robust policy is thus given by

$$\pi_1^{RO} = \min \left(a + \frac{1}{2} \min(0, b - c) - \frac{1}{2} \min(c - \bar{b} + \epsilon_b, \xi), a \right). \quad (\text{B.11})$$

On the other hand, by bidding truthfully bidder 1 gets

$$\pi_1^{TR} = \min \left(a + \frac{1}{2} \min(0, b - c) - \frac{1}{2} \min(c, a), a \right). \quad (\text{B.12})$$

Since $\min(c - \bar{b} + \epsilon_b, \xi) = c - \bar{b} + \epsilon_b \leq \min(c, a)$, we have that $\pi_1^{RO} \geq \pi_1^{TR}$ for all realization of $b_{-1} \in \mathcal{U}_{-1}$. \square

Example 3.5

Proof. Bidder 1's payoff function is:

$$\pi_1(x, y, b, c) = \begin{cases} 0 & \text{if } x + b < c \text{ and } y < c \\ \min \left(a, a - x + \frac{1}{2} \min(x, x + b - c) - \frac{1}{2} \min(c - y, 0) \right) & \text{if } c \leq x + b \text{ and } y \leq x + b \\ a - \max(b, c) & \text{if } x + b < y \text{ and } c \leq y. \end{cases} \quad (\text{B.13})$$

Since b and c are independent and distributed according to f_b and f_c given by (3.30) and (3.31), we have:

$$\begin{aligned} & \mathbb{E}[\pi_1(x, y, b, c)] \\ &= \frac{1}{324} \int_{\max(y-x, \bar{b}-\epsilon_b)}^{\bar{b}+\epsilon_b} \int_{\bar{c}-\epsilon_c}^{\min(x+b, \bar{c}+\epsilon_c)} \pi_1^1(x, y, b, c) (b-7)(13-c) \, dc \, db \\ &+ \frac{1}{324} \int_{\bar{b}-\epsilon_b}^{\min(y-x, \bar{b}+\epsilon_b)} \int_{\bar{c}-\epsilon_c}^{\min(y, \bar{c}+\epsilon_c)} \pi_1^2(b, c) (b-7)(13-c) \, dc \, db, \end{aligned}$$

where $\pi_1^1(x, y, b, c) = \left(a, a - x + \frac{1}{2} \min(x, x + b - c) - \frac{1}{2} \min(c - y, 0) \right)$ and $\pi_1^2(b, c) = a - \max(b, c)$. The expected-payoff maximizing policy b_1^{EM} can then be obtained by choosing (x, y) that maximizes $\mathbb{E}[\pi_1(x, y, b, c)]$ over \mathbb{R}_+^2 . \square

Proposition 3.11

Proof. (a) Bidder 1 has unique truthful allocation S_1 if and only if, for all $b_{-1} \in \mathcal{U}_{-1}$:

$$v_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1) \geq v_1(S_2) + w_{b_{-1}}(N \setminus 1, M \setminus S_2), \quad (\text{B.14})$$

and

$$v_1(S_1) + w_{b_{-1}}(N \setminus 1, M \setminus S_1) \geq w_{b_{-1}}(N \setminus 1, M), \quad \forall b_{-1} \in \mathcal{U}_{-1}. \quad (\text{B.15})$$

We will prove that similar to the single-minded case, the worst-case payoff of bidder 1 under the robust policy b_1^{RO} is at least as large as the upper bound provided by minimax inequality and thus b_1^{RO} must be optimal. First, given policy b_1^{RO} , we can partition uncertainty set \mathcal{U}_{-1} into two disjoint subsets $\mathcal{U}_{-1}^{(1)}$ and $\mathcal{U}_{-1}^{(2)}$ such that bidder 1 always wins S_i on the set $\mathcal{U}_{-1}^{(i)}$ for $i \in \{1, 2\}$. The worst-case payoff under policy b_1^{RO} is given by

$$\inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1^{RO}, b_{-1}) = \min \left(\inf_{b_{-1} \in \mathcal{U}_{-1}^{(1)}} (v_1(S_1) - p_1^{(1)}), \inf_{b_{-1} \in \mathcal{U}_{-1}^{(2)}} (v_1(S_2) - p_1^{(2)}) \right),$$

where $p_1^{(1)}$ and $p_1^{(2)}$ denote the corresponding payments for bidder 1. Note that by individual rationality, we have $p_1^{(1)} \leq \xi_1$ and $p_1^{(2)} \leq \xi_1$. In addition, $v_1(S_1) \leq v_1(S_2)$, so we have

$$\inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1^{RO}, b_{-1}) \geq v_1(S_1) - \xi_1.$$

Since $\pi_1^{MAXMIN} \geq \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1^{RO}, b_{-1})$, we have $\pi_1^{MAXMIN} \geq v_1(S_1) - \xi_1$. On the other hand, if we let $p_1^{VCG}(S_1)$ and $p_1^{VCG}(S_2)$ be bidder 1's VCG payment when she

wins S_1 and S_2 , respectively, then the minmax payoff is

$$\begin{aligned}\pi_1^{MINMAX} &= \inf_{b_{-1} \in \mathcal{U}_{-1}} \max(v_1(S_1) - p_1^{VCG}(S_1), v_1(S_2) - p_1^{VCG}(S_2)) \\ &= v_1(S_1) - \xi_1,\end{aligned}$$

where the last equality follows directly from the condition (B.14) that guarantees bidder 1 always win bundle S_1 under truthful bidding. Finally, by minimax inequality, we have $\pi_1^{MAXMIN} \leq \pi_1^{MINMAX}$. As a result, the minimax equality (3.13) must hold and b_1^{RO} is the optimal solution to (RO).

(b) We observe that under policy b_1^{RO} defined in (3.36), bidder 1 always wins S_2 . As a result, the optimality follows directly from minimax inequality argument analogous to the proof of Proposition 3.10. □

Lemma 3.2

When bidder 1 wins one item, her payment is given by (B.9). Recall that when $(x, y) \in \mathcal{U}'_1$, we have $x = \xi_1$. As a result, $\pi_1^{WO,1}(y)$ is given by

$$\begin{aligned}\pi_1^{WO,1}(y) &= \inf \min\{a, a - \xi_1 + \frac{1}{2} \min(\xi_1, \xi_1 + b - c) + \frac{1}{2}(y - c)^+\} \\ \text{s.t. } & b_{-1} \in \mathcal{U}_{-1} \\ & y \leq \xi_1 + b.\end{aligned} \tag{B.16}$$

The objective function on the right hand side is increasing in b and decreasing in c so the worst-case scenario corresponds to $b^* = \max(\bar{b} - \epsilon_b, y - \xi_1)$ and $c^* = \bar{c} + \epsilon_c$. Substituting these values into (B.16), we have

$$\pi_1^{WO,1}(y) = \min\{a, a - \xi_1 + \frac{1}{2} \min(\xi_1, (y - \bar{c} - \epsilon_c)^+) + \frac{1}{2}(y - \bar{c} - \epsilon_c)^+\}. \tag{B.17}$$

Similarly, when bidder 1 wins both items, her payment is $\max(b, c)$, so $\pi_1^{WO,2}(y)$ is given by

$$\begin{aligned}\pi_1^{WO,2}(y) &= \inf \quad \{a' - \max(b, c)\} \\ \text{s.t.} \quad & b_{-1} \in \mathcal{U}_{-1} \\ & \xi_1 + b < y.\end{aligned}\tag{B.18}$$

In this case, the worst-case bids are $b^* = \min(\bar{b} + \epsilon_b, y - \xi_1)$, $c^* = \bar{c} + \epsilon_c$ so we have

$$\pi_1^{WO,2}(y) = a' - \max\{\min(\bar{b} + \epsilon_b, y - \xi_1), \bar{c} + \epsilon_c\}.\tag{B.19}$$

We can see that $\pi_1^{WO,1}(y)$ is a piece-wise linear increasing function in y while $\pi_1^{WO,2}(y)$ is a piece-wise linear decreasing function in y . Furthermore, one can directly verify that $\pi_1^{WO,1}(y) = a$ for $y \geq \xi_1 + \bar{c} + \epsilon_c$ and $\pi_1^{WO,2}(y) = a' - \bar{c} - \epsilon_c$ for $y \leq \xi_1 + \bar{c} + \epsilon_c$.

Lemma 3.3

From (B.17) and (B.19), we have the following identities:

$$\begin{aligned}\pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b) &= \min(a, a - \xi_1 + \frac{1}{2} \min(\xi_1, \xi_1 - \bar{c} - \bar{b} + \epsilon_b - \epsilon_c) + \frac{1}{2}(\xi_1 - \bar{c} - \bar{b} + \epsilon_b - \epsilon_c)) \\ &= \min(a, a - \frac{1}{2}(\xi_1 + \bar{c} + \bar{b} - \epsilon_b + \epsilon_c)), \\ \pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b) &= a' - \max(\min(\bar{b} + \epsilon_b, \bar{b} - \epsilon_b), \bar{c} + \epsilon_c) \\ &= a' - \max(\bar{b} - \epsilon_b, \bar{c} + \epsilon_c), \\ \pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b) &= \min\left(a, a + \frac{1}{2} \min(\xi_1, \xi_1 - \xi + 2\epsilon_b) - \frac{1}{2}(\xi_1 + \xi - 2\epsilon_b)\right), \\ \pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b) &= a' - \max(\bar{b} + \epsilon_b, \bar{c} + \epsilon_c).\end{aligned}$$

There are three cases to consider:

1. $\bar{b} + \epsilon_b \leq \bar{c} + \epsilon_c$:

We have

$$\pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b) = a - (\bar{c} - \bar{b} + \epsilon_b + \epsilon_c),$$

$$\pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b) = a' - \bar{c} - \epsilon_c,$$

$$\pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b) = a - (\bar{c} - \bar{b} - \epsilon_b + \epsilon_c),$$

$$\pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b) = a' - \bar{c} - \epsilon_c.$$

2. $\bar{b} - \epsilon_b \leq \bar{c} + \epsilon_c < \bar{b} + \epsilon_b$:

We have

$$\pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b) = a - (\bar{c} - \bar{b} + \epsilon_b + \epsilon_c),$$

$$\pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b) = a' - \bar{c} - \epsilon_c,$$

$$\pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b) = a,$$

$$\pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b) = a' - \bar{b} - \epsilon_b.$$

3. $\bar{c} + \epsilon_c < \bar{b} - \epsilon_b$:

We have

$$\pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b) = a,$$

$$\pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b) = a' - \bar{b} + \epsilon_b,$$

$$\pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b) = a,$$

$$\pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b) = a' - \bar{b} - \epsilon_b.$$

In all three cases, one can verify that $\pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b) \leq \pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b)$ is equivalent to $a' \leq a + \bar{b} - \epsilon_b$ and $\pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b) \leq \pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b)$ is equivalent to $a' \leq a + \bar{b} + \epsilon_b$.

Proposition 3.12

Proof. We first consider the case when bidder 1 bids in the restricted policy space $\mathcal{U}'_1 = \{(x, y) \in \mathbb{R}_+^2 \mid x = \xi_1, y \geq x\}$. When $a' \in (a + \bar{b} - \epsilon_b, a + \bar{b} + \epsilon_b]$, by Lemma 3.3, we

have $\pi_1^{WO,1}(\xi_1 + \bar{b} - \epsilon_b) < \pi_1^{WO,2}(\xi_1 + \bar{b} - \epsilon_b)$ and $\pi_1^{WO,1}(\xi_1 + \bar{b} + \epsilon_b) \geq \pi_1^{WO,2}(\xi_1 + \bar{b} + \epsilon_b)$. Due to the monotonicity of $\pi_1^{WO,1}(y)$ and $\pi_1^{WO,2}(y)$ (Lemma 3.2), the optimal choice of y (with respect to the restricted space \mathcal{U}'_1) is at the intersection of $\pi_1^{WO,1}(y)$ and $\pi_1^{WO,2}(y)$. To determine the optimal worst-case payoff with respect to the restricted space \mathcal{U}'_1 , we recall that $\pi_1^{WO,1}(y)$ is constant for $y \geq \xi_1 + \bar{c} + \epsilon_c$ and $\pi_1^{WO,2}(y)$ is constant for $y \leq \xi_1 + \bar{c} + \epsilon_c$. Thus, at the intersection point of these two payoff functions, bidder 1's payoff is equal to $\min\{\pi_1^{WO,1}(\xi_1 + \bar{c} + \epsilon_c), \pi_1^{WO,2}(\xi_1 + \bar{c} + \epsilon_c)\} = \min(a, a' - \bar{c} - \epsilon_c)$ (see Figure 3.15c and 3.15d). Since this is a lower bound for the optimal worst-case payoff on the original policy space $\mathcal{U}_1 = \mathbb{R}_+^2$, one has that

$$\min(a, a' - \bar{c} - \epsilon_c) \leq \pi_1^{MAXMIN}. \quad (\text{B.20})$$

We now show that the above lower bound is the same with the upper bound on π_1^{MAXMIN} provided by the minimax inequality. For any $b_{-1} \in \mathcal{U}_{-1}$, bidder 1's best response payoff is her VCG payoff. Thus, $\sup_{b_1 \in \mathcal{U}_1} \pi_1(b_1, b_{-1}) = \max(a + b, a') - \max(b, c)$.

As a result, we have

$$\begin{aligned} \pi_1^{MINMAX} &= \inf_{b_{-1} \in \mathcal{U}_{-1}} \sup_{b_1 \in \mathcal{U}_1} \pi_1(b_1, b_{-1}) \\ &= \inf_{b_{-1} \in \mathcal{U}_{-1}} (\max(a + b, a') - \max(b, c)) \\ &= \min \left(\begin{array}{cc} \inf & a' - \max(b, c), & \inf & a - (c - b)^+ \\ \text{s.t.} & b < a' - a & \text{s.t.} & a' - a \leq b \\ & b_{-1} \in \mathcal{U}_{-1} & & b_{-1} \in \mathcal{U}_{-1} \end{array} \right) \\ &= \min\{a' - \max(a' - a, \bar{c} + \epsilon_c), a - (\bar{c} + \epsilon_c - a' + a)^+\} \\ &= \min(a, a' - \bar{c} - \epsilon_c), \end{aligned}$$

where the second last equality follows by substituting the minimizing values for b and c . By minimax inequality, we have $\pi_1^{MAXMIN} \leq \pi_1^{MINMAX} = \min(a, a' - \bar{c} - \epsilon_c)$.

Together with (B.20), we have $\pi_1^{MAXMIN} = \min(a, a' - \bar{c} - \epsilon_c)$ and the policy (x^*, y^*) is optimal. \square

Remark 3.6

Proof. Let $\pi_1^{WO,TR}$ be her worst-case payoff under truthful bidding. We want to show that if $a + \bar{b} - \epsilon_b < a'$ then this truthful worst-case payoff is the same with the optimal worst-case payoff from robust bidding, i.e.,

$$\pi_1^{WO,TR} = \pi_1^{MAXMIN}. \quad (\text{B.21})$$

First let us consider the case when $a + \bar{b} + \epsilon_b < a'$. By reporting truthfully, bidder 1 always wins two items, so her worst-case payoff is

$$\pi_1^{WO,TR} = a' - \max(\bar{b} + \epsilon_b, \bar{c} + \epsilon_c).$$

On the other hand, according to Proposition 3.11, we have $\pi_1^{MAXMIN} = a' - \max(\bar{b} + \epsilon_b, \bar{c} + \epsilon_c)$. Hence, $\pi_1^{WO,TR} = \pi_1^{MAXMIN}$.

Now let us assume that $a + \bar{b} - \epsilon_b < a' \leq a + \bar{b} + \epsilon_b$. By Proposition 3.12, $\pi_1^{MAXMIN} = \min(a, a' - \bar{c} - \epsilon_c)$. There are three possible cases to be considered:

1. $\bar{b} + \epsilon_b \leq \bar{c} + \epsilon_c$:

By reporting truthfully, bidder 1's worst-case payoffs when winning one and two items are respectively:

$$\begin{aligned} \pi_1^{WO1,TR} &= \min\left(a, \frac{1}{2} \min(a, a' - \bar{c} - \epsilon_c) + \frac{1}{2}(a' - \bar{c} - \epsilon_c)\right) \\ &= \min(a, a' - \bar{c} - \epsilon_c), \\ \pi_1^{WO2,TR} &= a' - \max(\min(a' - a, \bar{b} + \epsilon_b), \bar{c} + \epsilon_c) \\ &= a' - \bar{c} - \epsilon_c. \end{aligned}$$

Hence, bidder 1's worst-case payoff is

$$\pi_1^{WO,TR} = \min(\pi_1^{WO1,TR}, \pi_1^{WO2,TR}) = \min(a, a' - \bar{c} - \epsilon_c),$$

which is the same with π_1^{MAXMIN} .

2. $\bar{b} - \epsilon_b \leq \bar{c} + \epsilon_c < \bar{b} + \epsilon_b$:

Similar to the previous case, by reporting truthfully, bidder 1's worst-case payoffs when winning one and two items are respectively:

$$\pi_1^{WO1,TR} = \min(a, a' - \bar{c} - \epsilon_c),$$

$$\pi_1^{WO2,TR} = a' - \max(\min(a' - a, \bar{b} + \epsilon_b), \bar{c} + \epsilon_c).$$

Now we prove that $\min(\pi_1^{WO1,TR}, \pi_2^{WO2,TR}) = \min(a, a' - \bar{c} - \epsilon_c)$. If $a' - \bar{c} - \epsilon_c < a$ then $\pi_1^{WO1,TR} = \pi_1^{WO1,TR} = a' - \bar{c} - \epsilon_c$ so we indeed have

$$\min(\pi_1^{WO1,TR}, \pi_2^{WO2,TR}) = \min(a, a' - \bar{c} - \epsilon_c).$$

On the other hand, if $a \leq a' - \bar{c} - \epsilon_c$ then $\pi_1^{WO1,TR} = a$ and $\pi_1^{WO2,TR} = a' - \min(a' - a, \bar{b} + \epsilon_b) \geq a$, so $\min(\pi_1^{WO1,TR}, \pi_2^{WO2,TR}) = a = \min(a, a' - \bar{c} - \epsilon_c)$.

Therefore, $\pi_1^{WO,TR} = \pi_1^{MAXMIN}$.

3. $\bar{c} + \epsilon_c < \bar{b} - \epsilon_b$:

In this case, the worst-case payoffs of bidder 1 when winning one and two items are $\pi_1^{WO1,TR} = \pi_1^{WO2,TR} = a = \min(a, a' - \bar{c} - \epsilon_c)$. Thus, we also have $\pi_1^{WO,TR} = \pi_1^{MAXMIN}$, which completes our proof.

□

Proposition 3.13

Proof. (a) For each $b_{-1} \in \mathcal{U}_{-1}$, a best response of bidder 1 is to bid

$$b_1^{BEST} = b^{(m)} \mathbb{1}_{b^{(m)} \leq v_1}.$$

The maximum payoff of bidder 1 given b_{-1} is

$$\pi_1(b_1^{BEST}, b_{-1}) = (v_1 - b^{(m)}) \mathbb{1}_{b^{(m)} \leq v_1}. \quad (\text{B.22})$$

Minimizing the above payoff function over the uncertainty set \mathcal{U}_{-1} , one gets that $\pi_1^{MINMAX} = (v_1 - u(m))^+$, which is exactly the same with π_1^{MAXMIN} according to Proposition 3.1.

- (b) Without loss of generality, we assume that the common click-through rate is $\alpha = 1$. For each $b_{-1} \in \mathcal{U}_{-1}$, bidder 1's best response is to bid $b_1^{BEST} = b^{(m)}$ to win the m^{th} item if winning that item is profitable. Thus, the maximum payoff of bidder 1 given b_{-1} is

$$\pi_1(b_1^{BEST}, b_{-1}) = (v_1 - b^{(m)})^+.$$

By minimizing the above payoff over \mathcal{U}_{-1} , we get

$$\pi_1^{MINMAX} = (v_1 - u(m))^+,$$

which is exactly the same with π_1^{MAXMIN} , according to Proposition 3.2.

- (c) Minimax equality (3.13) follows directly from the proof of Proposition 3.10.

□

Appendix C

Proofs in Chapter 4

C.1 Proofs in Section 4.3

Proposition 4.1

Proof. Note that when $\tau_N \leq t \leq \tau_D \leq t + q$, the firm's total cost under normal and disrupted stages are respectively:

$$TC(t, q, \beta, T_0, \gamma_0, 0) = \frac{1}{2}h\lambda q + h\lambda(t - \tau_N) + \frac{K}{q},$$

$$TC(t, q, \beta, T_0, \gamma_0, 1) = \frac{1}{2q}h\lambda(t + q - \tau_D)^2 + \frac{1}{2q}b(1 + \beta)\lambda(\tau_D - t)^2 + \frac{K}{q}.$$

Thus, the firm's expected total cost before purchasing insurance is

$$\begin{aligned} \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)] &= (1 - \rho)\lambda \left[\frac{1}{2}hq + h(t - \tau_N) \right] \\ &\quad + \frac{\rho\lambda}{2q} \left[h(t + q - \tau_D)^2 + (1 + \beta)b(\tau_D - t)^2 \right] + \frac{K}{q}. \end{aligned}$$

The expected total cost $\mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]$ is convex in t and q because

$$\frac{\partial^2 \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial t^2} = \frac{\rho\lambda(h + b + b\beta)}{q} > 0,$$

$$\frac{\partial^2 \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial q^2} = \frac{2K + \rho\lambda(\tau_D - t)^2(h + b + b\beta)}{q^3} > 0,$$

and

$$\begin{aligned} & \frac{\partial^2 \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial t^2} \frac{\partial^2 \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial q^2} \\ & - \left(\frac{\partial^2 \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial t \partial q} \right)^2 \\ & = \frac{2K\rho\lambda(h + b + b\beta)}{q^4} > 0. \end{aligned}$$

Hence, $t^*(\beta)$ and $q^*(\beta)$ are uniquely determined by solving the first-order conditions:

$$\frac{\partial \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial t} = \frac{\lambda h[q + \rho(t - \tau_D)] + \rho\lambda b(1 + \beta)(t - \tau_D)}{q} = 0,$$

$$\frac{\partial \mathbb{E}_X[TC(t, q, \beta, T_0, \gamma_0, X)]}{\partial q} = \frac{\lambda h[q^2 - \rho(t - \tau_D)^2] - \rho\lambda b(1 + \beta)(t - \tau_D)^2 - 2K}{2q^2} = 0.$$

Solving these first-order conditions, we get $q^*(\beta)$ and $t^*(\beta)$ as given in (4.36) and (4.37), respectively. \square

Proposition 4.2

Proof. From the definition of $U(t, q, \beta, \gamma)$, we have:

$$\begin{aligned} U(t, q, \beta, \gamma) &= (1 - \rho)\lambda h \left[\frac{1}{2}q + t - \tau_N \right] + \frac{K}{q} \\ &+ \frac{\lambda\rho}{2q} \{ h(t + q - \tau_D)^2 + (1 + \beta\gamma)b(\tau_D - t)^2 \}. \end{aligned}$$

Therefore,

$$f(\beta) = U_\gamma(t^*(\beta), q^*(\beta), \beta, \gamma) = \rho \frac{\beta b \lambda}{2q^*(\beta)} (\tau_D - t^*(\beta))^2.$$

Similarly,

$$\begin{aligned}
f_0(\beta) &= U(t^*(\beta), q^*(\beta), \beta, 0) \\
&= (1 - \rho)h\lambda \left[\frac{1}{2}q^*(\beta) + t^*(\beta) - \tau_N \right] + \frac{K}{q^*(\beta)} \\
&\quad + \frac{\rho\lambda}{2q^*(\beta)} \left[h(t^*(\beta) + q^*(\beta) - \tau_D)^2 + b(\tau_D - t^*(\beta))^2 \right].
\end{aligned}$$

□

Lemma 4.2

Proof. From Proposition 4.2, we have

$$U_{\beta\gamma}(\beta, \gamma) = f'(\beta).$$

Recall that

$$\begin{aligned}
f(\beta) &= \rho \frac{\beta b \lambda}{2q^*(\beta)} (\tau_D - t^*(\beta))^2 \\
&= \frac{(\lambda K)^{1/2} b h^{3/2} \beta}{(2\rho)^{1/2} (h + b + b\beta)^{3/2} [\rho(h + b + b\beta) - h]^{1/2}}.
\end{aligned}$$

Hence, taking the first derivative with respect to β , we have:

$$f'(\beta) = d(\beta) \left[bh\beta - 2hb - 2h^2 + (4bh + 2h^2 + 2b^2 - 2b^2\beta^2)\rho \right], \quad (\text{C.1})$$

where $d(\beta)$ is defined as:

$$d(\beta) = \frac{bh^{3/2}(K\lambda)^{1/2}}{2^{3/2}\rho^{1/2}(h + b + b\beta)^{5/2}[\rho(h + b + b\beta) - h]^{3/2}} > 0, \quad \forall \beta > 0.$$

One can verify that for $\beta < \bar{\beta}_H$ (Assumption 4.2), the second term in the square bracket on the right hand side of (C.1) is positive. Hence, $U_{\beta\gamma}(\beta, \gamma) > 0$ for $\beta \in [\beta_L, \beta_H]$ and $\gamma \in [0, 1]$. □

Lemma 4.3

Proof. (i) Assume the opposite that $\gamma_L < \gamma_H$. Adding (IC-L) and (IC-H), we have

$$U(\beta_L, \gamma_L) + U(\beta_H, \gamma_H) \leq U(\beta_L, \gamma_H) + U(\beta_H, \gamma_L),$$

or equivalently

$$U(\beta_H, \gamma_H) - U(\beta_L, \gamma_H) \leq U(\beta_H, \gamma_L) - U(\beta_L, \gamma_L).$$

However, the single-crossing property (Lemma 4.2) implies that

$$U(\beta_H, \gamma_H) - U(\beta_L, \gamma_H) > U(\beta_H, \gamma_L) - U(\beta_L, \gamma_L),$$

which is a contradiction. Thus, we must have $\gamma_H \leq \gamma_L$.

(ii) Let $\Gamma^* = \{(T_L^*, \gamma_L^*), (T_H^*, \gamma_H^*)\}$ be an optimal solution of (4.19). Assume that (IR-L) does not bind. Let \bar{T}_L and \bar{T}_H be defined as follows:

$$\begin{aligned}\bar{T}_L &= U(\beta_L, \gamma_0) - U(\beta_L, \bar{\gamma}), \\ \bar{T}_H &= U(\beta_H, \gamma_0) - U(\beta_H, \bar{\gamma}).\end{aligned}$$

We have by (IR-L) that

$$T_L^* < U(\beta_L, \gamma_0) - U(\beta_L, \gamma_L^*) = U(\beta_L, \bar{\gamma}) - U(\beta_L, \gamma_L^*) + \bar{T}_L \leq U(\beta_H, \bar{\gamma}) - U(\beta_H, \gamma_L) + \bar{T}_L,$$

where the last inequality comes from the single-crossing property (Lemma 4.2). Consequently, we have

$$U(\beta_H, \gamma_L^*) + T_L^* < U(\beta_H, \bar{\gamma}) + \bar{T}_L \leq U(\beta_H, \bar{\gamma}) + \bar{T}_H = U(\beta_H, \gamma_0),$$

where the second inequality follows from the fact that $\bar{T}_L \leq \bar{T}_H$ (also from the single-crossing property). By (IC-H), we have

$$U(\beta_H, \gamma_H^*) + T_H^* \leq U(\beta_H, \gamma_L^*) + T_L^*.$$

Hence,

$$U(\beta_H, \gamma_H^*) + T_H^* < U(\beta_H, \gamma_0),$$

so (IR-H) also does not bind. This implies that there exists $\epsilon > 0$ such that the contract menu $\Gamma_\epsilon^* = \{(T_L^* + \epsilon, \gamma_L^*), (T_H^* + \epsilon, \gamma_H^*)\}$ is feasible and more profitable than the contract menu Γ^* , which is a contradiction.

(iii) Assume the opposite that (IC-H) does not bind at the optimal solution Γ^* . If (IR-H) does not bind also then there exists $\epsilon > 0$ such that the menu of contracts $\Gamma_\epsilon^* = \{(T_L^*, \gamma_L^*), (T_H^* + \epsilon, \gamma_H^*)\}$ is feasible and more profitable, which is a contradiction. Thus, (IR-H) must bind. But then by (IC-H) we have

$$U(\beta_H, \gamma_0) = U(\beta_H, \gamma_H^*) + T_H^* < U(\beta_H, \gamma_L^*) + T_L^*.$$

Therefore, we have

$$U(\beta_H, \gamma_0) - U(\beta_H, \gamma_L^*) < T_L^* = U(\beta_L, \gamma_0) - U(\beta_L, \gamma_L^*),$$

where the equality follows from (IR-L). This above inequality contradicts the single-crossing property (Lemma 4.2). Thus, (IC-H) binds.

(iv) The result follows directly from arguments in (ii) and (iii). \square

Proposition 4.3

Proof. The premium levels (4.43) and (4.44) are optimal due to Lemma 4.3. \square

Proposition 4.4

Proof. We first show that if a menu of contracts $\Gamma^* = \{(T_L^*, \gamma_L^*), (T_H^*, \gamma_H^*)\}$ has premium levels given by (4.43) and (4.44) and the contract qualities satisfy the monotonicity condition $\underline{\gamma} \leq \gamma_H^* \leq \gamma_L^* \leq \bar{\gamma}$ then it is feasible to the problem (4.19). We only need to check for (IC-L) and (IR-H), since (4.43) and (4.44) already ensure that (IR-L) and (IC-H) hold with equality. By replacing T_L^* and T_H^* with the respective expressions given in (4.43) and (4.44), we have that the constraint (IC-L) is

equivalent to

$$U(\beta_H, \gamma_H^*) - U(\beta_L, \gamma_H^*) \leq U(\beta_H, \gamma_L^*) - U(\beta_L, \gamma_L^*),$$

which holds due to the single-crossing property and the fact that $\gamma_H^* \leq \gamma_L^*$. Similarly, (IR-H) is equivalent to

$$U(\beta_H, \gamma_L^*) - U(\beta_L, \gamma_L^*) \leq U(\beta_H, \gamma_0) - U(\beta_L, \gamma_0) + U(\beta_H, \gamma_H^*) - U(\beta_L, \gamma_H^*).$$

We have

$$U(\beta_H, \gamma_L^*) - U(\beta_L, \gamma_L^*) \leq U(\beta_H, \gamma_0) - U(\beta_L, \gamma_0)$$

due to the single-crossing property and the fact that $\gamma_L \leq \gamma_0 = 1$. In addition, $U(\beta, \gamma)$ is increasing in β , so $0 \leq U(\beta_H, \gamma_H^*) - U(\beta_L, \gamma_H^*)$. Hence, (IR-H) holds. Combining the above result with those in Lemma 4.3, we have that the menu of contracts $\Gamma^* = \{(T_L^*, \gamma_L^*), (T_H^*, \gamma_H^*)\}$ is an optimal solution to the problem $\mathcal{IP}(\rho, \phi)$ if and only if the premium levels are given by (4.43) and (4.44) and the quality levels γ_L^* and γ_H^* are optimal to the following problem:

$$\begin{aligned} \max_{\gamma_L, \gamma_H} \quad & \left[-f(\beta_L) + (1 - \phi) \frac{f(\beta_L)}{\beta_L} + \phi f(\beta_H) \right] \gamma_L + \phi(1 - \beta_H) \frac{f(\beta_H)}{\beta_H} \gamma_H \quad (\text{C.2}) \\ \text{s.t.} \quad & \underline{\gamma} \leq \gamma_H \leq \gamma_L \leq \bar{\gamma}. \end{aligned}$$

The objective function of the above optimization problem is obtained by substituting the expression for T_L^* and T_H^* as given in (4.43) and (4.44) to the objective function (4.19). We now consider the following cases:

(i) If $\beta_H < \min(1, \bar{\beta}_H)$ then since $\phi(1 - \beta_H) \frac{f(\beta_H)}{\beta_H} > 0$ it is optimal to select $\gamma_H = \gamma_L$. In this case, the insurer's problem becomes

$$\begin{aligned} \max_{\gamma_L} \quad & \left[\left(\frac{f(\beta_L)}{\beta_L} - \frac{f(\beta_H)}{\beta_H} \right) (1 - \phi) - f(\beta_L) + \frac{f(\beta_H)}{\beta_H} \right] \gamma_L \\ \text{s.t.} \quad & \underline{\gamma} \leq \gamma_L \leq \bar{\gamma} \end{aligned}$$

Note that

$$\frac{f(\beta)}{\beta} = \frac{(\lambda K)^{1/2} b h^{3/2}}{(2\rho)^{1/2} (h + b + b\beta)^{3/2} [\rho(h + b + b\beta) - h]^{1/2}}$$

is decreasing in β , so $\frac{f(\beta_L)}{\beta_L} - \frac{f(\beta_H)}{\beta_H} > 0$. Furthermore, $f(\beta)$ is increasing in β , so

$$-f(\beta_L) + \frac{f(\beta_H)}{\beta_H} \geq -f(\beta_L) + f(\beta_H) > 0.$$

Therefore,

$$\left(\frac{f(\beta_L)}{\beta_L} - \frac{f(\beta_H)}{\beta_H} \right) (1 - \phi) - f(\beta_L) + \frac{f(\beta_H)}{\beta_H} > 0, \quad \forall \phi \in [0, 1],$$

so $\gamma_L^* = \gamma_H^* = \bar{\gamma}$

(ii) & (iii) Assume that $1 \leq \beta_H < \bar{\beta}_H$, then $\gamma_H^* = \underline{\gamma}$. The low type contract quality γ_L is determined by solving

$$\begin{aligned} \max_{\gamma_L} & \quad \left[\left(\frac{f(\beta_L)}{\beta_L} - f(\beta_H) \right) (1 - \phi) + f(\beta_H) - f(\beta_L) \right] \gamma_L \\ \text{s.t.} & \quad \underline{\gamma} \leq \gamma_L \leq \bar{\gamma}. \end{aligned}$$

Observe that if

$$\frac{f(\beta_L)}{\beta_L} - f(\beta_H) < 0,$$

and

$$\phi < \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}$$

then

$$\left(\frac{f(\beta_L)}{\beta_L} - f(\beta_H) \right) (1 - \phi) + f(\beta_H) - f(\beta_L) < 0$$

and $\gamma_L^* = \underline{\gamma}$. On the other hand, if either

$$\frac{f(\beta_L)}{\beta_L} - f(\beta_H) \geq 0$$

or

$$\frac{f(\beta_L)}{\beta_L} - f(\beta_H) < 0, \quad \text{and } \phi \geq \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}$$

then

$$\left(\frac{f(\beta_L)}{\beta_L} - f(\beta_H) \right) (1 - \phi) + f(\beta_H) - f(\beta_L) \geq 0$$

so $\gamma_L^* = \bar{\gamma}$. □

Proposition 4.5

Proof. Assume that $\rho = \rho_0$. By applying (4.21) repeatedly, we get

$$\pi_{i,1} = \frac{\pi_0 A_i(X_{1,1}, X_{2,1}, \dots, X_{i,1})}{\pi_0 A_i(X_{1,1}, X_{2,1}, \dots, X_{i,1}) + (1 - \pi_0) B_i(X_{1,1}, X_{2,1}, \dots, X_{i,1})}, \quad (\text{C.3})$$

where $A_i(X_{1,1}, X_{2,1}, \dots, X_{i,1})$ and $B_i(X_{1,1}, X_{2,1}, \dots, X_{i,1})$ are defined respectively by

$$A_i(X_{1,1}, X_{2,1}, \dots, X_{i,1}) = \prod_{j=1}^i \eta^{X_{j,1}},$$

$$B_i(X_{1,1}, X_{2,1}, \dots, X_{i,1}) = \prod_{j=1}^i \theta^{1-X_{j,1}},$$

with $\eta = \frac{\rho_1}{\rho_0}$ and $\theta = \frac{1-\rho_0}{1-\rho_1}$. Thus,

$$\begin{aligned} \mathbb{E}_0^\Gamma[\pi_{i,1}] &= \mathbb{E}_0^\Gamma \left[\frac{\pi_0 A_i(X_{1,1}, X_{2,1}, \dots, X_{i,1})}{\pi_0 A_i(X_{1,1}, X_{2,1}, \dots, X_{i,1}) + (1 - \pi_0) B_i(X_{1,1}, X_{2,1}, \dots, X_{i,1})} \right] \\ &= \mathbb{E}_0^\Gamma \left[\frac{1}{1 + \frac{1-\pi_0}{\pi_0} \theta^{\sum_{j=1}^i (1-X_{j,1})} \eta^{-\sum_{j=1}^i X_{j,1}}} \right] \\ &= \mathbb{E}_0^\Gamma \left[\frac{1}{1 + \frac{1-\pi_0}{\pi_0} \exp \left(\log \theta \sum_{j=1}^i (1 - X_{j,1}) - \log \eta \sum_{j=1}^i X_{j,1} \right)} \right]. \end{aligned}$$

Let L_i be the exponent in the denominator. For convenience, we write $\bar{\rho}_0 = 1 - \rho_0$ and $\bar{\rho}_1 = 1 - \rho_1$. We have

$$\begin{aligned}
L_i &= \sum_{j=1}^i (1 - X_{j,1}) \log \theta - \sum_{j=1}^i X_{j,1} \log \eta \\
&= \sum_{j=1}^i (1 - X_{j,1} - \bar{\rho}_0) \log \theta - \sum_{j=1}^i (X_{j,1} - \rho_0) \log \eta + i [\bar{\rho}_0 \log \theta - \rho_0 \log \eta] \\
&= \sum_{j=1}^i (\rho_0 - X_{j,1}) \log \theta - \sum_{j=1}^i (X_{j,1} - \rho_0) \log \eta + i [\bar{\rho}_0 \log \theta - \rho_0 \log \eta] \\
&= \sum_{j=1}^i (X_{j,1} - \rho_0) \log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} + i [\bar{\rho}_0 \log \theta - \rho_0 \log \eta]. \tag{C.4}
\end{aligned}$$

Let $\delta = |\rho_0 - \rho_1|$. Observe that,

$$\bar{\rho}_0 \log \theta - \rho_0 \log \eta = \bar{\rho}_0 \log \frac{1 - \rho_0}{1 - \rho_1} - \rho_0 \log \frac{\rho_1}{\rho_0} \geq 2(\rho_0 - \rho_1)^2 = 2\delta^2.$$

We now turn our attention to the first term in (C.4). Define $M_0 = 0$ and

$$M_i = \sum_{j=1}^i (X_{j,1} - \rho_0) \log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1}.$$

Let $\mathcal{F}_i = \sigma(X_{1,1}, X_{2,1}, \dots, X_{i,1})$ be the filtration induced by observations up to $X_{i,1}$ in decision epoch 1. Observe that $\{M_i\}$ is adapted to \mathcal{F}_i and since $X_{j,1} \in \{0, 1\}$:

$$\mathbb{E}_0^\Gamma |M_i| \leq \max \left(\left| i(1 - \rho_0) \log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} \right|, \left| -i\rho_0 \log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} \right| \right) < \infty.$$

Furthermore,

$$\begin{aligned}
\mathbb{E}_0^\Gamma [M_i | \mathcal{F}_{i-1}] &= \mathbb{E}_0^\Gamma \left[\log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} \sum_{j=1}^i (X_{j,1} - \rho_0) \mid \mathcal{F}_{i-1} \right] \\
&= \log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} \sum_{j=1}^{i-1} (X_{j,1} - \rho_0) + \mathbb{E}_0^\Gamma \left[\log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} (X_{i,1} - \rho_0) \mid \mathcal{F}_{i-1} \right] \\
&= M_{i-1},
\end{aligned}$$

where the last equation follows from $\mathbb{E}_0^\Gamma(X_{i,1} | \mathcal{F}_{i-1}) = \rho_0$. Hence, $\{M_i\}$ is a martingale with respect to \mathcal{F}_i . In addition, $\{M_i\}$ has bounded differences:

$$|M_i - M_{i-1}| = \left| \log \frac{\bar{\rho}_1 \rho_0}{\bar{\rho}_0 \rho_1} (X_{i,1} - \rho_0) \right| \leq \log \frac{\rho_{\max}(1 - \rho_{\min})}{\rho_{\min}(1 - \rho_{\max})},$$

where $\rho_{\max} = \max(\rho_0, \rho_1)$ and $\rho_{\min} = \min(\rho_0, \rho_1)$. Let $\epsilon = (\log \frac{\rho_{\max}(1 - \rho_{\min})}{\rho_{\min}(1 - \rho_{\max})})^2$ then by the Azuma-Hoeffding inequality, for some $\xi > 0$,

$$\mathbb{P}_0^\Gamma(|M_i| \geq i\xi) \leq 2 \exp\left(-\frac{1}{2\epsilon} \xi^2 i\right). \quad (\text{C.5})$$

Let $\mathcal{E}_{i,\xi} = \{|M_i| < \xi i\}$. We have

$$\begin{aligned} \mathbb{E}_0^\Gamma[\pi_{i,1}] &\leq \mathbb{E}_0^\Gamma \left[\frac{1}{1 + \frac{1-\pi_0}{\pi_0} \exp(M_i + 2i\delta^2)} \right] \\ &\leq \mathbb{E}_0^\Gamma \left[\frac{1}{1 + \frac{1-\pi_0}{\pi_0} \exp(M_i + 2i\delta^2)}; \mathcal{E}_{i,\xi} \right] \\ &\quad + \mathbb{E}_0^\Gamma \left[\frac{1}{1 + \frac{1-\pi_0}{\pi_0} \exp(M_i + 2i\delta^2)}; \mathcal{E}_{i,\xi}^c \right] \\ &\leq \frac{1}{1 + \frac{1-\pi_0}{\pi_0} \exp(-\xi i + 2i\delta^2)} + \mathbb{P}_0^\Gamma(|M_i| \geq i\xi) \\ &\leq \frac{1}{1 + \frac{1-\pi_0}{\pi_0} \exp(-\xi i + 2i\delta^2)} + 2 \exp\left(-\frac{1}{2\epsilon} \xi^2 i\right), \end{aligned}$$

where the second last step is from the definition of $\mathcal{E}_{i,\xi}$ and the last step follows from (C.5). Let $\mu_0 = 2 \max\{2, \pi_0/(1 - \pi_0)\}$ and $\mu_1 = \min\{2\delta^2 - \xi, \xi^2/2\epsilon\}$. We have

$$\begin{aligned} \mathbb{E}_0^\Gamma[\pi_{i,1}] &\leq \frac{\pi_0}{1 - \pi_0} \exp(\xi i - 2i\delta^2) + 2 \exp\left(-\frac{1}{2\epsilon} \xi^2 i\right) \\ &\leq \mu_0 \exp(-\mu_1 i). \end{aligned}$$

The proof for the case of $\mathbb{E}_1^\Gamma[1 - \pi_{i,1}] \leq \mu_0 \exp(\mu_1 i)$ can be obtained in a similar fashion. \square

Proposition 4.6

Proof. Assume that $\rho = \rho_0$. From Proposition 4.5, if I_1 is significantly large then under the MBP Ψ , $\pi_{i,n} = 0$ almost surely for $n \geq 2$ and $i = 1, 2, \dots, I_n$. As a result, $\Psi_n \equiv \Gamma_0^*$ for $n \geq 2$, i.e., under the MBP, the menu of contracts in decision epoch $n \geq 2$ is the same as the clairvoyant's optimal menu of contracts. Hence, for $N \geq 1$

$$\begin{aligned} \Delta_0^\Psi(N) &= \frac{1}{v_0(\Gamma_0^*)} [Nv_0(\Gamma_0^*) - V_0^\Psi(N)] = \frac{1}{v_0(\Gamma_0^*)} [v_0(\Gamma_0^*) - v_0(\Psi_1)] \\ &\leq \frac{1}{v_0(\Gamma_0^*)} [v_0(\Gamma_0^*) - \underline{v}_0], \end{aligned}$$

where \underline{v}_0 is a lower bound on $v_0(\Psi)$ for all Ψ that is feasible to $\mathcal{IP}(\rho_0, \phi)$. Let C_0 be the constant on the right hand side of the above inequality. We have $\Delta_0^\Psi(N) \leq C_0$ for $N \geq 1$. Similarly, when $\rho = \rho_1$, there exists a constant C_1 such that $\Delta_1^\Psi(N) \leq C_1$ for $N \geq 1$. Let $C = (C_0 + C_1)/2$. We have:

$$\Delta^\Psi(N) = \frac{1}{2} [\Delta_0^\Psi(N) + \Delta_1^\Psi(N)] \leq C, \quad N \geq 1.$$

□

Proposition 4.7

Proof. Assume without loss of generality that $\phi = \phi_0$. Let $\tilde{\mathcal{F}}_n = \sigma(Y_1, Y_2, \dots, Y_n)$ and note that $\tilde{\pi}_n$ is bounded so it is $\tilde{\mathcal{F}}_n$ -measurable. If $Y_n = -1$ then the offered menu of contracts is pooling so the insurer does not observe the customer's type from her contract choice. As a result,

$$\tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n \mid \tilde{\mathcal{F}}_{n-1}, Y_n = -1] = \tilde{\pi}_{n-1}. \quad (\text{C.6})$$

On the other hand, if $Y_n \geq 0$ then the insurer can observe the customer's type from her contract choice and update her belief using the Bayes' rule. Therefore, we have

$$\begin{aligned}
& \tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n \mid \tilde{\mathcal{F}}_{n-1}, Y_n \geq 0] \\
&= \tilde{\mathbb{E}}_0^\Gamma \left[\frac{\tilde{\pi}_{n-1} \phi_1^{Y_n} (1 - \phi_1)^{1-Y_n}}{\tilde{\pi}_{n-1} \phi_1^{Y_n} (1 - \phi_1)^{1-Y_n} + (1 - \tilde{\pi}_{n-1}) \phi_0^{Y_n} (1 - \phi_0)^{1-Y_n}} \mid \tilde{\mathcal{F}}_{n-1}, Y_n \geq 0 \right] \\
&= (1 - \phi_0) \frac{\tilde{\pi}_{n-1} (1 - \phi_1)}{\tilde{\pi}_{n-1} (1 - \phi_1) + (1 - \tilde{\pi}_{n-1}) (1 - \phi_0)} + \phi_0 \frac{\tilde{\pi}_{n-1} \phi_1}{\tilde{\pi}_{n-1} \phi_1 + (1 - \tilde{\pi}_{n-1}) \phi_0} \\
&= \tilde{\pi}_{n-1} - \frac{\tilde{\pi}_{n-1}^2 (1 - \tilde{\pi}_{n-1}) (\phi_0 - \phi_1)^2}{[\tilde{\pi}_{n-1} (1 - \phi_1) + (1 - \tilde{\pi}_{n-1}) (1 - \phi_0)] [\tilde{\pi}_{n-1} \phi_1 + (1 - \tilde{\pi}_{n-1}) \phi_0]}. \quad (\text{C.7})
\end{aligned}$$

Combining (C.6) and (C.7), we get:

$$\begin{aligned}
& \tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n \mid \tilde{\mathcal{F}}_{n-1}] \\
&= \tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n \mid \tilde{\mathcal{F}}_{n-1}, Y_n = -1] \tilde{\mathbb{P}}_0^\Gamma(Y_n = -1) + \tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n \mid \tilde{\mathcal{F}}_{n-1}, Y_n \geq 0] \tilde{\mathbb{P}}_0^\Gamma(Y_n \geq 0) \\
&\leq \tilde{\pi}_{n-1} \left[\tilde{\mathbb{P}}_0^\Gamma(Y_n = -1) + \tilde{\mathbb{P}}_0^\Gamma(Y_n \geq 0) \right] \\
&= \tilde{\pi}_{n-1}.
\end{aligned}$$

Thus, $\{\tilde{\pi}_n\}$ is a sub-martingale with respect to $\tilde{\mathcal{F}}_n$. A standard result in martingale theory implies that $\{\tilde{\pi}_n\}$ converges almost surely to a finite random variable $\tilde{\pi}_\infty$ (Williams, 1991). \square

Proposition 4.8

Proof. Assume that $\phi = \phi_0$. Since Γ is a self-selecting policy, $Y_n \geq 0$ a.s. for all n . Furthermore, $Y_n = 1$ if the customer is of type β_H and $Y_n = 0$ if the customer is of type β_L . By applying (4.24) repeatedly, we get

$$\tilde{\pi}_n = \frac{\tilde{\pi}_0 C_n(Y_1, Y_2, \dots, Y_n)}{\tilde{\pi}_0 C_n(Y_1, Y_2, \dots, Y_n) + (1 - \tilde{\pi}_0) D_n(Y_1, Y_2, \dots, Y_n)}, \quad (\text{C.8})$$

where

$$C_n(Y_1, Y_2, \dots, Y_n) = \tilde{\eta}^{\sum_{m=1}^n Y_m},$$

$$D_n(Y_1, Y_2, \dots, Y_n) = \tilde{\theta}^{\sum_{m=1}^n (1-Y_m)},$$

and $\tilde{\eta} = \frac{\phi_1}{\phi_0}$, $\tilde{\theta} = \frac{1-\phi_0}{1-\phi_1}$. Thus,

$$\begin{aligned} \tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n] &= \tilde{\mathbb{E}}_0^\Gamma \left[\frac{\tilde{\pi}_0 C_n(Y_1, Y_2, \dots, Y_n)}{\tilde{\pi}_0 C_n(Y_1, Y_2, \dots, Y_n) + (1 - \tilde{\pi}_0) D_n(Y_1, Y_2, \dots, Y_n)} \right] \\ &= \tilde{\mathbb{E}}_0^\Gamma \left[\frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \tilde{\theta}^{\sum_{m=1}^n (1-Y_m)} \tilde{\eta}^{-\sum_{m=1}^n Y_m}} \right] \\ &= \tilde{\mathbb{E}}_0^\Gamma \left[\frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp \left(\sum_{m=1}^n (1-Y_m) \log \tilde{\theta} - \sum_{m=1}^n Y_m \log \tilde{\eta} \right)} \right]. \end{aligned}$$

Let \tilde{L}_n be the exponent in the denominator. For convenience, we write $\bar{\phi}_0 = 1 - \phi_0$ and $\bar{\phi}_1 = 1 - \phi_1$. We have

$$\begin{aligned} \tilde{L}_n &= \sum_{m=1}^n (1 - Y_m) \log \tilde{\theta} - \sum_{m=1}^n Y_m \log \tilde{\eta} \\ &= \sum_{m=1}^n (1 - Y_m - \bar{\phi}_0) \log \tilde{\theta} - \sum_{m=1}^n (Y_m - \phi_0) \log \tilde{\eta} + n \left[\bar{\phi}_0 \log \tilde{\theta} - \phi_0 \log \tilde{\eta} \right] \\ &= \sum_{m=1}^n (\phi_0 - Y_m) \log \tilde{\theta} - \sum_{m=1}^n (Y_m - \phi_0) \log \tilde{\eta} + n \left[\bar{\phi}_0 \log \tilde{\theta} - \phi_0 \log \tilde{\eta} \right] \\ &= \sum_{m=1}^n (Y_m - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1} + n \left[\bar{\phi}_0 \log \tilde{\theta} - \phi_0 \log \tilde{\eta} \right]. \end{aligned} \tag{C.9}$$

Let $\tilde{\delta} = |\phi_0 - \phi_1| > 0$. Observe that,

$$\bar{\phi}_0 \log \tilde{\theta} - \phi_0 \log \tilde{\eta} = \bar{\phi}_0 \log \frac{1 - \phi_0}{1 - \phi_1} - \phi_0 \log \frac{\phi_1}{\phi_0} \geq 2(\phi_0 - \phi_1)^2 = 2\tilde{\delta}^2.$$

We now turn our attention to the first term in (C.9). Define $\tilde{M}_0 = 0$ and

$$\tilde{M}_n = \sum_{m=1}^n (Y_m - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1}.$$

Let $\tilde{\mathcal{F}}_n$ be the filtration induced by (Y_1, Y_2, \dots, Y_n) . $\{\tilde{M}_n\}$ is adapted to $\tilde{\mathcal{F}}_n$ and

$$\tilde{\mathbb{E}}_0^\Gamma |\tilde{M}_n| \leq \max \left(\left| n(1 - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1} \right|, \left| -n\phi_0 \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1} \right| \right) < \infty.$$

Furthermore,

$$\begin{aligned} \tilde{\mathbb{E}}_0^\Gamma [\tilde{M}_n | \tilde{\mathcal{F}}_{n-1}] &= \tilde{\mathbb{E}}_0^\Gamma \left[\sum_{m=1}^n (Y_m - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1} \mid \tilde{\mathcal{F}}_{n-1} \right] \\ &= \sum_{m=1}^{n-1} (Y_m - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1} + \tilde{\mathbb{E}}_0^\Gamma \left[(Y_n - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1} \mid \tilde{\mathcal{F}}_{n-1} \right] \\ &= \tilde{M}_{n-1}, \end{aligned}$$

where the last equality follows from $\tilde{\mathbb{E}}_0^\Gamma [Y_n | \tilde{\mathcal{F}}_{n-1}] = \phi_0$. Thus, $\{\tilde{M}_n\}$ is a martingale with respect to $\tilde{\mathcal{F}}_n$. Also,

$$|\tilde{M}_n - \tilde{M}_{n-1}| = |(Y_n - \phi_0) \log \frac{\bar{\phi}_1 \phi_0}{\bar{\phi}_0 \phi_1}| \leq \log \frac{\phi_{\max}(1 - \phi_{\min})}{\phi_{\min}(1 - \phi_{\max})},$$

where $\phi_{\max} = \max(\phi_0, \phi_1)$ and $\phi_{\min} = \min(\phi_0, \phi_1)$. Let $\epsilon = \left(\log \frac{\phi_{\max}(1 - \phi_{\min})}{\phi_{\min}(1 - \phi_{\max})} \right)^2$ then by the Azuma-Hoeffding inequality, for some $\xi > 0$,

$$\tilde{\mathbb{P}}_0^\Gamma (|\tilde{M}_n| \geq n\xi) \leq 2 \exp \left(-\frac{1}{2\epsilon} \xi^2 n \right). \quad (\text{C.10})$$

Let $\tilde{\mathcal{E}}_{n,\xi} = \{|\tilde{M}_n| < \xi n\}$. We have

$$\begin{aligned}
\tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n] &= \tilde{\mathbb{E}}_0^\Gamma \left[\frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp(\tilde{M}_n + 2n\tilde{\delta}^2)} \right] \\
&\leq \tilde{\mathbb{E}}_0^\Gamma \left[\frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp(\tilde{M}_n + 2n\tilde{\delta}^2)}; \tilde{\mathcal{E}}_{n,\xi} \right] \\
&\quad + \tilde{\mathbb{E}}_0^\Gamma \left[\frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp(\tilde{M}_n + 2n\tilde{\delta}^2)}; \tilde{\mathcal{E}}_{n,\xi}^c \right] \\
&\leq \frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp(-\xi n + 2n\tilde{\delta}^2)} + \tilde{\mathbb{P}}_0^\Gamma(|\tilde{M}_n| \geq n\xi) \\
&\leq \frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp(-\xi n + 2n\tilde{\delta}^2)} + 2 \exp\left(-\frac{1}{2\epsilon} \xi^2 n\right),
\end{aligned}$$

where the second last step is from the definition of $\tilde{\mathcal{E}}_{n,\xi}$ and the last step follows from (C.10). Let $\mu_0 = 2 \max\{2, \tilde{\pi}_0/(1 - \tilde{\pi}_0)\}$ and $\mu_1 = \min\{2\tilde{\delta}^2 - \xi, \xi^2/2\epsilon\}$. We have

$$\begin{aligned}
\tilde{\mathbb{E}}_0^\Gamma[\tilde{\pi}_n] &\leq \frac{1}{1 + \frac{1-\tilde{\pi}_0}{\tilde{\pi}_0} \exp(-\xi n + 2n\tilde{\delta}^2)} + 2 \exp\left(-\frac{1}{2\epsilon} \xi^2 n\right) \\
&\leq \frac{\tilde{\pi}_0}{1 - \tilde{\pi}_0} \exp(\xi n - 2n\tilde{\delta}^2) + 2 \exp\left(-\frac{1}{2\epsilon} \xi^2 n\right) \leq \mu_0 \exp(-\mu_1 n).
\end{aligned}$$

The proof for the case of $\tilde{\mathbb{E}}_1^\Gamma[1 - \tilde{\pi}_n]$ is similar. \square

Proposition 4.9

Proof. If we can show that $\Psi_{\epsilon,\alpha}$ is a self-selecting policy, then the convergence results follow directly from Proposition 4.8. By definition, $\Psi_{\epsilon,\alpha}$ is separating since the selected menu of contracts in each decision epoch satisfies the separating condition (4.50). Furthermore, each menu of contracts also satisfies (IC-L) and (IC-H). As a result, $\Psi_{\epsilon,\alpha}$ is a self-selecting policy. \square

Proposition 4.10

Proof. We only need to show the constant upper bound for $\tilde{\Delta}_0^{\Psi^{\epsilon,\alpha}}(N)$, since the bound for $\tilde{\Delta}_1^{\Psi^{\epsilon,\alpha}}(N)$ can be obtained in a similar fashion. From Proposition 4.4, if $\beta_H < \min(1, \bar{\beta}_H)$, or $1 \leq \beta_H < \bar{\beta}_H$ and $f(\beta_L) \geq \beta_L f(\beta_H)$, then the single-period decision is the same for both the clairvoyant and the insurer under $\Psi_{\epsilon,\alpha}$. Thus, $\tilde{\Delta}_0^{\Psi^{\epsilon,\alpha}}(N) = 0$.

For the rest of this proof, we assume that $1 \leq \beta_H < \bar{\beta}_H$ and $f(\beta_L) \geq \beta_L f(\beta_H)$. If ϕ_0 is such that

$$\phi_0 > \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}, \quad (\text{C.11})$$

then by Proposition 4.4, the clairvoyant selects a separating menu of contracts in each decision epoch with quality $\bar{\gamma}$ and $\underline{\gamma}$ for low type and high type customers, respectively. Let $\epsilon > 0$ be such that $\epsilon < |\bar{\gamma} - \underline{\gamma}|$. Also, let $\delta > 0$ be such that

$$\delta < |\phi_1 - \phi_0|^{-1} \left[\phi_0 - \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)} \right]. \quad (\text{C.12})$$

By the definition of $\tilde{\Delta}_0^{\Psi^{\epsilon,\alpha}}(N)$, we have

$$\begin{aligned} \tilde{\Delta}_0^{\Psi^{\epsilon,\alpha}}(N) &= \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \left[N\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{V}_0^{\Psi^{\epsilon,\alpha}}(N) \right] \\ &= \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \tilde{\mathbb{E}}_0^{\Psi^{\epsilon,\alpha}} [\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon,\alpha}^n)] \\ &= \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \tilde{\mathbb{E}}_0^{\Psi^{\epsilon,\alpha}} [\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon,\alpha}^n); \tilde{\pi}_{n-1} \geq \delta] \\ &\quad + \tilde{\mathbb{E}}_0^{\Psi^{\epsilon,\alpha}} [\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon,\alpha}^n); \tilde{\pi}_{n-1} < \delta]. \end{aligned} \quad (\text{C.13})$$

Consider the second expectation on the right hand side of the above expression. Note

that $\tilde{\pi}_{n-1} < \delta$ implies

$$|\hat{\phi}_n - \phi_0| = |\tilde{\pi}_{n-1}(\phi_1 - \phi_0)| < \delta|\phi_1 - \phi_0| < \phi_0 - \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}.$$

Hence, conditional on $\tilde{\pi}_{n-1} < \delta$, we must have

$$\hat{\phi}_n > \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}.$$

From Proposition 4.4 and from the fact that $\epsilon < |\bar{\gamma} - \underline{\gamma}|$, we have $\Psi_{\epsilon, \alpha}^n \equiv \tilde{\Gamma}_0^*$, i.e., the insurer's selection is the same with that of the clairvoyant. Therefore,

$$\tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}}[\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon, \alpha}^n); \tilde{\pi}_{n-1} < \delta] = 0.$$

We now turn our attention to the first expectation on the right hand side of (C.13). Let \underline{v}_0 be a lower bound on $\tilde{v}_0(\Gamma)$ over all Γ that is feasible to the problem $\mathcal{IP}_{\epsilon, \alpha}(\rho, \hat{\phi}_n)$ and over all $\hat{\phi}_n \in [0, 1]$. Using Markov's inequality, we have:

$$\begin{aligned} \tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}}[\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon, \alpha}^n); \tilde{\pi}_{n-1} \geq \delta] &\leq (\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0) \tilde{\mathbb{P}}_0^{\Psi^{\epsilon, \alpha}}(\tilde{\pi}_{n-1} \geq \delta) \\ &\leq (\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0) \frac{\tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}}[\tilde{\pi}_{n-1}]}{\delta}. \end{aligned}$$

As a result,

$$\begin{aligned} \tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N) &\leq \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \frac{(\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0)}{\delta} \tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}}[\tilde{\pi}_{n-1}] \\ &\leq \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \frac{(\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0)}{\delta} \mu_0 \exp(-\mu_1 n) \\ &= \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \frac{(\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0)}{\delta} \mu_0 \frac{1 - \exp(-\mu_1 N)}{1 - \exp(-\mu_1)}, \end{aligned}$$

where the second inequality follows from Proposition 4.8. The expression on the right hand side of the above inequality is increasing in N and converges to a constant as $N \rightarrow \infty$. Therefore, there exists $C > 0$ such that $\tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N) \leq C$ for all $N \geq 1$.

We now consider the case where

$$\phi_0 < \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}. \quad (\text{C.14})$$

From Proposition 4.4, the clairvoyant offers a pooling contract with quality level $\underline{\gamma}$ for both types of customers. Furthermore, the common premium level T_p is given by:

$$T_p = U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma}).$$

Let $\epsilon > 0$ be sufficiently small so that

$$\epsilon < \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)} - \phi_0. \quad (\text{C.15})$$

Also, let $\delta > 0$ be such that

$$\delta < \frac{\epsilon}{|\phi_1 - \phi_0|}. \quad (\text{C.16})$$

Similar to the previous case, we can express $\tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N)$ as follows:

$$\begin{aligned} \tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N) &= \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}} [\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon, \alpha}^n); \tilde{\pi}_{n-1} \geq \delta] \\ &\quad + \tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}} [\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon, \alpha}^n); \tilde{\pi}_{n-1} < \delta]. \end{aligned}$$

The first expectation can be bounded using Markov's inequality as in the previous case. For the second expectation, note that $\tilde{\pi}_{n-1} < \delta$ implies

$$|\hat{\phi}_n - \phi_0| = |\tilde{\pi}_{n-1}(\phi_1 - \phi_0)| < \delta|\phi_1 - \phi_0| < \epsilon. \quad (\text{C.17})$$

Hence,

$$\hat{\phi}_n < \phi_0 + \epsilon < \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)}. \quad (\text{C.18})$$

Thus, under $\Psi_{\epsilon, \alpha}$, at decision epoch n , the insurer offers a menu of contracts with quality levels $\underline{\gamma} + \epsilon_n(\alpha)$ and $\underline{\gamma}$ for low type and high type customers, respectively.

The premium levels for low type and high type customers are, respectively:

$$T_L = U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma} + \epsilon_n(\alpha)),$$

$$T_H = U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma} + \epsilon_n(\alpha)) + U(\beta_H, \underline{\gamma} + \epsilon_n(\alpha)) - U(\beta_H, \underline{\gamma}).$$

From Lemma C.1, there exists $C_1 \geq 0$ independent of n such that

$$\tilde{v}_0(\tilde{\Gamma}_0^*) - \tilde{v}_0(\Psi_{\epsilon, \alpha}^n) = C_1 \epsilon_n(\alpha) = C_1 \epsilon \exp(-\alpha n). \quad (\text{C.19})$$

Hence, we have

$$\begin{aligned} \tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N) &\leq \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \left[\frac{(\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0)}{\delta} \tilde{\mathbb{E}}_0^{\Psi^{\epsilon, \alpha}}[\tilde{\pi}_{n-1}] + C_1 \epsilon \exp(-\alpha n) \right] \\ &\leq \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \sum_{n=1}^N \left[\frac{(\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0)}{\delta} \mu_0 \exp(-\mu_1 n) + C_1 \epsilon \exp(-\alpha n) \right] \\ &= \frac{1}{\tilde{v}_0(\tilde{\Gamma}_0^*)} \left[\frac{(\tilde{v}_0(\tilde{\Gamma}_0^*) - \underline{v}_0)}{\delta} \mu_0 \frac{1 - \exp(-\mu_1 N)}{1 - \exp(-\mu_1)} + C_1 \epsilon \frac{1 - \exp(-\alpha N)}{1 - \exp(-\alpha)} \right], \end{aligned}$$

where the second inequality follows from Proposition 4.8. The expression on the right hand side of the above equation is increasing in N and converges to a constant as $N \rightarrow \infty$. Thus, we have that there exists a constant C such that $\tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N) \leq C$ for all $N \geq 1$.

Finally, for the special case where

$$\phi_0 = \frac{(\beta_L - 1)f(\beta_L)}{\beta_L f(\beta_H) - f(\beta_L)},$$

one can verify that $\tilde{v}_0(\tilde{\Gamma}_0^*) = \tilde{v}_0(\Psi_{\epsilon, \alpha}^n)$ for $n \geq 1$, so the proof of a constant upper bound on $\tilde{\Delta}_0^{\Psi^{\epsilon, \alpha}}(N)$ is trivial. \square

Lemma C.1. *For any $\epsilon > 0$, consider a pooling contract $\Gamma = \{(\underline{\gamma}, T_p)\}$ and a separating contract $\Gamma_\epsilon = \{(\underline{\gamma} + \epsilon, T_L), (\underline{\gamma}, T_H)\}$, where the premium levels T_p , T_L and*

T_H are given by

$$T_p = U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma}),$$

$$T_L = U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma} + \epsilon),$$

$$T_H = U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma} + \epsilon) + U(\beta_H, \underline{\gamma} + \epsilon) - U(\beta_H, \underline{\gamma}).$$

There exists a constant C (independent of ϵ) such that $\tilde{v}_0(\Gamma) - \tilde{v}_0(\Gamma_\epsilon) = C\epsilon$.

Proof. By definition of $\tilde{v}_0(\cdot)$, we have

$$\begin{aligned} \tilde{v}_0(\Gamma) &= (1 - \phi_0) [T_p - R_L(\underline{\gamma})] + \phi_0 [T_p - R_H(\underline{\gamma})], \\ &= T_p - (1 - \phi_0)R_L(\underline{\gamma}) - \phi_0 R_H(\underline{\gamma}), \\ &= U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma}) - (1 - \phi_0)\rho(1 - \underline{\gamma})IL_L - \phi_0\rho(1 - \underline{\gamma})IL_H, \end{aligned}$$

where $IL_L = IL(t^*(\beta_L), q^*(\beta_L))$ and $IL_H = IL(t^*(\beta_H), q^*(\beta_H))$. Similarly,

$$\begin{aligned} \tilde{v}_0(\Gamma_\epsilon) &= (1 - \phi_0) [T_L - R_L(\underline{\gamma})] + \phi_0 [T_H - R_H(\underline{\gamma})], \\ &= (1 - \phi_0)T_L + \phi_0 T_H - (1 - \phi_0)R_L(\underline{\gamma} + \epsilon) - \phi_0 R_H(\underline{\gamma}), \\ &= U(\beta_L, \gamma_0) - U(\beta_L, \underline{\gamma} + \epsilon) + \phi_0 [U(\beta_H, \underline{\gamma} + \epsilon) - U(\beta_H, \underline{\gamma})] \\ &\quad - (1 - \phi_0)\rho(1 - \underline{\gamma} - \epsilon)IL_L - \phi_0\rho(1 - \underline{\gamma})IL_H. \end{aligned}$$

Hence,

$$\begin{aligned} \tilde{v}_0(\Gamma) - \tilde{v}_0(\Gamma_\epsilon) &= U(\beta_L, \underline{\gamma} + \epsilon) - U(\beta_L, \underline{\gamma}) \\ &\quad - \phi_0 [U(\beta_H, \underline{\gamma} + \epsilon) - U(\beta_H, \underline{\gamma})] + (1 - \phi_0)\rho IL_L \epsilon, \\ &= [f(\beta_L) - \phi_0 f(\beta_H) + (1 - \phi_0)\rho IL_L] \epsilon. \end{aligned}$$

□

Bibliography

- Adachi, D., Nakata, H., Sawada, Y., and Sekiguchi, K. (2016), “Adverse selection and moral hazard in the corporate insurance market: Evidence from the 2011 Thailand floods,” Tech. rep., The Research Institute of Economy, Trade and Industry.
- AGCS (2015), “Global claims review 2015: Business interruption in focus,” Tech. rep.
- Aldous, D. and Diaconis, P. (1999), “Longest increasing subsequences: from patience sorting to the Baik-Deift-Johansson theorem,” *Bull. Amer. Math. Soc. (N.S.)*, 36, 413–432.
- Araman, V. F. and Caldentey, R. (2009), “Dynamic pricing for nonperishable products with demand learning,” *Operations Research*, 57, 1169–1188.
- Arlotto, A. and Steele, J. M. (2011), “Optimal sequential selection of a unimodal subsequence of a random sequence,” *Combin. Probab. Comput.*, 20, 799–814.
- Arlotto, A., Gans, N., and Steele, J. M. (2014), “Markov Decision Problems Where Means Bound Variances,” *Operations Research*, 62, 864–875.
- Arlotto, A., Nguyen, V. V., and Steele, J. M. (2015), “Optimal online selection of a monotone subsequence: a central limit theorem,” *Stochastic Processes and their Applications*, 125, 3596–3622.
- Ausubel, L. M. and Baranov, O. V. (2010), “Core-selecting auctions with incomplete information,” Working paper, University of Maryland.
- Ausubel, L. M. and Cramton, P. (2002), “Demand reduction and inefficiency in multi-unit auctions,” Working paper, University of Maryland.
- Ausubel, L. M. and Milgrom, P. (2002), “Ascending auctions with package bidding,” *Frontiers of Theoretical Economics*.
- Ausubel, L. M. and Milgrom, P. (2006), “The lovely but lonely Vickrey auction,” *Combinatorial Auctions*, 17, 22–26.

- Baik, J., Deift, P., and Johansson, K. (1999), “On the distribution of the length of the longest increasing subsequence of random permutations,” *J. Amer. Math. Soc.*, 12, 1119–1178.
- Bandi, C. and Bertsimas, D. (2014), “Optimal design for multi-item auctions: A robust optimization approach,” *Mathematics of Operations Research*, 39, 1012–1038.
- Beck, M. and Ott, M. (2013), “Incentives for overbidding in minimum-revenue core-selecting auctions,” Working paper, Stanford University.
- Belloni, A., Chen, B., and Sun, P. (2016), “Computation of optimal dynamic mechanism with participation requirements,” Working paper, Duke University.
- Ben-Tal, A. and Nemirovski, A. (2002), “Robust optimization—methodology and applications,” *Mathematical Programming*, 92, 453–480.
- Bertsimas, D. and Brown, D. B. (2009), “Constructing uncertainty sets for robust linear optimization,” *Operations Research*, 57, 1483–1495.
- Bertsimas, D. and Sim, M. (2004), “The price of robustness,” *Operations Research*, 52, 35–53.
- Besbes, O. and Zeevi, A. (2009), “Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms,” *Operations Research*, 57, 1407–1420.
- Bollobás, B. and Brightwell, G. (1992), “The height of a random partial order: concentration of measure,” *Ann. Appl. Probab.*, 2, 1009–1018.
- Bollobás, B. and Janson, S. (1997), “On the length of the longest increasing subsequence in a random permutation,” in *Combinatorics, geometry and probability (Cambridge, 1993)*, pp. 121–128, Cambridge Univ. Press, Cambridge.
- Bose, S., Ozdenoren, E., and Pape, A. (2006), “Optimal auctions with ambiguity,” *Theoretical Economics*, 1, 411–438.
- Boyd, S. and Vandenberghe, L. (2004), *Convex optimization*, Cambridge University Press.
- Broder, J. and Rusmevichientong, P. (2012), “Dynamic pricing under a general parametric choice model,” *Operations Research*, 60, 965–980.
- Brown, B. M. (1971), “Martingale central limit theorems,” *Ann. Math. Statist.*, 42, 59–66.

- Bruss, F. T. and Delbaen, F. (2001), “Optimal rules for the sequential selection of monotone subsequences of maximum expected length,” *Stochastic Process. Appl.*, 96, 313–342.
- Bruss, F. T. and Delbaen, F. (2004), “A central limit theorem for the optimal selection process for monotone subsequences of maximum expected length,” *Stochastic Process. Appl.*, 114, 287–311.
- Bruss, F. T. and Robertson, J. B. (1991), ““Wald’s lemma” for sums of order statistics of i.i.d. random variables,” *Adv. in Appl. Probab.*, 23, 612–623.
- Caldentey, R. and Haugh, M. (2006), “Optimal control and hedging of operations in the presence of financial markets,” *Mathematics of Operations Research*, 31, 285–304.
- Caldentey, R. and Haugh, M. (2009), “Supply contracts with financial hedging,” *Operations Research*, 57, 47–65.
- Chade, H. and Schlee, E. (2012), “Optimal insurance with adverse selection,” *Theoretical Economics*, 7, 571–607.
- Chod, J., Rudi, N., and Van Mieghem, J. A. (2010), “Operational flexibility and financial hedging: Complements or substitutes?” *Management Science*, 56, 1030–1045.
- Clarke, E. H. (1971), “Multipart pricing of public goods,” *Public Choice*, 11, 17–33.
- Day, R. and Cramton, P. (2012), “Quadratic core-selecting payment rules for combinatorial auctions,” *Operations Research*, 60, 588–603.
- Day, R. and Milgrom, P. (2008), “Core-selecting package auctions,” *International Journal of Game Theory*, 36, 393–407.
- den Boer, A. V. and Zwart, B. (2013), “Simultaneously learning and optimizing using controlled variance pricing,” *Management Science*, 60, 770–783.
- Ding, Q., Dong, L., and Kouvelis, P. (2007), “On the integration of production and financial hedging decisions in global markets,” *Operations Research*, 55, 470–489.
- Dong, L. and Tomlin, B. (2012), “Managing disruption risk: The interplay between operations and insurance,” *Management Science*, 58, 1898–1915.
- Dong, L., Tang, S. Y., and Tomlin, B. (2015), “Production Chain Disruptions: Inventory, Preparedness, and Insurance,” Working paper, Washington University, St. Louis.

- Edelman, B., Ostrovsky, M., and Schwarz, M. (2005), “Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords,” Tech. rep., National Bureau of Economic Research.
- Farias, V. F. and Van Roy, B. (2010), “Dynamic pricing with a prior on market response,” *Operations Research*, 58, 16–29.
- Flajolet, P. and Sedgewick, R. (2009), *Analytic combinatorics*, Cambridge University Press, Cambridge.
- Gaur, V. and Seshadri, S. (2005), “Hedging inventory risk through market instruments,” *Manufacturing & Service Operations Management*, 7, 103–120.
- Gilboa, I. and Schmeidler, D. (1989), “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18, 141–153.
- Gnedin, A. V. (1999), “Sequential selection of an increasing subsequence from a sample of random size,” *J. Appl. Probab.*, 36, 1074–1085.
- Goeree, J. K. and Lien, Y. (2016), “On the impossibility of core-selecting auctions,” *Theoretical Economics*, 11, 44–52.
- Groves, T. (1973), “Incentives in teams,” *Econometrica*, 41, 617–631.
- Guesnerie, R. and Laffont, J.-J. (1984), “A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm,” *Journal of public Economics*, 25, 329–369.
- Hall, P. and Heyde, C. C. (1980), *Martingale limit theory and its application*, Academic Press, Inc., New York-London, Probability and Mathematical Statistics.
- Hammersley, J. M. (1972), “A few seedlings of research,” in *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971), Vol. I: Theory of statistics*, pp. 345–394, Berkeley, CA, Univ. California Press.
- Harrison, J. M., Keskin, N. B., and Zeevi, A. (2012), “Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution,” *Management Science*, 58, 570–586.
- Hellwig, M. F. (2010), “Incentive problems with unidimensional hidden characteristics: A unified approach,” *Econometrica*, 78, 1201–1237.
- Iancu, D. A., Trichakis, N., and Tsoukalas, G. (2016), “Is operating flexibility harmful under debt?” *Management Science*, 63, 1730–1761.
- Jacquet, P. and Szpankowski, W. (1998), “Analytical de-Poissonization and its applications,” *Theoret. Comput. Sci.*, 201, 1–62.

- Jullien, B. (2000), “Participation constraints in adverse selection models,” *Journal of Economic Theory*, 93, 1–47.
- Keskin, N. B. and Birge, J. (2016), “Dynamic selling mechanisms for product differentiation and learning,” Working paper.
- Keskin, N. B. and Zeevi, A. (2014), “Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies,” *Operations Research*, 62, 1142–1167.
- Kim, J. H. (1996), “On increasing subsequences of random permutations,” *J. Combin. Theory Ser. A*, 76, 148–155.
- Kingman, J. F. C. (1973), “Subadditive ergodic theory,” *Ann. Probability*, 1, 883–909.
- Korevaar, J. (2004), *Tauberian theory: a century of developments*, vol. 329 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, Springer-Verlag, Berlin.
- Kouvelis, P., Li, R., and Ding, Q. (2013), “Managing storable commodity risks: The role of inventory and financial hedge,” *Manufacturing & Service Operations Management*, 15, 507–521.
- Krishna, V. (2009), *Auction theory*, Academic Press.
- Liberatore, M. J. (1979), “The EOQ model under stochastic lead time,” *Operations Research*, 27, 391–396.
- Lo, K. C. (1998), “Sealed bid auctions with uncertainty averse bidders,” *Economic Theory*, 12, 1–20.
- Lobo, M. S. and Boyd, S. (2003), “Pricing and learning with uncertain demand,” in *INFORMS Revenue Management Conference*.
- Logan, B. F. and Shepp, L. A. (1977), “A variational problem for random Young tableaux,” *Advances in Math.*, 26, 206–222.
- Maskin, E. and Riley, J. (1984), “Monopoly with incomplete information,” *The RAND Journal of Economics*, 15, 171–196.
- Matthews, S. and Moore, J. (1987), “Monopoly provision of quality and warranties: An exploration in the theory of multidimensional screening,” *Econometrica*, 55, 441–467.
- McLeish, D. L. (1974), “Dependent central limit theorems and invariance principles,” *Ann. Probab.*, 2, 620–628.

- Milgrom, P. and Weber, R. (1982), “A theory of auctions and competitive bidding,” *Econometrica*, 50, 1089–1122.
- Mussa, M. and Rosen, S. (1978), “Monopoly and product quality,” *Journal of Economic Theory*, 18, 301–317.
- Myerson, R. B. (1981), “Optimal auction design,” *Mathematics of Operations Research*, 6, 58–73.
- Myerson, R. B. (1982), “Optimal coordination mechanisms in generalized principal-agent problems,” *Journal of Mathematical Economics*, 10, 67–81.
- Nöldeke, G. and Samuelson, L. (2007), “Optimal bunching without optimal control,” *Journal of Economic Theory*, 134, 405–420.
- Odlyzko, A. M. and Rains, E. M. (2000), “On longest increasing subsequences in random permutations,” in *Analysis, geometry, number theory: the mathematics of Leon Ehrenpreis (Philadelphia, PA, 1998)*, vol. 251 of *Contemp. Math.*, pp. 439–451, Amer. Math. Soc., Providence, RI.
- Oum, Y., Oren, S., and Deng, S. (2006), “Hedging quantity risks with standard power options in a competitive wholesale electricity market,” *Naval Research Logistics*, 53, 697–712.
- Page, F. H. (1992), “Mechanism design for general screening problems with moral hazard,” *Economic Theory*, 2, 265–281.
- Pilpel, S. (1990), “Descending subsequences of random permutations,” *J. Combin. Theory Ser. A*, 53, 96–116.
- Rhee, W. and Talagrand, M. (1991), “A note on the selection of random variables under a sum constraint,” *J. Appl. Probab.*, 28, 919–923.
- Romik, D. (2014), *The Surprising Mathematics of Longest Increasing Subsequences*, Cambridge University Press, Cambridge.
- Salo, A. and Weber, M. (1995), “Ambiguity aversion in first-price sealed-bid auctions,” *Journal of Risk and Uncertainty*, 11, 123–137.
- Samuels, S. M. and Steele, J. M. (1981), “Optimal sequential selection of a monotone sequence from a random sample,” *Ann. Probab.*, 9, 937–947.
- Schmitt, A. J., Snyder, L. V., and Shen, Z.-J. M. (2010), “Inventory systems with stochastic demand and supply: Properties and approximations,” *European Journal of Operational Research*, 206, 313–328.
- Serpa, J. C. and Krishnan, H. (2016), “The strategic role of business insurance,” *Management Science*, 63, 384–404.

- Sion, M. (1958), “On general minimax theorems,” *Pacific J. Math*, 8, 171–176.
- Sphicas, G. P. (1982), “On the Solution of an Inventory Model with Variable Lead Times,” *Operations Research*, 30, 404–410.
- Sphicas, G. P. and Nasri, F. (1984), “An inventory model with finite-range stochastic lead times,” *Naval Research Logistics Quarterly*, 31, 609–616.
- Stark, R. M. and Rothkopf, M. H. (1979), “Competitive bidding: A comprehensive bibliography,” *Operations Research*, 27, 364–390.
- Stiglitz, J. E. (1977), “Monopoly, non-linear pricing and imperfect information: The insurance market,” *The Review of Economic Studies*, pp. 407–430.
- Tao, T. (2009), *Poincaré’s legacies, pages from year two of a mathematical blog. Part II*, American Mathematical Society, Providence, RI.
- Ulam, S. M. (1972), “Some Ideas and Prospects in Biomathematics,” *Ann. Rev. Biophys. Bioeng.*, 1, 277–292.
- Veršik, A. M. and Kerov, S. V. (1977), “Asymptotic behavior of the Plancherel measure of the symmetric group and the limit form of Young tableaux,” *Dokl. Akad. Nauk SSSR*, 233, 1024–1027.
- Vickrey, W. (1961), “Counterspeculation, auctions, and competitive sealed tenders,” *The Journal of Finance*, 16, 8–37.
- Wang, Z., Deng, S., and Ye, Y. (2014), “Close the gaps: A learning-while-doing algorithm for single-product revenue management problems,” *Operations Research*, 62, 318–331.
- Williams, D. (1991), *Probability with martingales*, Cambridge University Press.
- Winter, R. A. (2000), “Optimal insurance under moral hazard,” in *Handbook of Insurance*, pp. 155–183, Springer.

Biography

Van Vinh Nguyen was born on November 3, 1986 in Da Nang, the biggest city in Central Vietnam. In 2005, he received the Singapore Scholarship for an undergraduate study in Electrical and Computer Engineering at the National University of Singapore (NUS). In 2009, he graduated from NUS with a first-class honor. Right after that, he joined the Singapore-MIT Alliance program and received its Graduate Fellowship for a dual master's degree. He graduated in 2010 with a Master of Science degree in Computation for Design and Optimization from MIT and a Master of Engineering degree in Computational Engineering from NUS. In 2013, he enrolled in the PhD program at the Fuqua School of Business, Duke University, majoring in decision sciences. His PhD study was supported by the Fuqua School of Business Doctoral Fellowship. He is scheduled to complete his PhD in Business Administration in September, 2017.