UNIVERSITÉ DU LUXEMBOURG

PhD-FDEF-2025-027

Faculty of Law, Economics and Finance

DISSERTATION

Defence held on 25 September 2025 in Luxembourg
to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG
EN SCIENCES ECONOMIQUES

by

Rishikesh PARMA

Born on 14 January 1996 in Yavatmal, Maharashtra (India)

Resource Allocation Under Uncertainty

Dissertation Defence Committee

- Dr. Joachim Arts, Chair
  Professor - *University of Luxembourg*
- Dr. Çağıl Koçyiğit
  Professor - *University of Luxembourg*
- Dr. Benny Mantin
  Professor - *University of Luxembourg*
- Dr. Ilan Lobel
  Professor - *New York University*
- Dr. Wolfram Wiesemann
  Professor - *Imperial College London*

Affidavit / Statement of originality

I declare that this thesis:

- is the result of my own work. Any contribution from any other party, and any use of generative artificial intelligence technologies have been duly cited and acknowledged;
- is not substantially the same as any other that I have submitted, and;
- is not being concurrently submitted for a degree, diploma or other qualification at the University of Luxembourg or any other University or similar institution except as specified in the text.

With my approval I furthermore confirm the following:

- I have adhered to the rules set out in the University of Luxembourg's Code of Conduct and the Doctoral Education Agreement (DEA)[1], in particular with regard to Research Integrity.
- I have documented all methods, data, and processes truthfully and fully.
- I have mentioned all the significant contributors to the work.
- I am aware that the work may be screened electronically for originality.

I acknowledge that if any issues are raised regarding good research practices based on the review of the thesis, the examination may be postponed pending the outcome of any investigation of such issues. If a degree was conferred, any such subsequently discovered issues may result in the cancellation of the degree.

Approved on 2025-08-08

[1] If applicable (DEA is compulsory since August 2020)

During my four years of doctoral studies, I have learned that the most complex problems often have surprisingly simple solutions.

Acknowledgments

This thesis is not just the result of my efforts, but it is a reflection of the support, guidance, and kindness of so many remarkable people.

First and foremost, I wish to express my deepest gratitude to Prof. Dr. Çağıl Koçyiğit. Your guidance and endless patience have been pivotal in my PhD journey. Your constructive feedback has always inspired me to grow as a researcher. I am equally grateful to Prof. Dr. Benny Mantin. Thank you for the many opportunities you gave me to learn, explore, and contribute to research to the real world problems. To Prof. Dr. Joachim Arts, thank you for your insightful suggestions and for supporting me through my defence. I always deeply enjoyed our research discussions. I would also like to thank the members of my defence committee. Prof. Dr. Wolfram Wiesemann, I am especially grateful not only for your invaluable comments on my work but also for the generosity you showed while hosting me during my research visit. I am also truly honored to have Prof. Dr. Ilan Lobel as a committee member. Thank you for your time, your thoughtful feedback, and your invaluable contributions to my work.

This amazing journey would not have been the same without my colleagues. Neeraj, Poulad, Saber, Junxia, Tiffany, Matteo, Kristof, Jackie, Daria, Laura, Elshan, Daniel, Nihat, Ranit, Roozbeh, and Bikey —your camaraderie and friendship created an atmosphere of learning and laughter that I will always cherish.

A very special thanks goes to my dear friends Neeraj and Ranjani, and Anupam and Sweety: your unwavering emotional support turned the toughest days into lighter ones. I am grateful to all my teammates at the University of Luxembourg Cricket Club—Utsav, Beeshmoy, Ronak, and all others—for the countless unforgettable moments, both on and off the field.

Above all, I owe everything to my parents, Bharti and Sanjay. Your unconditional love and constant support have been the backbone of my journey. You have taught me to believe in hard work and to stand up after facing setbacks. Without your faith in me and your endless encouragement, none of this would have been possible.

Last but by no means least, I am deeply grateful to Prof. Dr. Nishant Mishra. Your recommendation opened the door that allowed me to begin this PhD journey at the University of Luxembourg.

Contents

List of Figures

List of Tables

Abstract

Effective resource allocation is a central challenge in operations research and economics. In practice, allocation decisions are often made under uncertainty, making it important to design mechanisms that perform well under different scenarios. From a practical point of view, it is also important that these mechanisms are easy to understand and implement in real-world settings. This dissertation explores decision-making in two problems of resource allocation under uncertainty in the domains of market design and revenue management. In the first part, we investigate an auction design problem where a seller aims to allocate a single item among multiple bidders with independent private values. In our setting, the seller has only limited information on the distribution of the bidders' values in that they only know an upper bound on their values. The goal is to design a deterministic auction mechanism that performs well across a wide class of distributions. We propose a second-price auction with a reserve price equal to half of the upper bound. While no deterministic mechanism can guarantee a positive fraction of the maximum achievable expected revenue across all possible distributions, we prove that our mechanism secures at least one-quarter of the maximum achievable expected revenue when bidders' values are independent, and the guarantee improves to one-half under the additional assumption that the bidders' values are identically distributed. Numerical experiments with randomly generated distributions show that our mechanism outperforms benchmarks from the literature in approximately 95% of the tested instances. We also account for estimation errors in the upper bound and demonstrate that a slight adjustment to the reserve price results in similar robustness guarantees. The second part studies the capacity allocation problem for a maritime asset provider who allocates vessel capacity between requests for capacity from contracted clients whose demand is uncertain but priced at pre-specified unit rates—and a spot market. Motivated by the operations of a major European ferry operator, we model the problem as an integer program with a concave objective function capturing expected revenues from both channels. We analyze a convex relaxation of the problem and characterize the optimality conditions, which reveal that optimal solutions equate the marginal revenues across contracted customers while ensuring these exceed or match the marginal revenue of the spot market. Building on this insight, we design a computationally efficient algorithm and prove that it yields optimal allocations. Using real data, we show that the algorithm consistently improves revenues in comparison to current practices of the ferry operator, with expected gains of 13% to 55% across sailings.

Chapter 1

Introduction

Whenever a scarce resource is to be allocated among multiple competing entities (agents), the owner typically has an objective in mind—such as maximizing expected revenue or balancing utilization targets—and seeks to allocate the asset, or portions thereof, accordingly. [1] The inherent uncertainty in agents' actions and the possible realizations of future demand over time make such allocation decisions particularly challenging. For instance, airlines need to decide how many seats to reserve for high-paying last-minute travelers versus early bookers; hospitals allocate intensive care unit beds while facing uncertainty in the arrival rate and severity of emergency cases; and governments face the problem of distributing a limited budget across ministries while anticipating uncertain future events (that may require additional expenses) and uncertain revenue generation from taxes. Making the right allocation decisions in these and similar scenarios can be critically important. In all of these contexts, resources are limited, the outcomes are uncertain, and the allocation decision has to be made in advance with only limited information about the environment. Such problems inevitably involve trade-offs, and, as a consequence, "suboptimal" allocation can lead to underutilized resources, forgone revenue, or unmet commitments.

But what does it mean to allocate resources in an "optimal" way? In many real-world settings, detailed information is simply not available to the decision maker, who must therefore rely on prior knowledge when making decisions. Consider first the case of allocating a single resource, such as selling an item through an auction. When the asset owner is selling the unit through an auction, they have only one opportunity to sell the item, and need to decide on a mechanism that maximizes expected revenue based on the available information (e.g., distribution of buyers' value for the item). However, when this information is limited or imprecise (e.g., only the first and second moments of distribution are known), the seller likely considers the worst-case scenarios

[1] The owner is not obliged to allocate the entire asset, as illustrated by the airline example: the overall load factor (the percentage of available passenger capacity that is actually sold and used) across passenger aircrafts was 83.5% in 2024 (https://www.iata.org/en/pressroom/2025-releases/2025-01-30-01/)

when making decisions. In such cases, the objective shifts toward designing a mechanism that guarantees some level of performance, even under adverse conditions. For multiple-resource allocation problems, the challenges are richer and more intricate. A seller may need to choose a mechanism to distribute resources among multiple customers without complete knowledge of their demand, or to determine how much capacity to retain for a spot market that could generate additional revenue in the future.

The choice of such a mechanism involves not only performance under the constraint of limited information but also simplicity. On the one hand, the seller seeks a mechanism that performs well—maximizing revenue, social welfare, or guarantees on revenue generation under uncertainty. On the other hand, mechanisms must also be easy to understand, transparent to participants, and simple to implement in practice. A highly sophisticated mechanism may promise superior theoretical guarantees but risk being impractical, therefore lacking the potential to make an actual impact. Conversely, overly simple mechanisms may be easy to implement but perform poorly in critical scenarios. Striking the right balance between these two objectives naturally raises the question: how can we determine a guiding principle for allocating resources through a practical way that aligns with the chosen performance objective?

This thesis discusses practically simple and easy-to-implement mechanisms and algorithms for resource allocation and decision-making under uncertainty in multi-agent environments. The overarching goal is to exploit structural properties of the problems to design computationally efficient allocation rules with rigorous performance guarantees.

1.1 Thesis Structure

In the following chapters, we illustrate both the theoretical foundations and the real-world challenges of resource allocation under uncertainty in two settings. Before turning to these contributions, it is useful to place them in the broader landscape of resource allocation problems. In general, such problems differ by the source of uncertainty (e.g., buyers' willingness to pay, magnitude of demand), the type of information available (e.g., bounds on the potential buyers' willingness to pay, historical observations), the nature of the environment (e.g., single-shot, repeated interactions), and the objectives pursued (e.g., maximize revenue, profit, welfare, or market share, sometimes under worst-case scenarios). Within this geography, the first problem we study involves a single-item auction design problem while facing uncertainty about bidders' willingness to pay, while the second problem focuses on capacity allocation in shipping under uncertain demand across two sales channels. We elaborate on these two below.

In Chapter 2, we study an auction design problem, where a seller aims to sell a single item to multiple bidders. Bidders' willingness to pay (values) are independent random variables, and the seller only knows an upper bound on these values, lacking specific distribution knowledge. The objective is to devise a deterministic auction mechanism effective across a broad set of distributions. We propose a second price auction with a reserve price set to half of the upper bound. Even though no deterministic mechanism can achieve a positive fraction of the maximum achievable expected revenue across all distributions, we define a distribution class $\mathscr{G}$, which is extensive and contains several distributions of practical importance with diverse structural properties, and we demonstrate that our mechanism achieves at least $\frac{1}{4}$ ($\frac{1}{2}$ under i.i.d. values) of the maximum expected revenue for distributions within $\mathscr{G}$. We conduct numerical experiments to evaluate our mechanism's performance beyond $\mathscr{G}$, under randomly generated distributions, demonstrating its superior performance in approximately 95% of the generated instances compared to benchmark mechanisms from the literature. Additionally, we illustrate numerically that our mechanism exhibits greater robustness against different correlations than the benchmarks when considering two non-independent bidders. We also consider the scenario where the estimated upper bound is subject to errors and show that appropriately lowering the reserve price based on estimation confidence ensures a constant positive fraction of the maximum expected revenue across $\mathscr{G}$. Traditional auction design strategies, under distributional ambiguity, often propose randomized mechanisms. While these can provide theoretical guarantees, such approaches are relatively rarely used in practice, and they often lack interpretability and transparency, potentially leading to trust issues among bidders and implementation challenges. Our deterministic mechanism is simple, requires minimal information, and is provably effective in various practical scenarios.

The contents of Chapter 2 can be found in the following paper:

- *Parma R., Koçyiğit Ç., Mantin B. Simple and Effective: A Deterministic Auction with Support Information. Under major revision.*

In Chapter 3, we study the capacity allocation problem for freight shipping faced by a maritime asset provider who allocates vessel capacity through contracts and the spot market under demand uncertainty. Motivated by the practices of our industry partner, a major ferry operator, we formulate this problem as an integer program with integer-valued capacity allocation decisions, aiming to maximize a concave revenue function representing the total expected revenue from contracts and the spot market. We first consider a convex relaxation of this problem and characterize its optimality conditions. These conditions reveal that, at optimality, the marginal revenues (roughly, the rates of change resulting from slightly perturbing the allocated capacity) are equalized across contracted customers and must be greater than or equal to that of the spot market. Leveraging this insight, we propose an easy-to-interpret and computationally efficient algorithm, which extends the Expected Marginal Seat Revenue (EMSR) framework from the airline revenue management literature, to solve the original integer capacity allocation problem. We show that the algorithm is exact in the sense that it produces the optimal allocations. We test the framework using synthetic and real data. Our real-data experiments demonstrate that our algorithm improves the revenues of the ferry operator consistently across all sailings considered, with improvements ranging from 13% to 55%.

The contents of Chapter 3 can be found in the following paper:

- *Parma R., Koçyiğit Ç., Mantin B. Capacity Allocation: Balancing Contract Commitments and Spot Market Opportunities. Working paper.*

Chapter 4 concludes with a summary of the main contributions of the thesis and a discussion of open questions that that can be explored in future research.

Chapter 2

Simple and Effective: A Deterministic Auction with Support Information

2.1 Introduction

Background and Motivation. Auctions have been a fundamental method for trade and exchange since ancient times. With records dating back to the Roman Empire, where auctions were used to sell property and estate goods, today they are employed to sell a wide variety of items, including artwork, collectibles, estates, financial instruments, and commodities. Despite their extensive use since ancient times, theoretical analysis of how to design an auction mechanism to sell an item only began in the twentieth century. Vickrey (1961) introduced and studied the second price auction, where the highest bidder wins the item but pays the second-highest bid. Myerson (1981) modelled bidders' willingness-to-pay (values) for the item as random variables governed by a known distribution and characterized the optimal mechanism that maximizes the expected revenue. If the bidders' values are independent and identically distributed (i.i.d.), this mechanism is a second price auction with an optimally chosen reserve price under a regularity assumption. The optimal mechanism, unsurprisingly, heavily depends on the distribution of values. The traditional mechanism design literature follows this practice and assumes that the distribution of the bidders' values for the item is known. In reality, however, it is difficult to justify precise knowledge of such a distribution. In fact, typically, there is very limited information about this distribution, which arguably is the reason for using an auction mechanism in the first place. There is, therefore, a need to design auctions that do not rely on distributional assumptions, referred to as the "Wilson doctrine" (Wilson, 1987).

In this chapter, we study an auction design problem assuming that the value distribution is unknown. We assume that the seller knows only an upper bound $\overline{v}$ on possible bidders' values; such a bound can be computed by looking at the prices paid in historical sales, by consulting with experts, using domain knowledge or common sense; refer to Cohen et al. (2021, Section 1.2) for a discussion on how to compute this bound in practical situations. Our setting is in particular relevant for scenarios where there is very limited information about demand, such as selling new products or rarely sold items like artwork, antiques, and collectibles. Our goal is to establish an easily understandable deterministic auction (i.e., the allocations are non-randomized) that provably generates high expected revenue across a large class of plausible distributions consistent with the knowledge of the seller. Our focus on deterministic auction mechanisms is motivated by the fact that they are interpretable and transparent, which are crucial for fostering trust among participants, and are, therefore, often preferred in practice.

Main Contributions. We propose a second price auction with a reserve price of $\frac{\bar{v}}{2}$, which is set at half of the upper bound. Our contributions can be summarized as follows.

- Aligned with the majority of the literature, we consider independent (though not necessarily identically distributed) bidders' values and regular distributions. Deterministic mechanisms cannot generate a positive fraction of the maximum expected revenue achievable across this class of distributions because any Dirac distribution that concentrates all probability mass on a single value scenario can be arbitrarily approximated by a sequence of regular beta distributions (see, e.g., Proposition 2.7). However, Dirac distributions represent unrealistic scenarios and are unlikely to occur in practice. We establish a subclass $\mathscr{G}$ of all independent and regular distributions and demonstrate that our proposed deterministic second-price auction generates at least $\frac{1}{4}$, which improves to $\frac{1}{2}$ when bidders' values are i.i.d., of the maximum expected revenue achievable across this distribution class (Theorem 2.8). This matches the guarantee of the (randomized) single-sample mechanism proposed by Dhangwatnotai et al. (2015) for independent and regular distributions, and it can be shown that the guarantee of the single-sample mechanism cannot be improved for $\mathscr{G}$. We also show that $\frac{1}{2}$ is a tight fraction and cannot be improved in the case of i.i.d. bidders' values (Proposition 2.13).

- We demonstrate that the distribution class $\mathscr{G}$ is extensive and contains several distributions of practical importance with diverse structural properties (Proposition 2.15). These include distributions where the marginal value distribution of each bidder belongs to the family of all triangular distributions, all twice-differentiable convex distributions, substantial subclasses of exponential, Pareto, and beta distributions, among others. To gain insight into the distribution class $\mathscr{G}$ and justify its practical relevance, we also provide examples using beta distributions, showing that distributions with sufficiently large expected values or dispersions are more likely to fall within $\mathscr{G}$. It is important to note that when both the expected value and the dispersion are low, the distribution concentrates most of its mass around lower values, resulting in a lower maximum achievable revenue. Such scenarios arguably represent situations where the seller would be less concerned about achieving revenues close to optimal. Furthermore, we characterize distribution-dependent bounds on the fraction of the expected revenue of our proposed mechanism compared to the maximum expected revenue achievable under different parametric distributions (Table 2.1).

- We extend our results regarding revenue guarantees to scenarios where estimating the true upper bound is challenging. Specifically, we assume that the estimated upper bound $\hat{v}$ is subject to estimation errors, and the true bound $\bar{v}$ falls within the range $[\hat{v}-\varepsilon, \hat{v}+\varepsilon]$, where $\varepsilon > 0$ represents the estimation confidence. In this case, we demonstrate that setting the reserve price equal to $\frac{\hat{v}-\varepsilon}{2}$ generates at least $\frac{\hat{v}-\varepsilon}{4(\hat{v}+\varepsilon)}$, which improves to $\frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)}$ in case of i.i.d. bidder values, times the maximum expected revenue achievable over $\mathscr{G}$ (Proposition 2.19).

- To investigate the performance of our mechanism beyond the distribution class $\mathscr{G}$, we conduct numerical experiments using randomly generated value distributions, not necessarily restricted to belong to $\mathscr{G}$ or regular distributions. Our numerical experiments demonstrate that our proposed mechanism outperforms two benchmark mechanisms—the single-sample mechanism and the second price auction without a reserve price—in approximately 95% of the generated instances. Additionally, we conduct a numerical study involving two non-independent bidders, which reveals that our proposed mechanism exhibits greater robustness against different correlations compared to the two benchmarks.

Related Literature. Our work contributes to the mechanism design literature focused on designing auctions under the assumption that the distribution of bidders' values is unknown and only partial information about this distribution is available. Various notions of partial information have been studied in the literature, including bounds on possible bidders' values (Bergemann and Schlag (2008), Bandi and Bertsimas (2014), Cohen et al. (2021), Anunrojwong et al. (2023), Rujeerapaiboon et al. (2023), Koçyiğit et al. (2024)), information about statistical moments of the value distribution (Carrasco et al. (2018), Koçyiğit et al. (2020), Chen et al. (2022), Suzdaltsev (2022), Giannakopoulos et al. (2023), Wang et al. (2024)), availability of a nominal distribution assumed to be close to the true value distribution (Bergemann and Schlag (2011), Chen et al. (2024)), and in the presence of samples from the value distribution (Dhangwatnotai et al. (2015), Allouah and Besbes (2020), Allouah et al. (2022)) or quantile observations (Azar et al. (2013), Allouah et al. (2023)).

To address the mechanism design problem under distributional ambiguity, two popular approaches have emerged: (distributionally) robust mechanism design and approximately optimal (also known as prior-independent) mechanism design. The literature on robust mechanism design evaluates mechanisms based on their performance under the most adverse distribution consistent with the available information. The objective typically involves either maximizing the worst-case expected revenue or minimizing the worst-case expected regret, which represents the difference between the hypothetical revenue under complete information and the actual revenue generated, in view of all distributions consistent with the available information (see, e.g., Chen et al. (2024), Koçyiğit et al. (2024) and references therein). The literature on approximately optimal mechanism design seeks to find mechanisms that consistently ensure an objective function value (such as expected revenue) as close as possible to a benchmark value derived under full information (such as the maximum expected revenue achievable), under any distribution consistent with the available knowledge (see, e.g., Dhangwatnotai et al. (2015), Allouah and Besbes (2020) and references therein).

The vast majority of the underlying literature proposes randomized mechanisms (i.e., the allocation and payment decisions are probabilistic and depend on random variables), whereas the study of deterministic mechanisms remains very limited. While randomized mechanisms may have inherent advantages such as better worst-case performance guarantees, they suffer from several drawbacks: implementability (e.g., the literature sometimes proposes random reserve prices, but it is unclear how to implement a random reserve price in practice), interpretability (randomized allocations and payments may be difficult to understand), and transparency (e.g., random reserve prices are only revealed ex-post). In addition, randomized mechanisms may suffer from time inconsistency: once the random outcome is realized, the seller may have an incentive to "roll the dice again" if the realized reserve price or allocation is unfavorable. These drawbacks can potentially lead to confusion and mistrust among bidders. Deterministic mechanisms, by contrast, are often more interpretable and transparent, and therefore more practically relevant.

Thus, our objective is aligned with the literature on approximately optimal mechanism design, but our work also draws inspiration from existing results in robust mechanism design. Most related to our work are (Cohen et al., 2021) and (Koçyiğit et al., 2024). Cohen et al. (2021) study a monopoly pricing problem under the assumption that the seller knows only an upper bound on the price that can be charged. They propose a deterministic posted price computed as if the demand curve was linear and equal to half of the upper bound if there is no cost to the seller. They show that this price yields good revenue compared to the maximum expected revenue achievable under several parametric distributions of demand. Koçyiğit et al. (2024), among other settings, study the minimax regret auction design problem over deterministic mechanisms. When the seller knows only an upper bound on the possible bidders' values, they show that the second price auction with a reserve price set to half of the upper bound minimizes the worst-case regret in view of all possible bidders' values. Note that this second price auction reduces to the posted price proposed by Cohen et al. (2021) when there is only one bidder. In this chapter, we study the same informational setting as in (Cohen et al., 2021) and (Koçyiğit et al., 2024), and we investigate the expected revenue of the second price auction with a reserve price set to half of the upper bound in comparison to the maximum expected revenue achievable in the multi-bidder auction design setting. Our results extend and generalize those of Cohen et al. (2021) to the auction setting and complement those of Koçyiğit et al. (2024), strengthening the performance guarantees of the proposed second price auction.

Our results also complement the study of deterministic mechanisms under different informational settings. Azar and Micali (2012) and Azar et al. (2013) propose deterministic auctions guaranteed to generate a constant fraction of the maximum expected revenue achievable in the presence of available statistical information such as mean, variance, and median. Giannakopoulos et al. (2023) consider the case where there is a single buyer, and the seller knows the mean values and an upper bound on the standard deviation. They, among other results, characterize the deterministic price that generates the maximum fraction of the achievable revenue in the worst-case distribution scenario consistent with the available information. In all of these papers, the proposed mechanisms rely on the statistical information (e.g., mean and variance) available. A positive fraction of the maximum expected revenue achievable is attainable in this case because of the additional statistical information. In our setting, there is no such information available. Note that without such information, no deterministic mechanism can attain a positive fraction of the maximum expected revenue across all distributions even if they are constrained to be regular.

Additionally, our results are related to Dhangwatnotai et al. (2015) who propose a mechanism, called the single-sample mechanism, guaranteed to generate at least a $\frac{1}{4}$ of the maximum expected revenue achievable throughout the class of all independent and regular distributions. Their proposed mechanism is randomized because it relies on (random) samples from bidders' value distribution, and the implementation of this mechanism requires that for any bidder, there exists at least one more bidder such that the values of these two bidders are i.i.d. We consider the single-sample mechanism as a benchmark mechanism in our analysis. In our setting, we require the upper bound on the bidders' values to be the same, but the value distribution of each bidder is allowed to be unique, permitting asymmetries across all bidders. Furthermore, we focus on deterministic mechanisms.

Structure of the chapter. The remainder of the chapter is structured as follows. Section 2.2 defines the problem and introduces preliminaries. Section 2.3 defines the distribution class $\mathscr{G}$ and studies the performance guarantees for our proposed mechanism across $\mathscr{G}$ (Section 2.3.1), the generality of the distribution class $\mathscr{G}$ (Section 2.3.2), distribution-dependent performance under different parametric distributions (Section 2.3.3), and the case of uncertainty in the estimated upper bound (Section 2.3.4). Section 2.4 presents the numerical experiments. Section 2.5 concludes. All proofs are deferred to the appendix.

Notation. For any $\boldsymbol{v} \in \mathbb{R}^N$, we write v_i to denote its i^{th} component and $\boldsymbol{v}_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_N)$ to denote its subvector excluding v_i. Random vectors are designated by tilde signs ($e.g., \tilde{\boldsymbol{v}}$) and their realizations are denoted by the same symbols without tildes ($e.g., \boldsymbol{v}$). For any Borel set $\mathscr{A}$, we use $\mathscr{P}_0(\mathscr{A})$ to represent the set of all probability distributions on $\mathscr{A}$. The family of all bounded Borel-measurable functions from a Borel set $\mathscr{A}$ to another Borel set $\mathscr{C}$ is denoted by $\mathscr{L}(\mathscr{A},\mathscr{C})$.

2.2 Problem Formulation and Preliminaries

We consider an auction design problem, where a seller aims to sell a single item to N bidders indexed by $i \in \mathscr{N} = \{1,\dots,N\}$. The bidders' values, $\tilde{v}_i$, $i \in \mathscr{N}$, are assumed to be independent random variables governed by a joint (cumulative) distribution $\mathbb{F} = \times_{i\in\mathscr{N}}\mathbb{F}_i$, where $\mathbb{F}_i$ denotes the distribution of the value $\tilde{v}_i$ of bidder i. We denote by $\tilde{\boldsymbol{v}} = (\tilde{v}_1,\dots,\tilde{v}_N)$ the vector of all bidders' values. We assume that the seller does not know the distribution $\mathbb{F}$ and only knows an upper bound $\overline{v}$ on the bidders' values. We denote by $\mathscr{V}^N$, where $\mathscr{V} = [0,\overline{v}]$, the set of possible bidders' values.

The seller aims to design an auction to sell the good that achieves a good performance across a wide class of distributions. An auction is a mechanism $M = (\boldsymbol{x},\boldsymbol{p})$, which consists of an allocation rule $\boldsymbol{x} \in \mathscr{L}(\mathscr{V}^N,[0,1]^N)$ and a payment rule $\boldsymbol{p} \in \mathscr{L}(\mathscr{V}^N,\mathbb{R}^N)$. Each bidder i, depending on his value v_i, simultaneously reports a bid $b_i \in \mathbb{R}$ to the seller. An allocation rule, given the vector of bids $\boldsymbol{b} \in \mathbb{R}^N$, determines the probability $x_i(b_i,\boldsymbol{b}_{-i})$ of allocation to bidder i, and the payment rule determines the amount $p_i(b_i,\boldsymbol{b}_{-i})$ charged to this bidder. Note that the allocation and payment rules depend only on the reported bids $\boldsymbol{b}$ and not on the true values $\boldsymbol{v}$. Given a mechanism M, bidders choose their bids strategically to maximize their utilities.

Definition 2.1 (Ex-post Utility)**.** The ex-post utility of bidder i with reported bid b_i and value v_i is defined as

$$u_i(b_i;v_i,\boldsymbol{b}_{-i}) = x_i(b_i,\boldsymbol{b}_{-i})v_i - p_i(b_i,\boldsymbol{b}_{-i}),$$

where $\boldsymbol{b}_{-i}$ represents the vector of bids reported by the others.

The ex-post utility denotes the bidder's expected payoff after all bids are revealed. We do not make any assumption on the bidders' beliefs about the value distribution and, for this reason, restrict attention to mechanisms that admit an ex-post Nash equilibrium bidding strategy.

Definition 2.2 (Ex-post Nash Equilibrium)**.** An N-tuple $\boldsymbol{\beta} = (\beta_1,\dots,\beta_N)$ of bidding strategies $\beta_i \in \mathscr{L}(\mathscr{V},\mathbb{R})$, $i \in \mathscr{N}$, is called an ex-post Nash equilibrium if

$$\begin{aligned} x_i(\beta_i(v_i),\boldsymbol{\beta}_{-i}(\boldsymbol{v}_{-i}))v_i - p_i(\beta_i(v_i),\boldsymbol{\beta}_{-i}(\boldsymbol{v}_{-i})) \geq x_i(b_i,\boldsymbol{\beta}_{-i}(\boldsymbol{v}_{-i}))v_i - p_i(b_i,\boldsymbol{\beta}_{-i}(\boldsymbol{v}_{-i}))\\ \forall i \in \mathscr{N}, \forall \boldsymbol{v} \in \mathscr{V}^N, \forall b_i \in \mathbb{R}. \end{aligned}$$

In plain words, under an ex-post Nash equilibrium, each bidder i maximizes their utility simultaneously for all value scenarios $\boldsymbol{v} \in \mathscr{V}^N$ and bidder i cannot improve their utility by deviating from their chosen bid, even after knowing the bids of all other bidders. The notion of ex-post Nash equilibrium does not rely on the value distribution $\mathbb{F}$ and is therefore natural to use by a seller who lacks knowledge of $\mathbb{F}$ and the bidders' beliefs about this distribution. One could show that an ex-post Nash equilibrium bidding strategy $\boldsymbol{\beta}$ is a Bayesian Nash equilibrium bidding strategy under any $\mathbb{G} = \times_{i\in\mathcal{N}}\mathbb{G}_i \in \mathscr{P}_0(\mathscr{V}^N)$, where $\tilde{v}_i \sim \mathbb{G}_i \in \mathscr{P}_0(\mathscr{V})$. This means that even if the bidders estimate the value distribution (accurately or inaccurately) and aim to maximize their expected utilities with respect to their estimated value distribution, the vector of bidding strategies $\boldsymbol{\beta}$ remains a Nash equilibrium.

The set of all mechanisms under which there exists an ex-post Nash equilibrium is vast and contains, for example, the first price auctions and the second price auctions with reserve prices. The class of all such mechanisms is referred as the class of dominant strategy incentive compatible (DSIC) mechanisms.

Definition 2.3 (Dominant Strategy Incentive Compatibility)**.** A mechanism $M = (\boldsymbol{x}, \boldsymbol{p})$ is called dominant strategy incentive compatible (DSIC) if

$$x_i(v_i, \boldsymbol{v}_{-i})v_i - p_i(v_i, \boldsymbol{v}_{-i}) \geq x_i(\hat{v}_i, \boldsymbol{v}_{-i})v_i - p_i(\hat{v}_i, \boldsymbol{v}_{-i}) \quad \forall i \in \mathcal{N}, \forall \boldsymbol{v} \in \mathscr{V}^N, \forall \hat{v}_i \in \mathscr{V}.$$

Under a DSIC mechanism, truthful reporting is an ex-post Nash equilibrium. In other words, each bidder i chooses to bid his true value v_i. By invoking the revelation principle (Myerson, 1981), the seller can focus on DSIC mechanisms without loss of generality. The revelation principle states that, given a mechanism and an equilibrium bidding strategy for that mechanism, there exists a mechanism in which (i) it is an equilibrium strategy for each bidder to report his true value, and (ii) the ex-post payoffs of all agents remain the same for every value profile $\boldsymbol{v} \in \mathscr{V}^N$.

Formally, we study the following set of mechanisms that adhere to the principles of incentive compatibility (IC), individual rationality (IR) and allocation feasibility (AF).

$$\mathscr{M} = \left\{ \begin{aligned} M = (\boldsymbol{x}, \boldsymbol{p}) : & x_i(v_i, \boldsymbol{v}_{-i})v_i - p_i(v_i, \boldsymbol{v}_{-i}) \geq x_i(\hat{v}_i, \boldsymbol{v}_{-i})v_i - p_i(\hat{v}_i, \boldsymbol{v}_{-i}) \\ & \qquad \forall i \in \mathcal{N}, \forall \boldsymbol{v} \in \mathscr{V}^N, \forall \hat{v}_i \in \mathscr{V} \quad \text{(IC)} \\ & x_i(v_i, \boldsymbol{v}_{-i})v_i - p_i(v_i, \boldsymbol{v}_{-i}) \geq 0 \quad \forall i \in \mathcal{N}, \forall \boldsymbol{v} \in \mathscr{V}^N \quad \text{(IR)} \\ & \sum_{i=1}^{N} x_i(v_i, \boldsymbol{v}_{-i}) \leq 1 \quad \forall \boldsymbol{v} \in \mathscr{V}^N \quad \text{(AF)} \end{aligned} \right\}$$

The first inequality (IC) ensures that the mechanism is DSIC and, therefore, truthful bidding is optimal for the bidders. The inequality (IC) can be enforced without loss of generality as discussed above. The second inequality (IR) ensures that the mechanism is ex-post individually rational in the sense that the bidders, by participating in the auction, receive non-negative payoffs for every value profile $\boldsymbol{v} \in \mathscr{V}^N$. The last inequality (AF) captures the allocation feasibility, i.e., it is not feasible to sell more than one good for every $\boldsymbol{v} \in \mathscr{V}^N$.

We denote by $\text{Rev}(M, \mathbb{F})$ the expected revenue of mechanism M under the value distribution $\mathbb{F}$, which is defined as

$$\text{Rev}(M, \mathbb{F}) = \mathbb{E}_{\mathbb{F}}\left[\sum_{i=1}^{N} p_i(\tilde{v}_i, \tilde{\boldsymbol{v}}_{-i})\right],$$

and denote by $\text{Rev}^\star(\mathbb{F})$ the optimal expected revenue, which is defined as

$$\text{Rev}^\star(\mathbb{F}) = \sup_{M \in \mathscr{M}} \mathbb{E}_{\mathbb{F}}\left[\sum_{i=1}^{N} p_i(\tilde{v}_i, \tilde{\boldsymbol{v}}_{-i})\right].$$

We next introduce the second price auctions with reserve prices, which belong to the class $\mathscr{M}$ of mechanisms.

Definition 2.4 (Second Price Auction with a Reserve Price)**.** The second price auction $M^{\text{sp}}(r) = (\boldsymbol{x}^{\text{sp}}, \boldsymbol{p}^{\text{sp}})$ with a reserve price r is defined as

$$x_i^{\text{sp}}(v_i, \boldsymbol{v}_{-i}) = \begin{cases} 1 & \text{if } i = \min \arg\max_{j \in \mathscr{N}} v_j \text{ and } v_i \geq r \\ 0 & \text{otherwise} \end{cases}$$

and

$$p_i^{\text{sp}}(v_i, \boldsymbol{v}_{-i}) = \begin{cases} \max\{\max_{j \neq i}(v_j), r\} & \text{if } x_i^{\text{sp}}(v_i, \boldsymbol{v}_{-i}) = 1 \\ 0 & \text{otherwise.} \end{cases}$$

A second price auction $M^{\text{sp}}(r)$ with a reserve price r allocates the item to the highest bidder if the highest bid is greater than or equal to the reserve price r, and the seller keeps the item if the highest bid falls below r. In Definition 2.4, we use a lexicographic tie-breaker to break the ties but our results do not depend on the particular choice of the tie-breaking rule. If the item is allocated, bidder i who has the highest bid pays the maximum of the second highest bid and the reserve price r, i.e., $\max\{\max_{j \neq i}(v_j), r\}$. It is well-known that a second price auction equipped with an appropriately chosen reserve price is optimal under the standard assumption of a known regular distribution $\mathbb{F}$ given that the bidders' values are independent and identically distributed (Myerson, 1981). Unfortunately, the optimal reserve price depends on the distribution $\mathbb{F}$ and is unidentifiable when $\mathbb{F}$ is unknown. Even though $\mathbb{F}$ is unknown to the seller in our setting, we will maintain the standard assumption on $\mathbb{F}$ and assume that it is regular. This class is essential to the fields of pricing and mechanism design, as it induces a concavity condition on the revenue function.

Definition 2.5 (Regular Distribution)**.** A distribution $\mathbb{F} = \times_{i \in \mathscr{N}} \mathbb{F}_i$ is called regular if the marginal density f_i of $\tilde{v}_i$ exists and is strictly positive on $(0,\overline{v})$, and the virtual value function

$$\varphi_i(v_i) = v_i - \frac{1-\mathbb{F}_i(v_i)}{f_i(v_i)}$$

is non-decreasing in v_i on $(0,\overline{v})$ for every $i \in \mathscr{N}$.

The class of regular distributions is extensive, and it contains all distributions with a monotonic hazard rate (MHR), such as truncated normal, exponential, and logistic, as well as non-MHR distributions, including log-normal and Pareto. For a summary of commonly used regular distributions, see Bagnoli and Bergstrom (2005) and Ewerhart (2013). Some of the highly regarded results in mechanism design require the regularity condition. For example, in Myerson (1981), a second price auction with a reserve price is a revenue-maximizing mechanism if the underlying value distribution is regular. From now on, we denote by $\mathscr{R}(\mathscr{V}^N)$ the set of all regular distributions on $\mathscr{V}^N$.

Throughout the chapter we will assess the performance of a mechanism via its competitive ratio that is the ratio of its expected revenue to the optimal expected revenue attainable.

Definition 2.6 (Competitive Ratio)**.** The competitive ratio $\gamma(M,\mathbb{F})$ of a mechanism M under distribution $\mathbb{F} \in \mathscr{R}(\mathscr{V}^N)$ is defined as

$$\gamma(M,\mathbb{F}) = \frac{\text{Rev}(M,\mathbb{F})}{\text{Rev}^\star(\mathbb{F})}.$$

The regularity condition on $\mathbb{F}$ ensures that $\text{Rev}^\star(\mathbb{F}) > 0$, and the competitive ratio $\gamma(M,\mathbb{F})$ always lies in $[0,1]$. The greater the ratio, the better the performance of the mechanism. We say that a mechanism M is an α-approximation under distribution $\mathbb{F}$ if $\gamma(M,\mathbb{F}) = \frac{1}{\alpha}$ and that the mechanism is an α-approximation over a set $\mathscr{P} \subseteq \mathscr{R}(\mathscr{V}^N)$ of distributions if $\gamma(M,\mathbb{F}) \geq \frac{1}{\alpha}$ for all $\mathbb{F} \in \mathscr{P}$.

Our goal is to characterize a simple and deterministic mechanism that is a constant-factor approximation of the optimal mechanism over a large class of distributions. We define the aforementioned distribution class in the next section and elaborate on its generality. In the next section we propose a second price auction with a deterministic reserve price and show that the proposed mechanism guarantees at least a $\frac{1}{4}$-fraction of the optimal revenue under every distribution in the aforementioned distribution class, that is, it is a 4-approximation. Furthermore, we show that if bidders' values are i.i.d., then our proposed mechanism is a 2-approximation.

2.3 Proposed Mechanism and Performance

We propose a second price auction where the reserve price is set to $\frac{\overline{v}}{2}$, denoted as $M^{\text{sp}}\left(\frac{\overline{v}}{2}\right)$. In this section, we study the performance of $M^{\text{sp}}\left(\frac{\overline{v}}{2}\right)$ in terms of the competitive ratio and show that the proposed second price auction is a 4-approximation (and 2-approximation for i.i.d. bidders' values) over a large class of distributions.

Specifically, we focus on the class $\mathscr{R}(\mathscr{V}^N)$ of regular distributions, as discussed earlier. It is recognized and easy to verify, that no deterministic mechanism can provide a strictly positive approximation guarantee under every distribution in $\mathscr{R}(\mathscr{V}^N)$. The following proposition formalizes this notion.

Proposition 2.7. *No deterministic mechanism $M \in \mathscr{M}$ can guarantee a strictly positive worst-case competitive ratio over the class $\mathscr{R}(\mathscr{V}^N)$ of regular distributions and $N \geq 1$. Specifically, for $N = 1$, we have $\inf_{\mathbb{F}\in\mathscr{R}(\mathscr{V})} \gamma(M,\mathbb{F}) = 0$ for any deterministic mechanism $M \in \mathscr{M}$.*

The proof of Proposition 2.7 relies on considering a sequence of regular beta distributions that converge to a Dirac point mass in the limit. Any such Dirac point mass is not particularly interesting and unrealistic to expect from a practical standpoint. We therefore introduce a new class of distributions under which our proposed deterministic mechanism provides a positive approximation guarantee.

$$\mathscr{G}(\mathscr{V}^N) = \left\{ \begin{array}{ll} \mathbb{F} \in \mathscr{P}_0(\mathscr{V}^N): & \mathbb{F} = \times_{i\in\mathscr{N}}\mathbb{F}_i \in \mathscr{R}(\mathscr{V}^N) \\ & \mathbb{F}_i(\frac{\overline{v}}{2}) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \leq 1 + \frac{1}{\overline{v}}\int_0^{\overline{v}} \mathbb{F}_i(v_i)\big(1-\mathbb{F}_i(v_i)\big)\,\mathrm{d}v_i \;\; \forall i \in \mathscr{N} \end{array} \right\} \tag{2.1}$$

The distribution class $\mathscr{G}(\mathscr{V}^N)$ is extensive and contains several distributions with different structural properties as we will discuss later. Some examples of distributions that belong to $\mathscr{G}([0,1])$ are shown in Figure 2.1. Independent coupling of any of these distributions belongs to $\mathscr{G}([0,1]^N)$.

In Section 2.3.1, we formally introduce our approximation result and prove it. In Section 2.3.2, we study the distribution class $\mathscr{G}(\mathscr{V}^N)$ and show that it contains several distributions of practical importance including triangular, the class of twice-differentiable convex and substantial subclasses of beta and concave distributions. In Section 3.3, we study distribution-dependent competitive ratios for well-known parametric distributions such as triangular, truncated exponential, and power distributions. In Section 3.4, we extend our approximation results to the case where the upper bound $\overline{v}$ is uncertain.

2.3.1 Revenue Approximation

The next theorem formalizes our main approximation result.

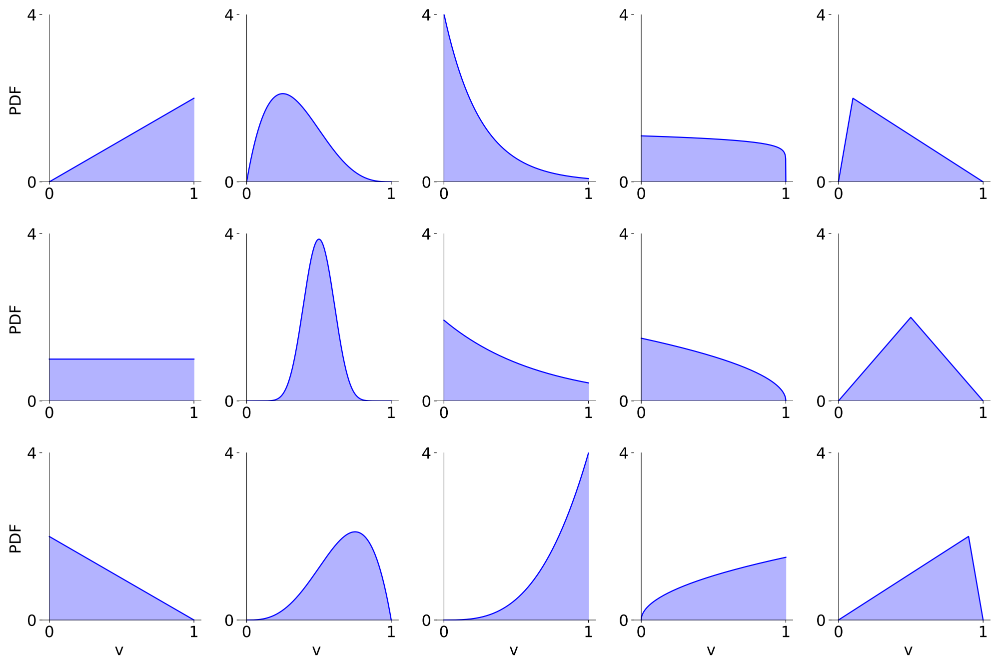

Figure 2.1: Examples of probability density functions whose cumulative distribution functions belong to the set $\mathscr{G}([0,1])$

Theorem 2.8. *For any number $N \geq 1$ of bidders and $\mathbb{F} \in \mathscr{G}(\mathscr{V}^N)$, the second price auction $M^{\text{sp}}(\frac{\bar{v}}{2})$ with reserve price $\frac{\bar{v}}{2}$ is at least a 4-approximation under $\mathbb{F}$, i.e., $\gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}), \mathbb{F}\big) \geq \frac{1}{4}$. If the bidders' values are i.i.d under $\mathbb{F}$, then $M^{\text{sp}}(\frac{\bar{v}}{2})$ is at least a 2-approximation, i.e., $\gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}), \mathbb{F}\big) \geq \frac{1}{2}$.*

The rest of this subsection is dedicated to outlining the proof idea behind Theorem 2.8 that relies on two established results from the literature. Theorem 2.9 demonstrates that if the bidders' values are i.i.d. and if their marginal value distributions are regular, then the seller benefits more by attracting an additional bidder and implementing a second price auction without a reserve price, compared to implementing the optimal mechanism with the original number of bidders.

Theorem 2.9 (Bulow and Klemperer (1996))**.** *For any number $N \geq 1$ of bidders and $\mathbb{F}_1 \in \mathscr{R}(\mathscr{V})$, we have* $\text{Rev}(M^{\text{sp}}(0), \mathbb{F}_1^{N+1}) \geq \text{Rev}^\star(\mathbb{F}_1^N)$.

Hartline and Roughgarden (2009) and Dhangwatnotai et al. (2015) presented an interesting interpretation of Theorem 2.9 stating that for a bidder with a value drawn from a regular distribution $\mathbb{F}_1$, the expected revenue of a posted price mechanism (which is a second price auction implemented for a single bidder, where the reserve price represents the posted price) where this price is randomly drawn from $\mathbb{F}_1$ is at least half of that from an optimal mechanism, i.e.,

$$\gamma\Big(M^{sp}(\tilde{r} \sim \mathbb{F}_1), \mathbb{F}_1\Big) \geq \frac{1}{2}. \tag{2.2}$$

To develop intuition, consider a second price auction with two bidders, where the values are independently and identically drawn from a distribution $\mathbb{F}_1$. In this setup, under a second price auction without reserve price, each bidder contributes an equal expected revenue and in some sense confronts a reserve price determined by the other bidder's value, specifically, a random reserve price drawn from $\mathbb{F}_1$. According to Theorem 2.9, the expected revenue generated from selling the good to a single bidder by posting a random price sampled from $\mathbb{F}_1$ is therefore guaranteed to be at least half of the revenue generated from selling the good to the same bidder by posting an optimally chosen price. In the proof of Theorem 2.8, we show that if $\mathbb{F}_1 \in \mathscr{G}(\mathscr{V})$, then our proposed mechanism $M^{\text{sp}}(\frac{\bar{v}}{2})$ outperforms $M^{\text{sp}}(\tilde{r} \sim \mathbb{F}_1)$ in the case of a single bidder and thus $M^{\text{sp}}(\frac{\bar{v}}{2})$ has a 2-approximation guarantee by (2.2).

To extend the approximation guarantee to the multi-bidder case, we study a variant of the second price auction. In second price auctions where different reserve prices are set for different bidders, two approaches emerge based on when these reserves come into play. Second price auctions with *eager* reserves eliminate bidders below their reserve right after bids are submitted, and then run the second price auction without a reserve price. On the other hand, second price auctions with *lazy* reserves do not eliminate these bids and offer the item to the highest bidder at a take-it-or-leave-it price set to the maximum of his reserve and the second-highest bid. When the reserve price is the same for all bidders, the two auction forms are identical. Note that in a second price auction with lazy reserves, the item can only be allocated to the highest bidder, whereas this is not true for eager reserves. In the following, we will only be interested in lazy reserves and denote by $\hat{M}^{\text{sp}}(\boldsymbol{r})$ the second price auction with lazy reserve prices $\boldsymbol{r} = (r_1, r_2 \ldots r_N)$.

Definition 2.10 (Second Price Auction with Lazy Reserve Prices)**.** Consider reserve price vector $\boldsymbol{r} = (r_1, r_2 \ldots r_N)$, where r_i denotes the bidder-specific lazy reserve price for bidder $i \in \mathscr{N}$.

The second price auction with lazy reserve prices $\boldsymbol{r}$ is defined as $\hat{M}^{\text{sp}}(\boldsymbol{r}) = (\hat{\boldsymbol{x}}^{\text{sp}}, \hat{\boldsymbol{p}}^{\text{sp}})$, where

$$\hat{x}_i^{\text{sp}}(v_i, \boldsymbol{v}_{-i}) = \begin{cases} 1 & \text{if } i = \min \arg\max_{j \in \mathscr{N}} v_j \text{ and } v_i \geq r_i \\ 0 & \text{otherwise} \end{cases}$$

and

$$\hat{p}_i^{\text{sp}}(v_i, \boldsymbol{v}_{-i}) = \begin{cases} \max\{\max_{j \neq i}(v_j), r_i\} & \text{if } \hat{x}_i^{\text{sp}}(v_i, \boldsymbol{v}_{-i}) = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Collecting bidders' values, this mechanism works as follows: (*i*) Determine the candidate winner, $i^\star = \min \arg\max_{i \in \mathscr{N}} v_i$ (we use a lexicographic tie-breaker to resolve ties, favoring the bidder with the smallest index), (*ii*) Allocate the good to bidder $i^\star$ if $v_{i^\star} \geq r_{i^\star}$, and charge the winning bidder $i^\star$ the maximum of $r_{i^\star}$ and the second-highest bid.

The following Theorems 2.11 and 2.12 from Dhangwatnotai et al. (2015) and Azar et al. (2013), respectively, along with the definition of $\hat{M}^{\text{sp}}(\boldsymbol{r})$ will help us to provide an approximation guarantee in the multi-bidder case.

Theorem 2.11. *(Dhangwatnotai et al., 2015, Theorem 3.17) For any number $N \geq 1$ of bidders and $\mathbb{F} = \times_{i \in \mathcal{N}} \mathbb{F}_i \in \mathscr{R}(\mathscr{V}^N)$, the second price auction with the bidder-specific lazy reserve prices $\boldsymbol{r}^\star = (r_1^\star, \ldots, r_N^\star)$ is at least a 2-approximation under $\mathbb{F}$, i.e., $\gamma\big(\hat{M}^{\text{sp}}(\boldsymbol{r}^\star), \mathbb{F}\big) \geq \frac{1}{2}$, where $r_i^\star \in \arg\max_{p \in \mathbb{R}} p(1 - \mathbb{F}_i(p))$ for all $i \in \mathcal{N}$.*

Theorem 2.11 shows that the second price auction with lazy reserve prices set to optimal monopoly prices is at least a 2-approximation under the regularity condition. Building upon this observation, the following theorem outlines a method for determining whether an arbitrary reserve price can offer an approximation guarantee.

Theorem 2.12. *(Azar et al., 2013, Theorem 3.1) Let $r \in \mathbb{R}_+$ be any deterministic reserve price. For any number $N \geq 1$ of bidders and $\mathbb{F} = \times_{i \in \mathcal{N}} \mathbb{F}_i \in \mathscr{R}(\mathscr{V}^N)$, if $\gamma\big(M^{\text{sp}}(r), \mathbb{F}_i\big) \geq \eta$ for all $i \in \mathcal{N}$ and $\gamma\big(\hat{M}^{\text{sp}}(\boldsymbol{r}^\star), \mathbb{F}\big) \geq \theta$, then we have $\gamma\big(M^{\text{sp}}(r), \mathbb{F}\big) \geq \eta\theta$.*

The proof of Theorem 2.12, which is omitted in this chapter, follows from the proof of Theorem 3.1 in Azar et al. (2013) even though its formal statement is slightly different. In particular, the proof shows that if $\gamma\big(M^{\text{sp}}(r), \mathbb{F}_i\big) \geq \eta$ for all $i \in \mathcal{N}$, then $M^{\text{sp}}(r)$ is an $\frac{1}{\eta}$-approximation of $\hat{M}^{\text{sp}}(\boldsymbol{r}^\star)$ for any number N of bidders. Then, the fact that $\gamma\big(M^{\text{sp}}(r), \mathbb{F}\big) \geq \eta\theta$ follows from the assumption that $\gamma\big(\hat{M}^{\text{sp}}(\boldsymbol{r}^\star), \mathbb{F}\big) \geq \theta$. Theorems 2.11 and 2.12 together imply that, under the regularity condition, if selling the item to any of the bidders using the posted price mechanism with price r earns at least η-fraction of the expected revenue that could have been achieved with the optimal price, then the second price auction with reserve price r is at least a $\frac{2}{\eta}$-approximation of the optimal revenue. As a special case, if bidders are i.i.d., the second price auction with reserve price r is at least an $\frac{1}{\eta}$-approximation to the optimal revenue. (This is because $\hat{M}^{\text{sp}}(\boldsymbol{r}^\star)$ is the optimal mechanism in case of i.i.d. bidders' values.)

The proof idea is to establish that the expected revenue of $M^{\text{sp}}(\frac{\bar{v}}{2})$ for the single-bidder case (with value distribution $\mathbb{F}_i$) is at least as high as that of the benchmark $M^{\text{sp}}(\tilde{r} \sim \mathbb{F}_i)$ if $\mathbb{F}_i \in \mathscr{G}(\mathscr{V})$. By (2.2), this implies that $M^{\text{sp}}(\frac{\bar{v}}{2})$ is at least a 2-approximation when $\mathcal{N} = \{i\}$. We then leverage Theorems 2.11 and 2.12 to extend the result to multiple-bidders case, ensuring a 4-approximation for non-i.i.d. bidders and a 2-approximation for i.i.d. bidders.

In the following, we consider the single-bidder case (i.e., $N = 1$) and show that the 2-approximation bound of Theorem 2.8 is tight. To this end, we illustrate the existence of a sequence $(\mathbb{F}_{1,k})_{k \in \mathbb{Z}_+}$ of distributions such that $\mathbb{F}_{1,k} \in \mathscr{G}([0,1])$ for all $k \in \mathbb{Z}_+$ and $\lim_{k \to \infty} \gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}), \mathbb{F}_{1,k}\big) = \frac{1}{2}$.

Proposition 2.13. *For $N = 1$ and the sequence of distributions $(\mathbb{F}_{1,k})_{k\in\mathbb{Z}_+}$ defined through*

$$\mathbb{F}_{1,k}(v_1) = \begin{cases} 0 & \text{if } v_1 < 0, \\ v_1^k & \text{if } v_1 \in [0,1], \\ 1 & \text{if } v_1 \geq 1, \end{cases}$$

we have $\mathbb{F}_{1,k} \in \mathscr{G}([0,1])$ for all $k \in \mathbb{Z}_+$ and $\lim_{k\to\infty} \gamma(M^{\text{sp}}\left(\frac{1}{2}\right), \mathbb{F}_{1,k}) = \frac{1}{2}$.

Proposition 2.13 provides a lower bound on the approximation guarantee of our proposed mechanism attainable over $\mathscr{G}(\mathscr{V}^N)$ and $N \geq 1$. This implies that it is impossible to improve the 2-approximation bound in the case of i.i.d. bidders. For the asymmetric bidders case, unfortunately, we are not able to show the tightness of 4-approximation bound, and it remains open whether this bound can be improved.

2.3.2 Generality of the Distribution Class $\mathscr{G}(\mathscr{V}^N)$

In this section, we will demonstrate that various practical distributions belong to the distribution class $\mathscr{G}(\mathscr{V}^N)$. To be included in this class, these distributions must satisfy the following inequality condition on their marginal distributions:

$$\mathbb{F}_i\left(\frac{\overline{v}}{2}\right) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \leq 1 + \frac{1}{\overline{v}}\int_0^{\overline{v}} \mathbb{F}_i(v_i)\big(1 - \mathbb{F}_i(v_i)\big)\,\mathrm{d}v_i \;\; \forall i \in \mathscr{N}. \tag{2.3}$$

Before delving into our exploration of distributions, we present a remark that offers a sufficient condition for the relatively complicated inequality condition on the marginal distribution $\mathbb{F}_i$, $i \in \mathscr{N}$, in equation (2.3). This straightforward, yet, insightful sufficient condition holds an easy interpretation and will be used for some of the results presented in this section. We then provide an intuition for the inequality condition (2.3).

Remark 2.14. The inequality condition on the marginal distribution $\mathbb{F}_i$, $i \in \mathscr{N}$, in equation (2.3) is implied by the following sufficient condition: $\mathbb{F}_i(\frac{\overline{v}}{2}) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \leq 1$.

This follows from the fact that the term $\frac{1}{\overline{v}}\int_0^{\overline{v}} \mathbb{F}_i(v_i)\big(1 - \mathbb{F}_i(v_i)\big)\,\mathrm{d}v_i$, which appears on the right-hand side of the inequality condition, is always non-negative. The sufficient condition relies only on two statistics: the cumulative distribution evaluated at $\frac{\overline{v}}{2}$ and the normalized expected value of $\tilde{v}_i$. Equivalently, this sufficient condition can be expressed as $\frac{1}{2}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \leq \frac{\overline{v}}{2}\left(1 - \mathbb{F}_i\left(\frac{\overline{v}}{2}\right)\right)$. Notably, this immediately implies that selling the good to bidder i at the price of $\frac{\overline{v}}{2}$ is at least a 2-approximation when compared to using the optimal price. This is evident as $\frac{\overline{v}}{2}\left(1 - \mathbb{F}_i\left(\frac{\overline{v}}{2}\right)\right)$ represents the expected revenue of the price $\frac{\overline{v}}{2}$, and the expected revenue of the optimal price is bounded from above by

$\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i]$, i.e., the expected revenue of changing the bidder i's true value. The sufficient condition $\mathbb{F}_i(\frac{\overline{v}}{2}) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \leq 1$ holds true if the $1 - \mathbb{F}_i\left(\frac{\overline{v}}{2}\right)$ probability of the event that bidder i's value exceeds our proposed reserve price $\frac{\overline{v}}{2}$ is at least as high as his normalized expected value $\frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i]$. Common probability distributions that satisfy this sufficient condition include twice-differentiable convex, a subclass of beta and (*not all but*) a subclass of triangular distributions—we will later show that all triangular distributions satisfy the original condition (2.3).

To gain intuition on condition (2.3), we will consider standard beta distributions as they are both expressive and practically relevant. Since not all beta distributions fall within the $\mathscr{G}(\mathscr{V}^N)$ class, and only a subset of beta distributions meet condition (2.3), this will enable us to study and interpret the criteria for beta distributions to be part of this distribution class.

Example 1. Assume $\tilde{v}_i$ follows a beta distribution on support $\mathscr{V} = [0,1]$ with its cumulative distribution function (cdf) given by

$$\mathbb{F}_i(v_i) = \frac{\int_0^{v_i} t^{\alpha-1}(1-t)^{\beta-1}\,\mathrm{d}t}{\int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\,\mathrm{d}t},$$

and its expected value $\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i]$ is given by $\frac{\alpha}{\alpha+\beta}$. (Note that the α used in the α-approximation is not the same α parameter utilized in the definition of beta distributions here. We have opted to employ this (repeated) notation to align with the existing literature.) For simplicity we set $\beta = 1$ in this example but our arguments below hold generally as demonstrated in Proposition 2.15. The cdf for this simplified case is equal to $\mathbb{F}_i(v_i) = v_i^{\alpha}$. Thus, we have

$$\mathbb{F}_i\left(\frac{\overline{v}}{2}\right) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] = \left(\frac{1}{2}\right)^{\alpha} + \frac{\alpha}{\alpha+1}.$$

Whenever $\alpha \geq 1$, this expression is smaller than or equal to 1 , therefore satisfying the sufficient condition for condition (2.3); see Remark 2.14. Note that the condition $\alpha \geq 1$ is equivalent to the requirement $\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \geq \frac{1}{2} = \frac{\overline{v}}{2}$ in this case. This means that as long as the expected value of $\tilde{v}_i$ is sufficiently large—specifically, greater than or equal to $\frac{\overline{v}}{2}$—condition (2.3) is satisfied. In general, this holds for any $\alpha \geq \beta$ without the assumption $\beta = 1$; see Proposition 2.15.

In Example 1, we illustrated that if the expected value of $\tilde{v}_i$ is sufficiently large, then its beta distribution $\mathbb{F}_i$ satisfies condition (2.3). In the following, we showcase that even if the expected value of $\tilde{v}_i$ is relatively low, a sufficiently large dispersion causes $\mathbb{F}_i$ to meet the criteria in (2.3).

Example 2. Assume $\tilde{v}_i$ follows a beta distribution with support $\mathscr{V} = [0,1]$. Suppose $\alpha = 1$. In this case, the cdf of $\tilde{v}_i$ is given by $\mathbb{F}_i(v_i) = 1 - (1-v_i)^{\beta}$. We thus have

$$\mathbb{F}_i\left(\frac{\overline{v}}{2}\right) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] = 1 - \left(\frac{1}{2}\right)^{\beta} + \frac{1}{1+\beta}.$$

Furthermore, we have

$$\int_0^1 \mathbb{F}_i(v_i)\,(1-\mathbb{F}_i(v_i))\,\mathrm{d}\,v_i = \int_0^1 (1-(1-v_i)^\beta)(1-v_i)^\beta\,\mathrm{d}\,v_i.$$

By letting $u = 1 - v_i$, we obtain

$$\int_0^1 (1-(1-v_i)^\beta)(1-v_i)^\beta\,\mathrm{d}\,v_i = \int_0^1 (1-u^\beta)u^\beta\,\mathrm{d}\,u = \int_0^1 (u^\beta - u^{2\beta})\,\mathrm{d}\,u = \frac{1}{\beta+1} - \frac{1}{2\beta+1}.$$

Note that as β grows, the expression $\frac{1}{\beta+1} - \frac{1}{2\beta+1}$ approaches zero. Additionally, as β grows, the beta distribution concentrates the probability mass towards zero, and its variance decreases to zero. In fact, there is a close relation between $\int_0^1 \mathbb{F}_i(v_i)\,(1-\mathbb{F}_i(v_i))\,\mathrm{d}\,v_i$ and the distribution's dispersion. The product $\mathbb{F}_i(v_i)\,(1-\mathbb{F}_i(v_i))$ is the product of the probabilities that $\tilde{v}_i$ falls below v_i and that $\tilde{v}_i$ exceeds v_i. Both this product and the integral take higher values when the distribution is well spread out. A larger value for this integral indicates a higher degree of variability or dispersion in the distribution. Through the above derivations, we obtain

$$\begin{aligned} &\mathbb{F}_i\left(\frac{\overline{v}}{2}\right) + \frac{1}{\overline{v}}\mathbb{E}_{\mathbb{F}_i}[\tilde{v}_i] \leq 1 + \int_0^1 \mathbb{F}_i(v_i)\big(1-\mathbb{F}_i(v_i)\big)\,\mathrm{d}\,v_i \\ \iff\ & 1 - \left(\frac{1}{2}\right)^\beta + \frac{1}{1+\beta} \leq 1 + \frac{1}{\beta+1} - \frac{1}{2\beta+1} \quad \iff \quad 2^\beta \leq 2\beta + 1. \end{aligned}$$

This implies that condition (2.3) holds if $2^\beta \leq 2\beta+1$, which is the case, for example, even if $\beta > \alpha = 1$ in which case the expected value of $\tilde{v}_i$ is smaller than $\frac{\overline{v}}{2}$ as long as β remains sufficiently small, i.e., $\lessapprox 2.66$, representing a high-dispersion scenario.

In Examples 1 and 2, we illustrated the connection between condition (2.3) and the distribution's expected value and dispersion through beta distributions. Specifically, we illustrated that distributions of this form with either high expected value or high dispersion satisfy condition (2.3). Converse to this, distributions with low expected value and low dispersion may not belong to $\mathscr{G}(\mathscr{V})$ and from a practical point of view, when both the expected value and dispersion are low, which indicates a concentrated probability mass at low values, the resulting expected revenue is also low with respect to $\overline{v}$. This is arguably a less interesting and less realistic scenario from the seller's point of view.

The next proposition demonstrates that $\mathscr{G}(\mathscr{V})$ contains a diverse range of distributions. By definition of $\mathscr{G}(\mathscr{V}^N)$, note that if $\mathbb{F}_i \in \mathscr{G}(\mathscr{V})$ for all $i \in \mathscr{N}$, then the independent coupling $\mathbb{F} = \times_{i\in\mathscr{N}}\mathbb{F}_i \in \mathscr{G}(\mathscr{V}^N)$.

Proposition 2.15. $\mathbb{F}_1 \in \mathscr{G}(\mathscr{V})$ *if either one of the following conditions is met:*

(i) $\mathbb{F}_1$ *is regular and convex on* $\mathscr{V}$.

(ii) $\mathbb{F}_1$ *is regular and concave on* $\mathscr{V}$, *and its median is at least* $\frac{\sqrt{3}-1}{2}\overline{v}$.

(iii) $\mathbb{F}_1$ *is a triangular distribution on* $\mathscr{V}$.

(iv) $\mathbb{F}_1$ *is a regular scaled beta distribution on* $\mathscr{V}$, *where* $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \geq \frac{\overline{v}}{2}$.

Remark 2.16. A twice differentiable and convex distribution $\mathbb{F}_1 \in \mathscr{G}(\mathscr{V})$ on the interval $\mathscr{V} = [0, \overline{v}]$ is regular. Particularly, we will demonstrate that a twice differentiable convex distribution $\mathbb{F}_1$ has non-decreasing hazard rate, which is a sufficient condition for regularity. The hazard rate of $\mathbb{F}_1$ is given by

$$h(v_1) = \frac{f_1(v_1)}{1 - \mathbb{F}_1(v_1)},$$

where f_1 denotes the density function of $\mathbb{F}_1$. Note that f_1 exists and is strictly positive because $\mathbb{F}_1$ is twice differentiable, which means that $\mathbb{F}_1$ is increasing. To demonstrate that the hazard rate is non-decreasing, we need to show that $h'(t) \geq 0$ for all t in the interval $[0, \overline{v}]$, where $h'(t)$ denotes the derivative of the hazard rate. We can express $h'(t)$ as

$$h'(t) = \frac{f_1'(t)(1 - \mathbb{F}_1(t)) + f_1(t)\mathbb{F}_1'(t)}{(1 - \mathbb{F}_1(t))^2}.$$

Recall that $\mathbb{F}_1$ is twice differentiable, which means that both $\mathbb{F}_1$ and f_1 are differentiable. Since $f_1'(t)$ for a convex $\mathbb{F}_1$ is non-negative, we argue that the numerator is always non-negative, and hence $\mathbb{F}_1$ has a non-decreasing hazard rate and is therefore regular.

Remark 2.17. Regular concave distributions include exponential, substantially large subclasses of Pareto and beta distributions, and truncated regular distributions (Ewerhart, 2013).

Remark 2.18. The probability density function (pdf) of the scaled beta distribution $\mathbb{F}_1$ on $[0, \overline{v}]$ with $\alpha, \beta \geq 0$ is defined as

$$f_1(v_1) = \frac{1}{B(\alpha, \beta, \overline{v})} \left(\frac{v_1}{\overline{v}}\right)^{\alpha-1} \left(1 - \frac{v_1}{\overline{v}}\right)^{\beta-1},$$

where $B(\alpha, \beta, \overline{v})$ is the Beta function (throughout this chapter, to distinguish between beta distributions and Beta functions the former are written with lowercase b and the latter with uppercase B) given by

$$B(\alpha, \beta, \overline{v}) = \int_0^{\overline{v}} \left(\frac{t}{\overline{v}}\right)^{\alpha-1} \left(1 - \frac{t}{\overline{v}}\right)^{\beta-1} dt.$$

The mean μ and variance σ^2 of the scaled beta distribution can be expressed as

$$\mu = \frac{\alpha}{\alpha+\beta}\overline{v} \quad \text{and} \quad \sigma^2 = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}\overline{v}^2.$$

Note that $\mu \geq \frac{\bar{v}}{2}$ if and only if $\alpha \geq \beta$.

Scaled beta distributions with $\alpha, \beta \geq 1$ are regular. In fact, they can be constructed via a linear transformation from the standard beta distribution with support $[0,1]$ and the same α and β, which is known to have a log-concave density function (Bagnoli and Bergstrom, 2005, Table 1). By Bagnoli and Bergstrom (2005, Corollary 2), distributions with log-concave density functions are regular, and by Bagnoli and Bergstrom (2005, Theorem 7), under the linear transformation, the log-concavity of the density function is preserved. This implies that the scaled beta distributions with $\alpha, \beta \geq 1$ have log-concave density functions and are therefore regular.

Proposition 2.15 shows that $\mathscr{G}(\mathscr{V})$ contains various forms of distributions, including convex, concave, and S-shaped distributions (beta and triangular). The distributions in Proposition 2.15 can display a wide range of expected values and dispersions. For instance, if $\mathbb{F}_1$ is convex, then the expected value is high exceeding $\frac{\bar{v}}{2}$, whereas if $\mathbb{F}_1$ is concave, then expected value is low falling below $\frac{\bar{v}}{2}$. For triangular $\mathbb{F}_1$, the expected value varies depending on the mode of the distribution. In cases where the expected value is low, a sufficently large dispersion may categorize the distribution into $\mathscr{G}(\mathscr{V})$.

Even though Proposition 2.15 considers the single-bidder case ($\mathbb{F} = \mathbb{F}_1$), recall that in the case of $N > 1$, the independent coupling $\mathbb{F} = \times_{i \in \mathscr{N}} \mathbb{F}_i$, where $\mathbb{F}_i \in \mathscr{G}(\mathscr{V})$ for all $i \in \mathscr{N}$, belongs to $\mathscr{G}(\mathscr{V}^N)$ by definition of $\mathscr{G}(\mathscr{V}^N)$. We next present an example to illustrate this relation to the multi-bidder case.

Example 3. Consider a case of two independent bidders whose values for the item follow two different distributions in $\mathscr{G}(\mathscr{V})$, where $\mathscr{V} = [0,1]$. The marginal distribution for the first bidder is $\mathbb{F}_1(v_1) = \frac{1-e^{-v_1}}{1-e^{-1}}$, representing a truncated exponential distribution with $\lambda = 1$. This distribution is regular and concave with median of 0.38 $\left(\text{which is greater than} \frac{\sqrt{3}-1}{2}\right)$ and thus belongs to $\mathscr{G}(\mathscr{V})$. For the second bidder, the marginal distribution is $\mathbb{F}_2(v_2) = 2v_2 - v_2^2$, which is a triangular distribution with $c = 0$ and thus belongs to $\mathscr{G}(\mathscr{V})$. Since the bidders' values are independent, the joint distribution of their values can be expressed as $\mathbb{F} = \mathbb{F}_1 \times \mathbb{F}_2$. This joint distribution belongs to $\mathscr{G}(\mathscr{V}^2)$.

An intuitive explanation for the satisfactory performance of $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ across the wide range of distributions presented in this section is closely related to the choice of the reserve price at $\frac{\bar{v}}{2}$. With this reserve price, the seller only needs a sufficient probability that any one of the bidders' values exceeds $\frac{\bar{v}}{2}$ to effectively capitalize on bidders with high values. In contrast, the seller faces the risk of not being able to capitalize on such bidders by setting a high reserve price and the risk of forfeiting profit by setting a low reserve price. Our arguments above indicate that a value distribution in $\mathscr{G}(\mathscr{V}^N)$ likely exhibits high expected values or high dispersions, in which case there is a good chance that any one of the bidders' values exceeds the reserve price $\frac{\bar{v}}{2}$. In practice, high dispersion signifies less confidence in foreseeing the bidders' values, and the presence of this type of uncertainty is arguably the reason for holding auctions in the first place. In fact, high dispersion is only natural to expect in practical situations.

In the next section, we assess the distribution-dependent performance of our proposed mechanism under some commonly used parametric distributions, demonstrating that it can offer better performance under certain cases.

2.3.3 Distribution-dependent Performance under Parametric Distributions

We now study the performance of $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ under some (regular) parametric distribution classes (beta, triangular, power and exponential), which are discussed widely in the literature (Bagnoli and Bergstrom, 2005; Ewerhart, 2013). We characterize the distribution-dependent competitive ratios of this mechanism under the aforementioned distributions. When characterizing the distribution-dependent competitive ratios, we focus on the case of a single bidder ($N = 1$). Note that in this case, our mechanism is a 2-approximation over $\mathscr{G}(\mathscr{V})$. We will discuss how our distribution-dependent competitive ratio results can be extended to the multi-bidder case through an example later in this section.

Table 2.1: Distribution-dependent performance of $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$

Distribution	CDF	Competitive ratio
Beta	$\frac{\int_0^{v_1}\left(\frac{t}{\bar{v}}\right)^{\alpha-1}\left(1-\frac{t}{\bar{v}}\right)^{\beta-1}dt}{B(\alpha,\beta)}$	No closed form expression
Triangular	$\begin{cases} \frac{v_1^2}{\bar{v}c} & \text{for } 0 \leq v_1 \leq c, \\ 1-\frac{(\bar{v}-v_1)^2}{\bar{v}(\bar{v}-c)} & \text{for } c < v_1 \leq \bar{v}. \end{cases}$	$\begin{cases} \frac{27}{32} & \text{for } 0 < c \leq \frac{\bar{v}}{3}, \\ \frac{3\sqrt{3}\bar{v}^2}{16\sqrt{\bar{v}c}(\bar{v}-c)} & \text{for } \frac{\bar{v}}{3} < c \leq \frac{\bar{v}}{2}, \\ \frac{3\sqrt{3}\bar{v}(4c-\bar{v})}{16\sqrt{\bar{v}}c^{\frac{3}{2}}} & \text{for } \frac{1}{2} < c < \bar{v}. \end{cases}$
Power	v_1^k	$\frac{\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right)}{\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)}$
Exponential	$\frac{1-e^{-\lambda v_1}}{1-e^{-\lambda\bar{v}}}$	$\frac{\lambda\bar{v}}{2\left(1-W(e^{1-\lambda\bar{v}})\right)}\left(\frac{e^{\frac{-\lambda\bar{v}}{2}}-e^{-\lambda\bar{v}}}{e^{\left(W(e^{1-\lambda\bar{v}})-1\right)}-e^{-\lambda\bar{v}}}\right)$

Notes: (*i*) All distributions except power distribution are defined on the general support $[0,\bar{v}]$. For power, $\bar{v} = 1$. (*ii*) $\alpha \in [1,\infty)$ and $\beta \in [1,\infty)$ are parameters of beta distribution (see Remark 2.18). (*iii*) $c \in [0,\bar{v}]$ is the mode of triangular distribution. (*iv*) $k \in \mathbb{R}_+$ for power distribution. (*v*) $\lambda \in \mathbb{R}_+$ for exponential distribution and $W(\cdot)$ is the Lambert W function.

Table 2.1 defines the distributions considered and shows the respective competitive ratio formulas whenever available in closed form. Figure 2.2 displays the numerical values of the competitive ratios for the four distribution classes with different parameter values, where $\mathscr{V} = [0,1]$. From Figure 2.2a, we observe that for the class of regular beta distributions, previously shown to belong to $\mathscr{G}([0,1])$ when the expected value exceeds $\frac{1}{2}$ (i.e., $\alpha \geq \beta \geq 1$), our mechanism can guarantee a 2-approximation even when $\beta > \alpha$. From Figure 2.2b, we notice that the competitive ratio is actually much higher than $\frac{1}{2}$ and exceeds 0.8 across all choices of the mode c for triangular distributions. In Figure 2.2c, we observe that, for power distributions, the competitive ratio is very high for small values of k and converges to $\frac{1}{2}$ as k increases. As k grows, the probability mass of the power distribution concentrates towards the upper bound $\bar{v} = 1$, in which case the optimal price tends to the same value as well, explaining this convergence. One can argue that as k increases, the corresponding value distributions become less practically relevant, and for more practical situations where k is relatively small, we can expect our mechanism to perform very well. From Figure 2.2d, we observe that, for exponential distributions, $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ guarantees at least 2-approximation of optimal revenue for the seller when the distribution parameter $\lambda \leq 5.1$. One can show that if $\lambda \lessapprox 3.94$ then the exponential distribution belongs to $\mathscr{G}([0,1])$. However, it is worth emphasizing that the exponential distribution with $\lambda \gtrapprox 3.94$ does not belong to $\mathscr{G}([0,1])$, showcasing that our proposed mechanism can perform well also outside of $\mathscr{G}([0,1])$. As λ grows, the competitive ratio converges to zero. However, note that these situations hold very little practical relevance and appeal because the optimal revenue with the knowledge of the distribution also tends to zero in this case.

We can use the competitive ratios in Table 2.1 to establish a bound on the approximation guarantee for the multi-bidder case. The following example illustrates the way to do this.

Example 4. Suppose that there are two bidders, with $\mathbb{F}_1(v_1) = \frac{1-e^{-2v_1}}{1-e^{-2}}$, $\mathbb{F}_2(v_2) = 2v_2 - v_2^2$, and $\mathbb{F} = \mathbb{F}_1 \times \mathbb{F}_2$. Theorems 2.11 and 2.12 together imply that, if selling the item to any of the bidders using mechanism $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ earns at least η-fraction of the expected revenue that could have been achieved with the optimal monopoly price, then $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ earns at least $\frac{\eta}{2}$-fraction of the optimal mechanism in the case of two bidders. By using the formulas in Table 2.1, we can verify that $\gamma\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}_1\right) = 0.92$ and $\gamma\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}_2\right) = 0.84$, and we therefore obtain $\gamma\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}\right) \geq \frac{\min\{0.92, 0.84\}}{2} = 0.42$ for the two-bidders case. This example could be extended beyond two bidder by considering the minimum of their the single bidder case competitive ratios.

So far, we have characterized the revenue guarantees associated with $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ when the seller has the exact knowledge of the upper bound $\bar{v}$. However, it may be difficult to know the exact value of this upper bound in reality. In the next section, we study the case where the seller only knows a noisy estimator of the upper bound.

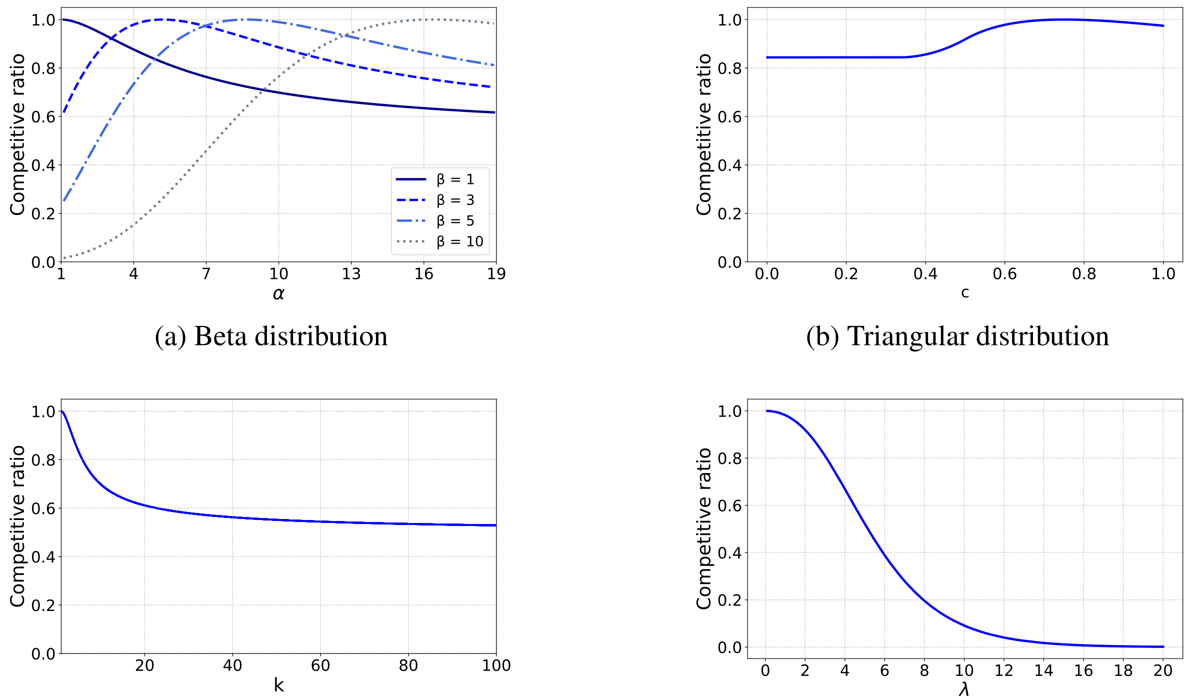

(a) Beta distribution

(b) Triangular distribution

(c) Power distribution

(d) Truncated exponential distribution

Figure 2.2: Competitive ratios for parametric distributions

2.3.4 Uncertainty in Upper Bound

In practice, the seller may not be able to precisely estimate the upper bound $\overline{v}$ and can only rely on a noisy estimator $\hat{v}$. Assume that $\overline{v}$ lies within a range of $[\hat{v}-\varepsilon, \hat{v}+\varepsilon]$, where $\varepsilon \geq 0$ reflects the level of uncertainty. When $\varepsilon = 0$, $\overline{v}$ becomes equal to $\hat{v}$, and the seller possesses a perfect knowledge of the upper bound. As ε increases, the seller's uncertainty about the estimate $\hat{v}$ also grows.

The following proposition demonstrates that, for any $\varepsilon \in [0, \hat{v}]$ and $\overline{v} \in [\hat{v}-\varepsilon, \hat{v}+\varepsilon]$, by implementing a second price auction with the deterministic reserve price $\frac{\hat{v}-\varepsilon}{2}$, the seller can secure a revenue of at least $\frac{\hat{v}-\varepsilon}{4(\hat{v}+\varepsilon)}$ times the optimal expected revenue across $\mathscr{G}(\mathscr{V}^N)$.

Proposition 2.19. *For any $\varepsilon \in [0, \hat{v}]$, $\overline{v} \in [\hat{v}-\varepsilon, \hat{v}+\varepsilon]$, and $\mathbb{F} \in \mathscr{G}(\mathscr{V}^N)$, the second price auction with reserve price $\frac{\hat{v}-\varepsilon}{2}$ is at least a $\frac{4(\hat{v}+\varepsilon)}{\hat{v}-\varepsilon}$-approximation under $\mathbb{F}$, i.e., $\gamma\left(M^{\mathrm{sp}}\left(\frac{\hat{v}-\varepsilon}{2}\right), \mathbb{F}\right) \geq \frac{\hat{v}-\varepsilon}{4(\hat{v}+\varepsilon)}$. If the bidders' values are i.i.d. under $\mathbb{F}$, then $M^{\mathrm{sp}}\left(\frac{\hat{v}-\varepsilon}{2}\right)$ is at least a $\frac{2(\hat{v}+\varepsilon)}{\hat{v}-\varepsilon}$-approximation, i.e., $\gamma\left(M^{\mathrm{sp}}\left(\frac{\hat{v}-\varepsilon}{2}\right), \mathbb{F}\right) \geq \frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)}$.*

Note that when $\varepsilon = 0$, Proposition 2.19 reduces to Theorem 2.8. As the value of ε increases, the approximation constant $\frac{4(\hat{v}+\varepsilon)}{\hat{v}-\varepsilon}$ also increases, which results in a weaker revenue approximation guarantee.

Thus far, we have demonstrated the analytical revenue guarantees offered by our proposed mechanism. In the next section, we showcase its performance through a series of numerical experiments.

2.4 Numerical Experiments

We now design numerical experiments to demonstrate the performance of the proposed mechanism in environments that go beyond the scope of the subclass of regular distributions, $\mathscr{G}$, as defined in (2.1), and to further compare this performance to multiple benchmark mechanisms. To this end, we generate random distributions from which the bidders' values, $\tilde{\boldsymbol{v}}$, are drawn. We then carry out numerical experiments to assess the performance of the proposed mechanism and compare this performance to three other well-established benchmarks. This section proceeds as follows: In Subsection 2.4.1, we describe the construction of the random distributions and elaborate on the choice of benchmark mechanisms. In Subsection 2.4.2, we present the results of our numerical experiments which are divided into two parts. In the first, we assume that the bidders' values are independent, whereas in the second, we consider the case of non-independent bidders' values.

2.4.1 Experimental Setup

Below we outline the generation of random distributions followed by the description of the benchmark mechanisms.

2.4.1.1 Generation of Random Distributions:

In our process of generating random distributions $\mathbb{F}$, we first generate marginal distributions $\mathbb{F}_i$ for every $i \in \mathscr{N}$. Subsequently, we combine these distributions via independent coupling to create the joint distribution represented by $\mathbb{F} = \times_{i \in \mathscr{N}} \mathbb{F}_i$. Generating marginal random distributions, however, is not a trivial task. Our generation process is inspired by that employed by Cohen et al. (2021). In their approach, which is in the context of inverse demand functions, they seek to associate quantities with corresponding price levels. In other words, they generate random prices $P(\cdot)$ as a function of demand quantities that are assumed to fall within $[0, Q_{max}]$, satisfying $P(0) = P_{max}$ and $P(Q_{max}) = 0$. To do so, the range of demanded quantities, $[0, Q_{max}]$, is split into a fixed number S of segments of identical size. They then generate $S-1$ random values y_s, $s = 1, \ldots, S-1$, independently drawn from a uniform distribution over $[0,1]$. They set the prices $P_0 = P_{max}$, $P_1 = y_1 P_{max}$, $P_s = y_s P_{s-1}$, for $s = 2, ..., S-1$, and $P_S = 0$, where P_s denotes the price at the right boundary of segment s, $s = 1, \ldots, S$. They construct the price function $P(\cdot)$ as a piece-wise linear function consistent with the computed prices P_s. Naturally, this approach is simple and it ensures a consistent number of price and demand quantity points for all bidders. One limitation of this approach is that since it fixes the segment size, it may miss out on a large range of alternative random distributions.

Our proposed approach addresses this issue through a modification. We adapt their procedure to our setting and modify it in two ways. First, instead of using equal intervals for the range of values, we associate each interval with a random variable, denoted as x_s, $s = 1,2\dots S$, and normalize it to determine the interval length as $\frac{x_s}{\sum_j x_j}\bar{v}$ and the value point $v_s \in [0,\bar{v}]$ (which is the right boundary of interval s) is then given by $v_s = v_{s-1} + \frac{x_s}{\sum_j x_j}\,\bar{v}$, for $s = 1,2\dots S$, with $v_0 = 0$. Second, we link each value point with another i.i.d. random variable, labeled y_s, drawn from uniform distribution over $[0,1]$ and normalize it to calculate the corresponding marginal probability mass $f_i(v_s) = \frac{y_s}{\sum_j y_j}$ at value point v_s, $s = 1,\dots,S$. We denote by $\mathbb{F}_i(v_s) = \sum_{t\le v_s} f_i(t)$, $s = 1,2\dots S$, with $\mathbb{F}_i(0) = 0$, the marginal cumulative distribution of the value $\tilde{v}_i$. We repeat the same procedure for each bidder $i \in \mathcal{N}$. The resulting joint distribution for the values $\tilde{\boldsymbol{v}}$ of all bidders is given by $\mathbb{F} = \times_{i\in\mathcal{N}}\mathbb{F}_i$.

In summary, our method advances the generation of marginal random distributions, enhancing versatility across scenarios. Improving upon Cohen et al. (2021), our approach introduces dynamically determined interval lengths for flexible value ranges. This approach introduces further variability, potentially leading to more realistic value distributions through random value points generation for each bidder, leading to varying positions of value points within $[0,\bar{v}]^N$, affecting the distribution's shape.

We illustrate our method with an example. Let $\bar{v} = 1$ and $S = 5$. Assume the first set of randomly drawn numbers are $x_1 = 0.234$, $x_2 = 0.562$, $x_3 = 0.754$, $x_4 = 0.123$, and $x_5 = 0.879$, the sum of which is $\sum_j x_j = 2.552$. Calculating the value points, we have $v_0 = 0$, $v_1 = \frac{0.234}{2.552}\bar{v} = 0.092$, $v_2 = 0.092 + \frac{0.562}{2.552} = 0.312$, and similarly $v_3 = 0.607$, $v_4 = 0.656$, and $v_5 = 1$. Now, randomly drawing the second set of random numbers, we have $y_1 = 0.287$, $y_2 = 0.482$, $y_3 = 0.156$, $y_4 = 0.893$, and $y_5 = 0.648$, with a sum of $\sum_j y_j = 2.466$. Accordingly, we determine the marginal probabilities for the value points as: $f_i(0.092) = \frac{0.287}{2.466} = 0.116$, $f_i(0.312) = 0.196$, $f_i(0.607) = 0.063$, $f_i(0.656) = 0.362$, and $f_i(1) = 0.263$. We repeat this procedure for each bidder. Figure 2.3 illustrates the randomly chosen value points and associated probability masses for the bidder from the above example (Bidder 1) along with an another bidder (Bidder 2).

2.4.1.2 Benchmark Mechanisms:

We consider three benchmarks:

(*i*) The optimal mechanism, M^{opt}, which maximizes expected revenue with a complete knowledge of the distribution $\mathbb{F}$. This mechanism serves as a criterion for calculating the competitive ratio, a performance metric defined in Definition 2.6, quantifying the relative performance of other mechanisms, such as mechanisms (*ii*) and (*iii*), which are outlined in the next paragraphs, as well as our proposed mechanism in comparison to the optimal mechanism. The competitive ratio enables

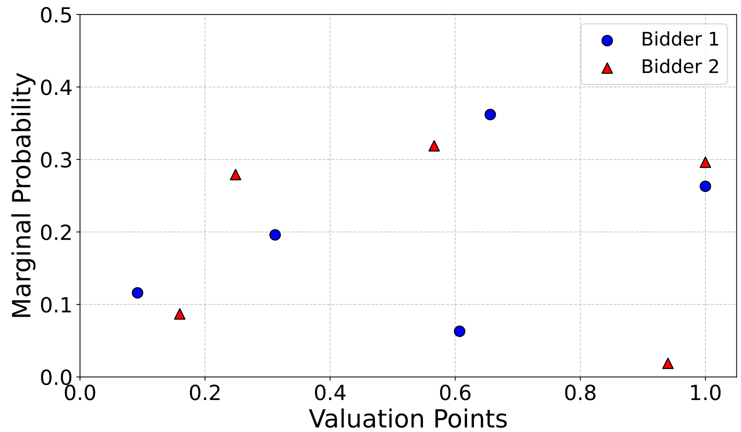

Figure 2.3: Illustration of randomly generated value points and corresponding probability masses

us to assess the proportion of the maximum achievable revenue that each mechanism can attain.

(*ii*) A second price auction equipped with lazy bidder-specific random reserve prices $\tilde{r}_i \sim \mathbb{F}_i$, where $i \in \mathcal{N}$. This mechanism, denoted as $M^{\mathrm{sp}}\,(\tilde{\boldsymbol{r}} \sim \mathbb{F})$, conceptually builds upon ideas presented in Dhangwatnotai et al. (2015), which assumes that $\mathbb{F}$ is regular (note that the randomly generated distributions are not regular). Under this mechanism, if the bidders' values are i.i.d., drawing a single sample from this i.i.d. distribution to serve as the common reserve price, half of the optimal revenue shall be obtained in expectation. If bidders are not identical, the mechanism draws an individual reserve price for each of the bidders from their respective marginal distributions.

(*iii*) A second price auction without a reserve price, denoted as M^{sp}. This mechanism is of particular interest thanks to its relation to the Bulow-Klemperer theorem, which provides a revenue guarantee for the underlying mechanism. Specifically, the expected revenue of M^{sp} with $n+1$ bidders, each with value i.i.d. drawn from $\mathbb{F}_1$, is at least that of a revenue-maximizing auction with n such bidders. The significance of the Bulow-Klemperer theorem lies in the ability of M^{sp} to outperform the optimal auction revenue with the addition of just one more bidder. M^{sp} offers a straightforward way to maximize revenue in the i.i.d. setting.

2.4.2 Results

The experiments are divided into two parts. In the first part, we assume that the bidders' values are independent (Section 2.4.2.1). In the second part, we examine the case of non-independent bidders' values (Section 2.4.2.2).

2.4.2.1 Independent Bidders' Values:

To numerically evaluate the performance of our proposed mechanism, we generated various distributions for bidder values. Specifically, we create instances of joint probability distributions for multiple bidders with each bidder's potential values ranging within the interval [0,1]. In total, we generate 1000 instances.

Figure 2.4 illustrates the competitive ratio distribution for $M^{\text{sp}}(\frac{\bar{v}}{2})$ in comparison to M^{sp}, and $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ for the case of two bidders (i.e., $N = 2$). Specifically, for each of the mechanisms $M^{\text{sp}}(\frac{\bar{v}}{2})$, M^{sp}, and $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$, we evaluate the performance based on competitive ratio, computed for 1000 random distribution instances, and Figure 2.4 shows histograms of these ratios. Panel (a) shows the histograms of the competitive ratio of $M^{\text{sp}}(\frac{\bar{v}}{2})$ compared to that of $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$, while the histograms in Panel (b) shows the competitive ratio of $M^{\text{sp}}(\frac{\bar{v}}{2})$ compared to that of M^{sp}. The Y-axis represents frequency, indicating how many times a particular competitive ratio value occurred out of 1000 instances, while the X-axis represents the range of competitive ratio values.

From these figures, we see that mechanism $M^{\text{sp}}(\frac{\bar{v}}{2})$ has a competitive ratio that is always greater than 0.5 and can reach values close to 0.9. However, the competitive ratio for $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ can be as low as 0.44 and it never exceeds 0.82. Mechanism M^{sp} is able to offer better performance than $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ in certain instances, but it comes at the cost of broader variance and a worst-case competitive ratio as low as 0.36.

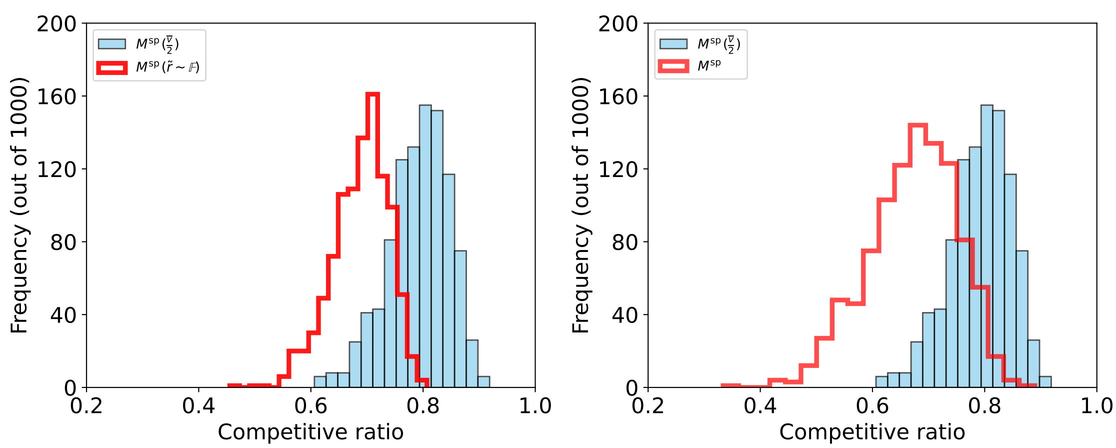

(a) $M^{\text{sp}}(\frac{\bar{v}}{2})$ vs. $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$

(b) $M^{\text{sp}}(\frac{\bar{v}}{2})$ vs. M^{sp}

Figure 2.4: Competitive ratio comparison for $M^{\text{sp}}(\frac{\bar{v}}{2})$, $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ and M^{sp}

To gain a better understanding of the relative performance of the mechanisms in each distribution scenario, for each of the 1000 instances, we calculate the revenue ratio for the mechanisms mentioned above. Figure 2.5(a) displays the distribution of the revenue ratio between $M^{\text{sp}}(\frac{\bar{v}}{2})$ and $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$, while Figure 2.5(b) shows the distribution of the revenue ratio between $M^{\text{sp}}(\frac{\bar{v}}{2})$ and M^{sp}. From this figure, we see that $M^{\text{sp}}(\frac{\underline{v}}{2})$ outperforms the two benchmark mechanisms across the vast majority of the instances. Notably, in 957 out of 1000 instances, $M^{\text{sp}}(\frac{\underline{v}}{2})$ outperforms $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$, and in 947 out of 1000 instances, $M^{\text{sp}}(\frac{\underline{v}}{2})$ outperforms M^{sp}.

In addition to the above demonstration, we also calculate two additional performance measures. The first measure is the average of the revenue ratios, across all 1000 instances. The average revenue ratio of $M^{\text{sp}}(\frac{\underline{v}}{2})$ to $M^{\text{sp}}(\tilde{r} \sim \mathbb{F})$ is $\frac{1}{1000}\sum_{i=1}^{1000} \frac{\text{Rev}(M^{\text{sp}}(\frac{\underline{v}}{2}),\mathbb{F}_i)}{\text{Rev}(M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_i),\mathbb{F}_i)} = 1.15$, and to M^{sp} is $\frac{1}{1000}\sum_{i=1}^{1000} \frac{\text{Rev}(M^{\text{sp}}(\frac{\underline{v}}{2}),\mathbb{F}_i)}{\text{Rev}(M^{\text{sp}},\mathbb{F}_i)} = 1.19$. These two values represent the means of the distributions depicted in Figure 2.5. The second is the ratio of the average revenue obtained by the $M^{\text{sp}}(\frac{\underline{v}}{2})$ to that obtained by either of the benchmark mechanisms. Specifically, compared with $M^{\text{sp}}(\tilde{r} \sim \mathbb{F})$, we have $\frac{\frac{1}{1000}\sum_{i=1}^{1000}\text{Rev}(M^{\text{sp}}(\frac{\underline{v}}{2}),\mathbb{F}_i)}{\frac{1}{1000}\sum_{i=1}^{1000}\text{Rev}(M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_i),\mathbb{F}_i)} = 1.08$, and compared with M^{sp}, we have $\frac{\frac{1}{1000}\sum_{i=1}^{1000}\text{Rev}(M^{\text{sp}}(\frac{\underline{v}}{2}),\mathbb{F}_i)}{\frac{1}{1000}\sum_{i=1}^{1000}\text{Rev}(M^{\text{sp}},\mathbb{F}_i)} = 1.18$. The average revenues under $M^{\text{sp}}\left(\frac{\underline{v}}{2}\right)$, M^{sp}, and $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ are 0.50, 0.43, and 0.46, respectively, (see the instance of the two bidder case in Figure 2.6, which will be discussed in more detail in the next paragraph) producing these ratios. Considering either of the performance measures, $M^{\text{sp}}(\frac{\underline{v}}{2})$ consistently achieves higher revenue compared to both $M^{\text{sp}}(\tilde{r} \sim \mathbb{F})$ and M^{sp}.

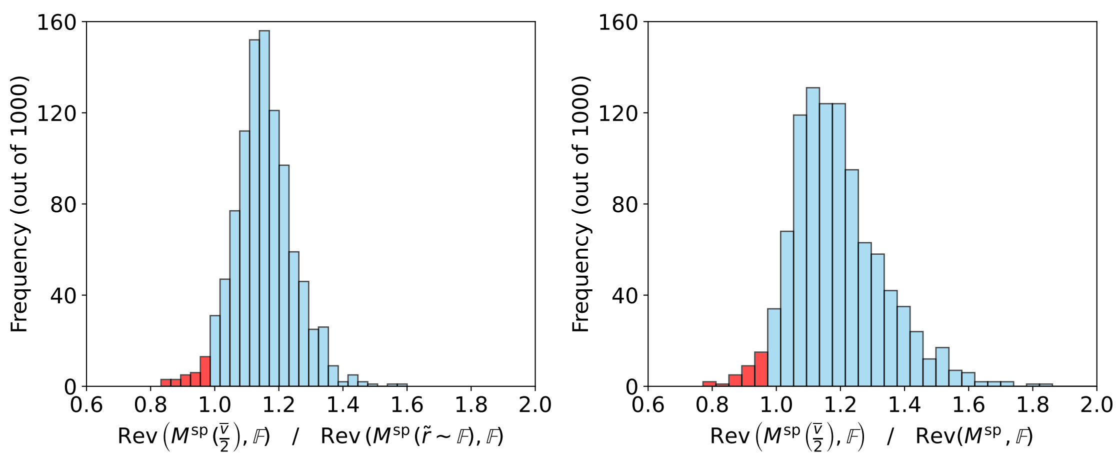

(a) $M^{\text{sp}}(\frac{\underline{v}}{2})$ vs. $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$

(b) $M^{\text{sp}}(\frac{\underline{v}}{2})$ vs. M^{sp}

Figure 2.5: Instance-specific revenue ratio comparison for $M^{\text{sp}}(\frac{\underline{v}}{2})$, $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ and M^{sp}

As we increase the number of bidders from 2 to 3 or more, computing the optimal revenue becomes computationally challenging. While we are able to compute and compare the competitive ratio for two bidders, for cases with 3 to 10 bidders, we focus solely on comparing the average revenue without calculating the competitive ratio. In this experiment, $M^{\text{sp}}(\frac{\underline{v}}{2})$ consistently outperforms both $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ and M^{sp} on average. Figure 2.6 demonstrates the average revenue obtained by the three discussed mechanisms across 1000 iterations for different number of bidders. Note that we generate 1000 instances for each case of the number of bidders separately. The X-axis represents the number of bidders, while the Y-axis denotes the average revenue obtained by each mechanism over the course of 1000 instances. Mechanism $M^{\text{sp}}(\frac{\bar{v}}{2})$ generates at least as much average revenue as $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ and M^{sp} across all considered number of bidders.

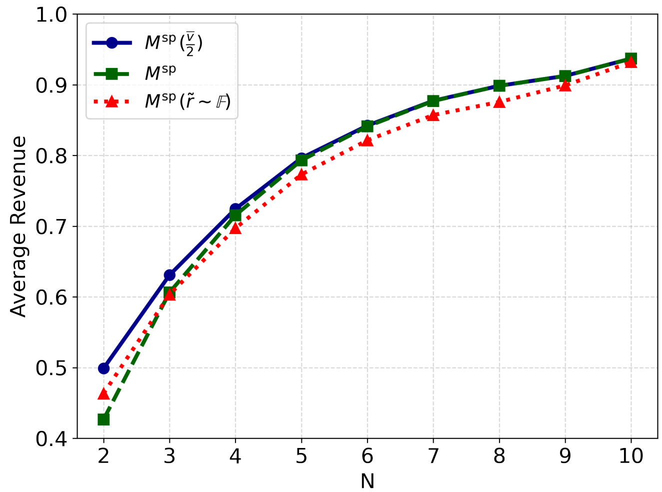

Figure 2.6: Revenue comparison for $N \in \{2, \ldots, 10\}$

2.4.2.2 Non-independent Bidders' Values:

In this section, we focus on the case of two bidders, i.e., $N = 2$, studying a bivariate relationship between the bidders' values to simplify the analysis. Working with two bidders enables us to explore a wide range of correlation levels and parameters in a computationally efficient manner and may potentially serve as a prerequisite for later expanding the analysis to include additional bidders. We model the bidders' values through affine functions that depend on two independent signals $\tilde{s}_1$ and $\tilde{s}_2$ and a correlation parameter $\rho \in [0, 1]$. These variables are generated randomly similarly to before.

Bidder 1's value equals to the signal $\tilde{s}_1$ whereas bidder 2's value involves $\tilde{s}_1$ and $\tilde{s}_2$ and ρ. For positive correlations ($\rho > 0$), $\tilde{v}_2$ is influenced positively by $\tilde{s}_1$ reflected by the term involving max function in equation (2.4). Conversely, for negative correlations ($\rho < 0$), $\tilde{v}_2$ is influenced negatively by $\tilde{s}_1$ reflected by the term involving $\bar{v} - \tilde{s}_1$ in (2.4). This setup allows us to observe distinct behaviors at extreme correlation values, shaping bidder 2's value based on the interplay between $\tilde{s}_1$, $\tilde{s}_2$, and ρ. Specifically, the value $\tilde{v}_1$ of bidder 1 is given by

$$\tilde{v}_1 = \tilde{s}_1$$

and the value $\tilde{v}_2$ of bidder 2 is given by

$$\tilde{v}_2 = \tilde{s}_2\,(1 - |\rho|) + \max\{0, \rho\}\tilde{s}_1 - \min\{0, \rho\}\,(\bar{v} - \tilde{s}_1)\,. \tag{2.4}$$

The equalities above are to be understood to hold almost surely. This arrangement implies that the sum of three terms on the right-hand side of (2.4) can never exceed $\bar{v}$.

We break down each individual term in (2.4) to understand its contribution and meaning:

1. $\tilde{s}_2(1-|\rho|)$: The magnitude of the coefficient $(1-|\rho|)$ of $\tilde{s}_2$ increases as the absolute value of ρ decreases, in other words, as it approaches to zero. As $\tilde{s}_2$ is independent from $\tilde{s}_1 = \tilde{v}_1$, this means that as $|\rho|$ decreases to zero, $\tilde{v}_2$ correlates less with $\tilde{v}_1$.

2. $\max\{0,\rho\}\tilde{s}_1$: This term adds a value proportional to $\tilde{s}_1$ if ρ is positive, creating a positive correlation with $\tilde{s}_1 = \tilde{v}_1$.

3. $\min\{0,\rho\}(\bar{v}-\tilde{s}_1)$: This term adds a value that is proportional to $\bar{v}-\tilde{s}_1$ if ρ is negative. This creates a negative correlation because the term added is larger if $\tilde{s}_1 = \tilde{v}_1$ is small, and it is smaller if $\tilde{s}_1 = \tilde{v}_1$ is large.

To enhance the understanding further, we also consider the following extreme cases for correlation parameter ρ. If $\rho = 0$, then $\tilde{v}_2 = \tilde{s}_2$, and bidder 2's value is independent from $\tilde{s}_1 = \tilde{v}_1$. If $\rho = 1$, then $\tilde{v}_2 = \tilde{s}_1 = \tilde{v}_1$, indicating perfect correlation. If $\rho = -1$, $\tilde{v}_2 = \bar{v}-\tilde{s}_1 = \bar{v}-\tilde{v}_1$. In this case, bidder 2's value grows as bidder 1's value decreases. In summary, this way of modeling the bidders' values allows us to construct various scenarios for different correlation structures.

Figure 2.7 presents a comparison of competitive ratios across three mechanisms at varying correlation levels. The y-axis illustrates the average competitive ratio across 1000 instances, while the x-axis represents the correlation parameter ρ. We observe that the competitive ratio of the $M^{\text{sp}}(\frac{\bar{v}}{2})$ remains relatively stable despite changing correlation parameter, whereas this is not the case for the other two mechanisms. The $M^{\text{sp}}(\frac{\bar{v}}{2})$ outperforms both alternatives when ρ is negative or positive but sufficiently small. In case of negative ρ, the scenarios where one bidder's value is high whereas the other's value is low become more likely. In this case, the deterministic reserve price of our mechanism ensures that we charge a reasonable amount as long as one of the bidders has a value that exceeds this reserve price, which becomes more likely to happen as ρ decreases. On the other hand, M^{sp} charges the second highest value, which is likely to be much lower than the highest one, and $M^{\text{sp}}(\tilde{\boldsymbol{r}} \sim \mathbb{F})$ can end up sampling a small reserve price and charge the second highest value in a similar fashion. When ρ is positive, the bidders' values are aligned in the sense that they are likely to be closer. In this case, our mechanism ensures that we charge at least the deterministic reserve price in scenarios where this alignment is not strong and one bidder's value falls above the reserve price, whereas the other's value falls below. For scenarios where bidders' values are close and larger than the reserve price, our mechanism is also able to capitalize on the second highest value. As ρ increases to 1, however, our mechanism is outperformed by M^{sp} because M^{sp} is able to extract almost full surplus in this case because the bidders values are very likely to be the almost same.

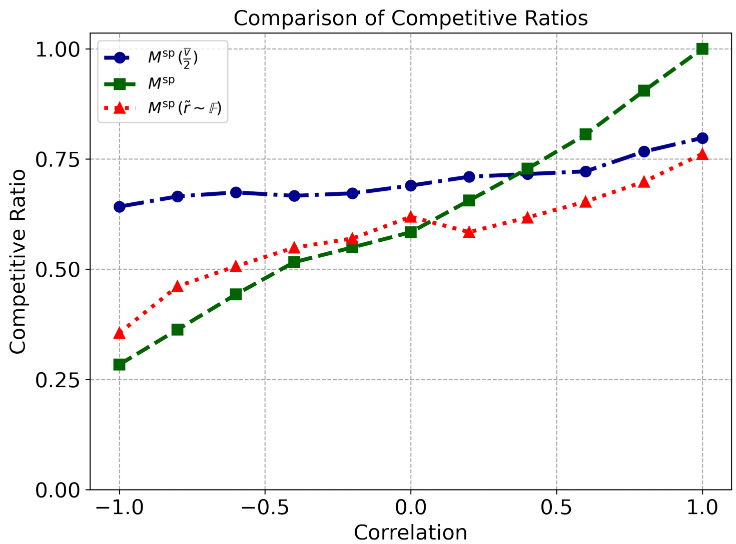

Figure 2.7: Competitive ratio comparison for 2 non-independent bidders

2.5 Conclusion

We propose a second price auction with a reserve price set at half of the upper bound on potential bidders' values, which is the only information available to the seller. Our mechanism is easy to understand, deterministic, requires minimal information, and is, therefore, appealing from a practical standpoint. On the theoretical side, we show that our proposed mechanism achieves a 4-approximation, which can be improved to a 2-approximation in the case of i.i.d. bidders' values, across an extensive class of distributions that include several practically relevant distributions with varying structural properties. In our numerical analysis, we demonstrate that our mechanism outperforms two benchmarks from the literature for the vast majority of randomly generated distributions of bidders' values.

Our approach to studying the performance of our mechanism from a theoretical standpoint differs from the common approach used in the literature, which considers distribution classes such as all distributions consistent with the available information or their restrictions to regular distributions. Recognizing that no deterministic mechanism can offer a constant-factor approximation guarantee under all such distributions with only upper bound information, we choose to study a specific subclass. Our discussion indicates that this subclass is extensive and eliminates unrealistic and uninteresting scenarios, such as distribution mass concentrated at lower values with low variance. Our work thus raises questions about the potential conservatism of considering all distributions or regular distributions, which may include some distributions that are less practically relevant.

In this chapter, we focus on a single-item auction setting and assume that the bidders' values are independent. An interesting future avenue is to study the proposed second price auction and its robustness in the case of non-independent bidders' values from a theoretical perspective. Our numerical study involving two non-independent bidders suggests that the proposed mechanism is likely to demonstrate robustness against different correlation structures. Another future direction is to consider a multi-item setting and study the following question: How good are deterministic mechanisms for selling two or more goods and how much of the optimal revenue is guaranteed when using them? One can potentially leverage separation results, i.e., the revenue guarantee of a mechanism that sells each item separately, (e.g., Hart and Nisan (2017), Carroll (2017), Koçyiğit et al. (2024)) from the literature to generalize the results to this setting.

2.6 Appendix

Proof of Proposition 2.7. We consider $N = 1$, in which case the class of deterministic mechanisms in $\mathscr{M}$ is equal to the class of deterministic posted-price mechanisms; see, e.g., Koçyiğit et al. (2021, Section 4). We will show that for any fixed deterministic posted price $p \in \mathbb{R}_+$, we have that

$$\inf_{\mathbb{F}\in\mathscr{R}(\mathscr{V})} \gamma(M^{\mathrm{sp}}(p),\mathbb{F}) = \inf_{\mathbb{F}\in\mathscr{R}(\mathscr{V})} \frac{p(1-\mathbb{F}(p))}{\sup_r\{r(1-\mathbb{F}(r))\}} = 0, \tag{2.5}$$

where the first equality holds because the expected revenue attained from the posted price p is $p(1-\mathbb{F}(p))$, and from Myerson (1981) we know that the optimal price that maximizes the expected revenue is deterministic under any fixed $\mathbb{F} \in \mathscr{R}(\mathscr{V})$.

Since the worst-case competitive ratio $\inf_{\mathbb{F}\in\mathscr{R}(\mathscr{V})} \gamma(M^{\mathrm{sp}}(p),\mathbb{F})$ is lower bounded by 0, it is sufficient to show that there exists a subclass $\mathscr{B}(\mathscr{V}) \subset \mathscr{R}(\mathscr{V})$ of regular distributions for which the worst-case competitive ratio of p over $\mathscr{B}(\mathscr{V})$ equals zero. This will allow us to conclude that $\inf_{\mathbb{F}\in\mathscr{R}(\mathscr{V})} \gamma(M^{\mathrm{sp}}(p),\mathbb{F}) = 0$ because

$$0 \le \inf_{\mathbb{F}\in\mathscr{R}(\mathscr{V})} \frac{p(1-\mathbb{F}(p))}{\sup_r\{r(1-\mathbb{F}(r))\}} \le \inf_{\mathbb{F}\in\mathscr{B}(\mathscr{V})} \frac{p(1-\mathbb{F}(p))}{\sup_r\{r(1-\mathbb{F}(r))\}} = 0.$$

We define $\mathscr{B}(\mathscr{V})$ as the class of all beta distributions with $\alpha \ge 1$ and $\beta \ge 1$, scaled to support $\mathscr{V}$, which belong to $\mathscr{R}(\mathscr{V})$ (refer to Remark 2.18). For some fixed $\alpha \ge 1$, $\beta \ge 1$ and $\mu = \frac{\alpha}{\alpha+\beta}\bar{v} \in (0,\bar{v})$, we construct a sequence $(\mathbb{F}^B_{\alpha,\beta,k})_{k\in\mathbb{Z}_+}$ of distributions where, for any $k \in \mathbb{N}$, $\mathbb{F}^B_{\alpha,\beta,k} \in \mathscr{B}(\mathscr{V})$ has a pdf defined by

$$f_{\alpha,\beta,k}(x) = \frac{1}{B(k\alpha,k\beta)} \left(\frac{x}{\bar{v}}\right)^{k\alpha-1} \left(1-\frac{x}{\bar{v}}\right)^{k\beta-1},$$

where $B(k\alpha,k\beta)$ is the beta function. Note that the expected value μ under any $\mathbb{F}^B_{\alpha,\beta,k}$ is constant and equal to $\mu = \frac{k\alpha}{k\alpha+k\beta}\bar{v} = \frac{\alpha}{\alpha+\beta}\bar{v}$ irrespective of the value of k. Note also that for any $\mu \in (0,\bar{v})$, we can find some $\alpha \ge 1$ and $\beta \ge 1$ such that $\mu = \frac{\alpha}{\alpha+\beta}\bar{v}$.

In contrast, the variance $\sigma_k^2 = \frac{k^2\alpha\beta}{(k(\alpha+\beta))^2(k(\alpha+\beta)+1)}\bar{v}^2$ under $\mathbb{F}^B_{\alpha,\beta,k}$ decreases as k increases and converges to zero. In particular, we have

$$\lim_{k\to\infty} \sigma_k^2 = \lim_{k\to\infty} \frac{\alpha\beta}{(\alpha+\beta)^2(k(\alpha+\beta)+1)}\bar{v}^2 = 0.$$

This means that as k grows, $\mathbb{F}^B_{\alpha,\beta,k}$ concentrates the probability mass around μ, eventually converging to a Dirac point mass, i.e., $\lim_{k\to\infty}\mathbb{F}^B_{\alpha,\beta,k}(x) = 0$ if $x < \mu$ and $\lim_{k\to\infty}\mathbb{F}^B_{\alpha,\beta,k}(x) = 1$ if $x \ge \mu$.

We now return to showing that $\inf_{\mathbb{F}^B \in \mathscr{B}(\mathscr{V})} \frac{p(1-\mathbb{F}^B(p))}{\sup\limits_r \{r(1-\mathbb{F}^B(r))\}} = 0$ for any fixed $p \in \mathbb{R}_+$. If $p = 0$, then we have $p\left(1-\mathbb{F}^B(p)\right) = 0$ for any $\mathbb{F}^B \in \mathscr{B}(\mathscr{V})$, whereas $\sup\limits_r \{r\left(1-\mathbb{F}^B(r)\right)\} > 0$ because $\mathbb{F}^B \in \mathscr{R}(\mathscr{V})$ and is continuous. The claim thus follows for $p = 0$. Consider now any $p \in (0, \overline{v}]$ (Note that $p > \overline{v}$ generates an expected revenue of zero). Since for any $\alpha, \beta \geq 1$ and $k \in \mathbb{Z}_+$ we have $\mathbb{F}^B_{\alpha,\beta,k} \in \mathscr{B}(\mathscr{V})$, we obtain

$$\begin{aligned}
\inf_{\mathbb{F}^B \in \mathscr{B}(\mathscr{V})} \frac{p(1-\mathbb{F}^B(p))}{\sup\limits_r \{r(1-\mathbb{F}^B(r))\}} &\leq \inf_{\alpha,\beta \geq 1} \inf_{k \in \mathbb{Z}_+} \frac{p\left(1-\mathbb{F}^B_{\alpha,\beta,k}(p)\right)}{\sup\limits_r \left\{r\left(1-\mathbb{F}^B_{\alpha,\beta,k}(r)\right)\right\}} \\
&\leq \inf_{\alpha,\beta \geq 1} \lim_{k \to \infty} \frac{p\left(1-\mathbb{F}^B_{\alpha,\beta,k}(p)\right)}{\sup\limits_r \left\{r\left(1-\mathbb{F}^B_{\alpha,\beta,k}(r)\right)\right\}}.
\end{aligned}$$

We next show that the right-hand side of the above inequality, $\inf\limits_{\alpha,\beta \geq 1} \lim\limits_{k \to \infty} \frac{p\left(1-\mathbb{F}^B_{\mu,k}(p)\right)}{\sup\limits_r \left\{r\left(1-\mathbb{F}^B_{\mu,k}(r)\right)\right\}}$ equals zero, which immediately implies that the left-hand side, $\inf_{\mathbb{F}^B \in \mathscr{B}(\mathscr{V})} \frac{p(1-\mathbb{F}^B(p))}{\sup\limits_r \{r(1-\mathbb{F}^B(r))\}}$, is also equal to zero and proves our claim. To this end, fix any $\alpha', \beta' \geq 1$ such that $\mu' = \frac{\alpha'}{\alpha'+\beta'}\overline{v} \in (0, p)$ and $\delta \in (0, \mu')$. Hence, we have $(\mu' - \delta) > 0$ and that $(\mu' - \delta)(1 - \mathbb{F}^B_{\alpha',\beta',k}(\mu' - \delta)) \leq \sup\limits_r \, r(1 - \mathbb{F}^B_{\alpha',\beta',k}(r))$. We thus obtain

$$\begin{aligned}
\inf_{\alpha,\beta \geq 1} \lim_{k \to \infty} \frac{p\left(1-\mathbb{F}^B_{\alpha,\beta,k}(p)\right)}{\sup\limits_r \left\{r\left(1-\mathbb{F}^B_{\alpha,\beta,k}(r)\right)\right\}} &\leq \lim_{k \to \infty} \frac{p(1-\mathbb{F}^B_{\alpha',\beta',k}(p))}{\sup\limits_r \{r(1-\mathbb{F}^B_{\alpha',\beta',k}(r))\}} \\
&\leq \lim_{k \to \infty} \frac{p(1-\mathbb{F}^B_{\alpha',\beta',k}(p))}{(\mu'-\delta)(1-\mathbb{F}^B_{\alpha',\beta',k}(\mu'-\delta))}.
\end{aligned}$$

We now apply the limit operation separately to the numerator and denominator, which is possible because the two limits exists and the limit of the denominator is nonzero. We have

$$\begin{aligned}
\lim_{k \to \infty} \frac{p(1-\mathbb{F}^B_{\alpha',\beta',k}(p))}{(\mu'-\delta)(1-\mathbb{F}^B_{\alpha',\beta',k}(\mu'-\delta))} &= \frac{\lim\limits_{k \to \infty} p(1-\mathbb{F}^B_{\alpha',\beta',k}(p))}{\lim\limits_{k \to \infty} (\mu'-\delta)(1-\mathbb{F}^B_{\alpha',\beta',k}(\mu'-\delta))} \\
&= \frac{p(1-\lim\limits_{k \to \infty} \mathbb{F}^B_{\alpha',\beta',k}(p))}{(\mu'-\delta)(1-\lim\limits_{k \to \infty} \mathbb{F}^B_{\alpha',\beta',k}(\mu'-\delta))} = 0,
\end{aligned}$$

where the last equality follows from our previous argument, i.e., $\lim_{k \to \infty} \mathbb{F}^B_{\alpha',\beta',k}(x) = 0$ if $x < \mu'$

and $\lim_{k\to\infty}\mathbb{F}^B_{\alpha',\beta',k}(x)=1$ if $x\geq\mu'$. The claim thus follows. $\square$

Proof of Theorem 2.8. We first consider the single-bidder case with $\tilde{v}_1\sim\mathbb{F}_1$ and show that the expected revenue of our proposed mechanism $M^{\text{sp}}(\frac{\bar{v}}{2})$ is at least as high as that of $M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_1)$ when $\mathbb{F}_1\in\mathscr{G}(\mathscr{V})$. To this end, we will show that $\text{Rev}\left(M^{\text{sp}}\left(\tilde{r}\sim\mathbb{F}_1\right),\mathbb{F}_1\right)\leq\text{Rev}\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right),\mathbb{F}_1\right)$ is equivalent to the condition that

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)+\frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]\leq 1+\frac{1}{\bar{v}}\int_0^{\bar{v}}\mathbb{F}_1(x)(1-\mathbb{F}_1(x))\mathrm{d}x, \tag{2.6}$$

which is the inequality condition in the definition of $\mathscr{G}(\mathscr{V})$. In other words, any $\mathbb{F}_1\in\mathscr{G}(\mathscr{V})$ satisfies (2.6), and therefore the revenue inequality holds for every $\mathbb{F}_1\in\mathscr{G}(\mathscr{V})$.

We start by re-expressing $\text{Rev}\left(M^{\text{sp}}\left(\tilde{r}\sim\mathbb{F}_1\right),\mathbb{F}_1\right)$ as

$$\begin{aligned}
\text{Rev}\left(M^{\text{sp}}\left(\tilde{r}\sim\mathbb{F}_1\right),\mathbb{F}_1\right) &= \int_0^{\bar{v}} r(1-\mathbb{F}_1(r))f_1(r)\mathrm{d}r\\
&= \int_0^{\bar{v}}\int_0^{r}\left(1-\mathbb{F}_1(x)-xf_1(x)\right)\mathrm{d}x f_1(r)\mathrm{d}r\\
&= \int_0^{\bar{v}}\int_x^{\bar{v}}\left(1-\mathbb{F}_1(x)-xf_1(x)\right)f_1(r)\mathrm{d}r\mathrm{d}x\\
&= \int_0^{\bar{v}}\left(1-\mathbb{F}_1(x)-xf_1(x)\right)(1-\mathbb{F}_1(x))\mathrm{d}x\\
&= \int_0^{\bar{v}}(1-\mathbb{F}_1(x))^2\mathrm{d}x-\int_0^{\bar{v}}x(1-\mathbb{F}_1(x))f_1(x)\mathrm{d}x\\
&= \int_0^{\bar{v}}(1-\mathbb{F}_1(x))^2\mathrm{d}x-\text{Rev}(M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_1),\mathbb{F}_1),
\end{aligned}$$

where the first equality holds because the expected revenue of selling the item to a bidder with value $\tilde{v}_1\sim\mathbb{F}_1$ using a posted price mechanism with price r is given by $r(1-\mathbb{F}_1(r))$, and because we randomize the price $\tilde{r}\sim\mathbb{F}_1$. The second equality follows from the fundamental theorem of calculus (note that $\mathbb{F}_1$ is monotone non-decreasing and therefore differentiable almost everywhere). Remaining equalities follow from simple algebra, and the last equality is a result of the observation that $\int_0^{\bar{v}}x(1-\mathbb{F}_1(x))f_1(x)\mathrm{d}x=\text{Rev}(M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_1),\mathbb{F}_1)$, which can be seen from the first equality. By the derivation above, we have

$$\text{Rev}(M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_1),\mathbb{F}_1)=\frac{1}{2}\int_0^{\bar{v}}(1-\mathbb{F}_1(x))^2dx.$$

Noting that $\text{Rev}\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right),\mathbb{F}_1\right)=\frac{\bar{v}}{2}\left(1-\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)$, we can now write the inequality

$\text{Rev}\left(M^{\text{sp}}\left(\tilde{r}\sim\mathbb{F}_1\right),\mathbb{F}_1\right)\le\text{Rev}\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right),\mathbb{F}_1\right)$ equivalently as

$$\begin{aligned}
&\frac{1}{2}\int_0^{\bar{v}}(1-\mathbb{F}_1(x))^2\mathrm{d}x\le\frac{\bar{v}}{2}\left(1-\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)\\
\iff&\int_0^{\bar{v}}(1-\mathbb{F}_1(x))\mathrm{d}x-\int_0^{\bar{v}}\mathbb{F}_1(x)(1-\mathbb{F}_1(x))\mathrm{d}x\le\bar{v}\left(1-\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)\\
\iff&\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)+\frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]\le 1+\frac{1}{\bar{v}}\int_0^{\bar{v}}\mathbb{F}_1(x)(1-\mathbb{F}_1(x))\mathrm{d}x,
\end{aligned}$$

where the first if and only if condition follows from simple algebra, and the second if and only if condition follows from the observation that $\int_0^{\bar{v}}(1-\mathbb{F}_1(x))\mathrm{d}x=\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]$, which holds because $\tilde{v}$ is non-negative. Recalling the definition of $\mathscr{G}(\mathscr{V})$ and noting that the last inequality above coincides with the condition in the definition of $\mathscr{G}(\mathscr{V})$, we can now conclude that $\text{Rev}\left(M^{\text{sp}}\left(\tilde{r}\sim\mathbb{F}_1\right),\mathbb{F}_1\right)\le\text{Rev}\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right),\mathbb{F}_1\right)$ if $\mathbb{F}_1\in\mathscr{G}(\mathscr{V})$.

Recall that the competitive ratio of $M^{\text{sp}}(\tilde{r}\sim\mathbb{F}_1)$ is at least $\frac{1}{2}$ under any $\mathbb{F}_1\in\mathscr{R}(\mathscr{V})$ and note that $\mathscr{G}(\mathscr{V})\subseteq\mathscr{R}(\mathscr{V})$. For every $\mathbb{F}_1\in\mathscr{G}(\mathscr{V})$, we thus have

$$\text{Rev}\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right),\mathbb{F}_1\right)\ge\text{Rev}\left(M^{\text{sp}}\left(\tilde{r}\sim\mathbb{F}_1\right),\mathbb{F}_1\right)\ge\frac{1}{2}\text{Rev}^{\star}(\mathbb{F}_1),$$

which implies that the competitive ratio of $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ is at least $\frac{1}{2}$ under any $\mathbb{F}_1\in\mathscr{G}(\mathscr{V})$, i.e., $\gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}),\mathbb{F}_1\big)\ge\frac{1}{2}$.

Consider now the case of $N\ge 2$ bidders and any $\mathbb{F}=\times_{i\in\mathscr{N}}\mathbb{F}_i\in\mathscr{G}(\mathscr{V}^N)$. Note that $\mathbb{F}_i\in\mathscr{G}(\mathscr{V})$ for all $i\in\mathscr{N}$ by the definition of $\mathscr{G}$. Thus, we have $\gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}),\mathbb{F}_i\big)\ge\frac{1}{2}$ for all $i\in\mathscr{N}$ by our previous arguments in this proof. According to Theorem 2.11, $\gamma\big(\hat{M}^{\text{sp}}\left(\boldsymbol{r}^{\star}\right),\mathbb{F}\big)\ge\frac{1}{2}$ when bidders' values are independent and not necessarily identically distributed, and $\gamma\big(\hat{M}^{\text{sp}}\left(\boldsymbol{r}^{\star}\right),\mathbb{F}\big)=1$ when bidders' values are i.i.d. By Theorem 2.12, we can now conclude that under $\mathbb{F}\in\mathscr{G}(\mathscr{V}^N)\subseteq\mathscr{R}(\mathscr{V}^N)$, we have $\gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}),\mathbb{F}\big)\ge\frac{1}{2}\cdot\frac{1}{2}=\frac{1}{4}$ for independent but not necessarily identical bidders' values, meaning $M^{\text{sp}}\left(\frac{\bar{v}}{2}\right)$ guarantees at least a 4-approximation. For the case of i.i.d. bidders' values, similarly we have $\gamma\big(M^{\text{sp}}(\frac{\bar{v}}{2}),\mathbb{F}\big)\ge\frac{1}{2}\cdot 1=\frac{1}{2}$. This completes the proof. □

Proof of Proposition 2.13. For any $k\in\mathbb{Z}_+$, it follows from our subsequent Proposition 2.15 that $\mathbb{F}_{1,k}\in\mathscr{G}([0,1])$ as $\mathbb{F}_{1,k}$ is convex on $[0,1]$. We next prove that $\lim_{k\to\infty}\gamma(M^{\text{sp}}\left(\frac{1}{2}\right),\mathbb{F}_{1,k})=\frac{1}{2}$. To this end, consider any $k\in\mathbb{Z}_+$ and note that we can express the expected revenue $\text{Rev}\left(M^{\text{sp}}\left(\frac{1}{2}\right),\mathbb{F}_{1,k}\right)$ of mechanism $M^{\text{sp}}\left(\frac{1}{2}\right)$ as $\frac{1}{2}\left(1-\mathbb{F}_{1,k}\left(\frac{1}{2}\right)\right)$ and the optimal expected revenue $\text{Rev}^{\star}(\mathbb{F}_{1,k})$ as $\max_{r\in\mathbb{R}}r\left(1-\mathbb{F}_{1,k}(r)\right)$. Let $r^{\star}$ be the optimal solution to this maximization problem, and note that it can be obtained by solving the equation $r^{\star}=\frac{1-\mathbb{F}_{1,k}(r^{\star})}{f_{1,k}(r^{\star})}$; see Myerson (1981) for details about the derivation of $r^{\star}$. One can show that $r^{\star}=\left(\frac{1}{k+1}\right)^{\frac{1}{k}}$ solves the underlying equation for $\mathbb{F}_{1,k}$. We thus have

$$\gamma\left(M^{\rm sp}\left(\frac{1}{2}\right), \mathbb{F}_{1,k}\right) = \frac{\mathrm{Rev}(M^{\rm sp}(\frac{1}{2}), \mathbb{F}_{1,k})}{\mathrm{Rev}^{\star}(\mathbb{F}_{1,k})} = \frac{\frac{1}{2}\left(1-\mathbb{F}_{1,k}\left(\frac{1}{2}\right)\right)}{r^{\star}\left(1-\mathbb{F}_{1,k}(r^{\star})\right)} = \frac{\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right)}{\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)},$$

where the last equality follows from the definition of $\mathbb{F}_{1,k}$ and $r^{\star}$. We can now write $\lim_{k\to\infty}\gamma(M^{\rm sp}\left(\frac{1}{2}\right), \mathbb{F}_{1,k})$ as

$$\lim_{k\to\infty}\frac{\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right)}{\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)} = \frac{\lim\limits_{k\to\infty}\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right)}{\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)},$$

where the equality holds because the limits of the numerator and the denominator exist and are positive, as we show below. Consider the numerator, we have $\lim\limits_{k\to\infty}\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right) = \frac{1}{2}$, and we will next show that $\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right) = 1$. Note that this implies that $\lim_{k\to\infty}\gamma(M^{\rm sp}\left(\frac{1}{2}\right), \mathbb{F}_{1,k}) = \frac{1}{2}$. To this end, note that $\lim\limits_{k\to\infty}\left(\frac{k}{k+1}\right) = 1$. We will show that $\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k} = 1$ as well. We introduce an auxiliary variable y such that $y = \left(\frac{1}{k+1}\right)^{1/k} = (k+1)^{-\frac{1}{k}}$. By taking the log on both sides, we have $\log(y) = -\frac{\log(k+1)}{k}$, which implies that $e^{\log(y)} = y = e^{-\frac{\log(k+1)}{k}}$. We now have $\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k} = \lim\limits_{k\to\infty} e^{-\frac{\log(k+1)}{k}} = e^{-\lim\limits_{k\to\infty}\frac{\log(k+1)}{k}}$. By using the L'Hôpital's rule, we can now verify that $\lim\limits_{k\to\infty}\frac{\log(k+1)}{k} = 0$, which implies that $\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k} = e^{-\lim\limits_{k\to\infty}\frac{\log(k+1)}{k}} = 1$. As $\lim\limits_{k\to\infty}\left(\frac{k}{k+1}\right) = 1$ and $\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k} = 1$, we can write the denominator limit $\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)$ as

$$\lim_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right) = \lim_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k}\lim_{k\to\infty}\left(\frac{k}{k+1}\right) = 1.$$

We thus have

$$\lim_{k\to\infty}\gamma\left(M^{\rm sp}\left(\frac{1}{2}\right), \mathbb{F}_{1,k}\right) = \frac{\lim\limits_{k\to\infty}\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right)}{\lim\limits_{k\to\infty}\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)} = \frac{1}{2},$$

and the claim follows. □

Proof of Proposition 2.15. We prove each item separately.

(i) Let $\mathbb{F}_1$ be a regular and convex distribution on the interval $\mathcal{V} = [0, \bar{v}]$. We next will prove that

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]}{\bar{v}} \leq 1,$$

which is a sufficient condition for $\mathbb{F}_1$ to satisfy the inequality (2.3) that appears in the definition of $\mathcal{G}(\mathcal{V})$ (see Remark 2.14). According to the Hermite-Hadamard theorem and because $\mathbb{F}_1$ is convex, we have

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) \leq \frac{1}{\bar{v}}\int_0^{\bar{v}} \mathbb{F}_1(v_1)\mathrm{d}v_1.$$

We transform the above inequality to the desired form.

$$\begin{aligned} &\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) \leq \frac{1}{\bar{v}}\int_0^{\bar{v}} (\mathbb{F}_1(v_1)-1+1)\mathrm{d}v_1 \\ \iff &\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) \leq -\frac{\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]}{\bar{v}} + \frac{1}{\bar{v}}\int_0^{\bar{v}} \mathrm{d}v_1. \end{aligned}$$

The if and only if condition follows from $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] = \int_0^{\bar{v}}(1-\mathbb{F}_1(v_1))\mathrm{d}v_1$.

Rearranging these terms, we obtain the inequality

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]}{\bar{v}} \leq 1,$$

and claim (i) follows.

(ii) Let $\mathbb{F}_1$ be a regular and concave distribution on the interval $\mathscr{V} = [0,\bar{v}]$. Let M denote the median of the distribution. We will show that $\mathbb{F}_1$ with median $M \geq \left(\frac{\sqrt{3}-1}{2}\right)\bar{v}$ belongs to $\mathscr{G}(\mathscr{V})$. In other words, we show that

$$\left(\frac{\sqrt{3}-1}{2}\right)\bar{v} \leq M \implies \mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \leq 1 + \frac{1}{\bar{v}}\int_0^{\bar{v}} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1.$$

To begin with, we have the following set of equivalent conditions

$$\begin{aligned} \left(\frac{\sqrt{3}-1}{2}\right)\bar{v} \leq M &\iff 0 \leq M - \left(\frac{\sqrt{3}-1}{2}\right)\bar{v} \\ &\iff 0 \leq 4\left(M + \left(\frac{\sqrt{3}+1}{2}\right)\bar{v}\right)\left(M - \left(\frac{\sqrt{3}-1}{2}\right)\bar{v}\right) \\ &\iff 0 < 4M^2 + 4\bar{v}M - 2\bar{v}^2 \\ &\iff 0 < 4M\left(\bar{v}+M\right) - 2\bar{v}^2. \end{aligned}$$

In the second step, we multiply both sides of the inequality $0 \leq M - \left(\frac{\sqrt{3}-1}{2}\right)\bar{v}$ by $4\left(M + \left(\frac{\sqrt{3}+1}{2}\right)\bar{v}\right)$, which is positive, resulting in a quadratic inequality in M. The rearrangement of the quadratic inequality yields

$$\left(\frac{\sqrt{3}-1}{2}\right)\overline{v} \le M \iff \overline{v}\frac{\overline{v}}{4M} \le \frac{\overline{v}+M}{2}. \tag{2.7}$$

Since for a concave function $\mathbb{F}_1$, the median M is less than or equal to $\frac{\overline{v}}{2}$ (as illustrated in the Figure

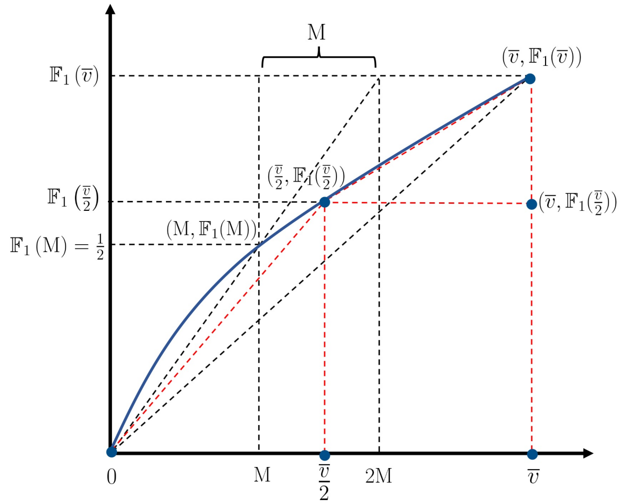

Figure 2.8: A concave distribution curve (the area enclosed by dots are for visual representation of the area of the quadrilateral discussed in the proof)

2.8). By exploiting the concavity of $\mathbb{F}_1$, we can also show that the slope of a secant line between $(M, \mathbb{F}_1(M))$ and $(\frac{\overline{v}}{2}, \mathbb{F}_1(\frac{\overline{v}}{2}))$ is less than or equal to $\frac{1}{2M}$ and that this implies $\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{\overline{v}}{4M}$ as follows:

$$\frac{\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) - \mathbb{F}_1(M)}{\frac{\overline{v}}{2} - M} \le \frac{\mathbb{F}_1(M) - \mathbb{F}_1(0)}{M-0} \iff \frac{\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) - \frac{1}{2}}{\frac{\overline{v}}{2} - M} \le \frac{1}{2M} \iff \mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{\overline{v}}{4M}$$

where the first inequality follows from the concavity of $\mathbb{F}_1$. Replacing $\frac{\overline{v}}{4M}$ in (2.7) with its lower bound $\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)$, we obtain

$$\begin{aligned}\left(\frac{\sqrt{3}-1}{2}\right)\overline{v} \le M \implies & \overline{v}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{\overline{v}+M}{2} \\ \iff & \overline{v}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{\overline{v}}{2} + \frac{M}{4} + \frac{M}{4}.\end{aligned} \tag{2.8}$$

Figure 2.8 illustrates that

$$\frac{M}{4} \le \int_0^M \mathbb{F}_1(x)\mathrm{d}x \quad \text{and} \quad \frac{M}{4} \le \int_M^{\overline{v}} (1-\mathbb{F}_1(x))\,\mathrm{d}x,$$

where $\frac{M}{4}$ is the area of the triangle with corner points $(0,0),(M,\mathbb{F}_1(M)),(M,0)$ contained in the area under the curve $\mathbb{F}_1(x)$ from 0 to M. Value $\frac{M}{4}$ is also the area of the triangle with corner points $(M,\mathbb{F}_1(M)),(M,\mathbb{F}_1(\overline{v})),(2M,\mathbb{F}_1(\overline{v}))$ that is smaller than or equal to $\int_M^{\overline{v}}(1-\mathbb{F}_1(x))\,\mathrm{d}x$. The last inequality in (2.8) thus implies

$$\overline{v}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{\overline{v}}{2} + \int_0^M \mathbb{F}_1(x)\mathrm{d}x + \int_M^{\overline{v}} (1-\mathbb{F}_1(x))\,\mathrm{d}x. \tag{2.9}$$

Dividing both sides of (2.9) by 2, and by (2.8), we have

$$\left(\frac{\sqrt{3}-1}{2}\right)\overline{v} \le M \implies \frac{\overline{v}}{2}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{\overline{v}}{4} + \int_0^M \frac{\mathbb{F}_1(x)}{2}\mathrm{d}x + \int_M^{\overline{v}} \left(\frac{1-\mathbb{F}_1(x)}{2}\right)\mathrm{d}x. \tag{2.10}$$

Now consider the quadrilateral in Figure 2.8 with corner points $(0,0),\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),(\overline{v},\mathbb{F}_1(\overline{v})),(\overline{v},0)$. Figure 2.8 illustrates that the area under $\mathbb{F}_1(x)$ over the interval $[0,\overline{v}]$, computed as $\int_0^{\overline{v}}\mathbb{F}_1(x)\,\mathrm{d}x$, is greater than or equal to the area of this quadrilateral. The area of the quadrilateral is equal to the sum of the area of the triangle formed by $(0,0),(\frac{\overline{v}}{2},0),\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right)$, the area of the rectangle formed by $(\frac{\overline{v}}{2},0),\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),\left(\overline{v},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),(\overline{v},0)$, and the area of the triangle formed by $\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),\left(\overline{v},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),(\overline{v},\mathbb{F}_1(\overline{v}))$. This area constitutes a lower bound for $\int_0^{\overline{v}}\mathbb{F}_1(x)\,\mathrm{d}x$:

$$\begin{aligned} \frac{\overline{v}}{4}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) + \frac{\overline{v}}{2}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) + \frac{\overline{v}}{4}\left(1-\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right) &\le \int_0^{\overline{v}} \mathbb{F}_1(x)\mathrm{d}x \\ \iff \frac{\overline{v}}{2}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) + \frac{\overline{v}}{4} &\le \int_0^{\overline{v}} \mathbb{F}_1(x)\mathrm{d}x, \end{aligned} \tag{2.11}$$

where the three terms at the left-hand side of the first inequality in (2.11) represent the area of triangle formed by $(0,0),(\frac{\overline{v}}{2},0),\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right)$, the area of rectangle formed by $(\frac{\overline{v}}{2},0),\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),\left(\overline{v},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),(\overline{v},0)$ and the area of triangle formed by $\left(\frac{\overline{v}}{2},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),\left(\overline{v},\mathbb{F}_1\left(\frac{\overline{v}}{2}\right)\right),(\overline{v},\mathbb{F}_1(\overline{v}))$, respectively. We add the final inequality to the second inequality in (2.10) to get

$$\begin{aligned} &\overline{v}\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \int_0^{\overline{v}} \mathbb{F}_1(x)\mathrm{d}x + \int_0^M \frac{\mathbb{F}_1(x)}{2}\mathrm{d}x + \int_M^{\overline{v}}\left(\frac{1-\mathbb{F}_1(x)}{2}\right)\mathrm{d}x \\ \iff &\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) \le \frac{1}{\overline{v}}\int_0^M \left(\mathbb{F}_1(x) + \frac{\mathbb{F}_1(x)}{2}\right)\mathrm{d}x + \frac{1}{\overline{v}}\int_M^{\overline{v}}\left(\mathbb{F}_1(x) + \frac{1-\mathbb{F}_1(x)}{2}\right)\mathrm{d}x. \end{aligned} \tag{2.12}$$

Consider the terms $\frac{\mathbb{F}_1(x)}{2}$ and $\frac{1-\mathbb{F}_1(x)}{2}$ in (2.12). We next show that these terms are bounded from above by $\mathbb{F}_1(x)-\mathbb{F}_1(x)^2$ when x is in $[0,M]$ and $[M,\overline{v}]$, respectively. Since $\mathbb{F}_1(x)\in[0,1]$, we always have $\mathbb{F}_1(x)-\mathbb{F}_1(x)^2\ge 0$. Furthermore, if $x\in[0,\mathbb{F}_1^{-1}(\frac{1}{2})]=[0,M]$, in which case $\mathbb{F}_1(x)\in[0,\frac{1}{2}]$, we have $\frac{\mathbb{F}_1(x)}{2} \le \mathbb{F}_1(x) - \mathbb{F}_1(x)^2$. Similarly, if $x \in [\mathbb{F}_1^{-1}(\frac{1}{2}), \bar{v}] = [M, \bar{v}]$, in which case $\mathbb{F}_1(x) \in [\frac{1}{2}, 1]$, and we have $\frac{1-\mathbb{F}_1(x)}{2} \le \mathbb{F}_1(x) - \mathbb{F}_1(x)^2$. By replacing the two terms by their upper bound $\mathbb{F}_1(x) - \mathbb{F}_1(x)^2$ in (2.12), we obtain

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) \le \frac{1}{\bar{v}} \int_0^M \left(\mathbb{F}_1(x) + \frac{\mathbb{F}_1(x)}{2}\right) \mathrm{d}x + \frac{1}{\bar{v}} \int_M^{\bar{v}} \left(\mathbb{F}_1(x) + \frac{1 - \mathbb{F}_1(x)}{2}\right) \mathrm{d}x$$
$$\implies \mathbb{F}_1\left(\frac{\bar{v}}{2}\right) \le \frac{1}{\bar{v}} \left(\int_0^M \left(\mathbb{F}_1(x) + \mathbb{F}_1(x) - \mathbb{F}_1(x)^2\right) \mathrm{d}x + \int_M^{\bar{v}} \left(\mathbb{F}_1(x) + \mathbb{F}_1(x) - \mathbb{F}_1(x)^2\right) \mathrm{d}x\right).$$

Noting that $\int_0^{\bar{v}} \mathbb{F}_1(v_1) \mathrm{d}v_1 = \bar{v} - \mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]$, we thus showed that

$$\left(\frac{\sqrt{3}-1}{2}\right)\bar{v} \le M \implies \mathbb{F}_1\left(\frac{\bar{v}}{2}\right) \le \frac{1}{\bar{v}} \int_0^{\bar{v}} \left(\mathbb{F}_1(x) + \mathbb{F}_1(x) - \mathbb{F}_1(x)^2\right) \mathrm{d}x$$
$$\iff \mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}} \mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \le 1 + \frac{1}{\bar{v}} \int_0^{\bar{v}} \mathbb{F}_1(x)(1 - \mathbb{F}_1(x)) \mathrm{d}x.$$

This concludes the proof of *(ii)*.

(iii) Let $\mathbb{F}_1$ be a triangular distribution on the interval $\mathscr{V} = [0, \bar{v}]$. We will show that $\mathbb{F}_1$ is regular and satisfies condition (2.3) and thus $\mathbb{F}_1 \in \mathscr{G}(\mathscr{V})$. To this end, we first demonstrate that $\mathbb{F}_1$ is regular. Particularly, we will show that $\mathbb{F}_1$ has non-decreasing virtual value. Let c be the mode, then by the definition of triangular distribution, for v_1 in the interval $[0, c]$ and in $(c, \bar{v}]$, the pdf and cdf are defined as

$$f_1(v_1) = \begin{cases} \frac{2v_1}{\bar{v}c} & \text{for } 0 \le v_1 \le c \\ \frac{2(\bar{v}-v_1)}{\bar{v}(\bar{v}-c)} & \text{for } c < v_1 \le \bar{v} \end{cases} \quad \text{and} \quad \mathbb{F}_1(v_1) = \begin{cases} \frac{v_1^2}{\bar{v}c} & \text{for } 0 \le v_1 \le c \\ 1 - \frac{(\bar{v}-v_1)^2}{\bar{v}(\bar{v}-c)} & \text{for } c < v_1 \le \bar{v}. \end{cases} \tag{2.13}$$

Next, we examine the virtual value for v_1 within the interval $(0, \bar{v})$:

$$\varphi_1(v_1) = \begin{cases} \frac{3v_1}{2} - \frac{\bar{v}c}{2v_1} & \text{for } 0 < v_1 \le c \\ \frac{3v_1}{2} - \frac{\bar{v}}{2} & \text{for } c < v_1 < \bar{v}. \end{cases}$$

In both cases above, one can notice that the derivative with respect to v_1 is non-negative, indicating that $\varphi_1(v_1)$ is indeed non-decreasing in each of the two intervals. Since $\varphi_1(v_1)$ is continuous at c, we can conclude that the triangular distribution is regular.

We next prove that $\mathbb{F}_1$ satisfies (2.3) by a two-step process. In the first step, we consider the case $c > \frac{\bar{v}}{2}$ and show that $\mathbb{F}_1$ satisfies the condition in Remark 2.14, which is a sufficient condition for (2.3). In the second step, we consider the case $c \le \frac{\bar{v}}{2}$ and show that $\mathbb{F}_1$ satisfies the condition (2.3). Since both conditions require the values of $\mathbb{F}_1(\frac{\bar{v}}{2})$ and $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]$, we evaluate $\mathbb{F}_1(\frac{\bar{v}}{2})$ by substituting

$v_1 = \frac{\bar{v}}{2}$ into the cdf

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) = \begin{cases} \frac{\bar{v}}{4c} & \text{for } 0 \leq \frac{\bar{v}}{2} \leq c \\ 1 - \frac{\bar{v}}{4(\bar{v}-c)} & \text{for } c < \frac{\bar{v}}{2} \leq \bar{v}, \end{cases}$$

and the expected value is $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] = \frac{\bar{v}+c}{3}$.

Step 1: We consider $c > \frac{\bar{v}}{2}$ and is given by $c = \frac{\bar{v}}{2} + \delta$, where $\delta \in (0, \frac{\bar{v}}{2}]$. Substituting this value of c into $\mathbb{F}_1(\frac{\bar{v}}{2})$ and $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]$, we obtain:

$$\begin{aligned}\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] &= \frac{\bar{v}}{2\bar{v}+4\delta} + \frac{3\bar{v}+2\delta}{6\bar{v}} = \frac{6\bar{v}^2 + (2\bar{v}+4\delta)(3\bar{v}+2\delta)}{12\bar{v}^2+24\bar{v}\delta} \\ &= \frac{12\bar{v}^2+16\bar{v}\delta+8\delta^2}{12\bar{v}^2+24\bar{v}\delta} \leq \frac{12\bar{v}^2+16\bar{v}\delta+8\bar{v}\delta}{12\bar{v}^2+24\bar{v}\delta} = 1.\end{aligned}$$

The last equality follows from $\delta \in (0, \frac{\bar{v}}{2}]$ and, therefore, $\bar{v} > \delta$. Thus, $\mathbb{F}_1$ fulfils the sufficient condition in Remark 2.14.

Step 2: We now consider $c \leq \frac{\bar{v}}{2}$, where $c = \frac{\bar{v}}{2} - \delta$ for a $\delta \in [0, \frac{\bar{v}}{2}]$. We have

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] = \frac{\bar{v}+4\delta}{2\bar{v}+4\delta} + \frac{3\bar{v}-2\delta}{6\bar{v}}.$$

Next, we simplify the integral $\frac{1}{\bar{v}}\int_0^{\bar{v}} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1$ by integrating the expression by separating it over the intervals $[0,c] = [0, \frac{\bar{v}}{2} - \delta]$ and $[c,\bar{v}] = [\frac{\bar{v}}{2} - \delta, \bar{v}]$.

$$\begin{aligned}&\frac{1}{\bar{v}}\int_0^{\bar{v}} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1 \\ &= \frac{1}{\bar{v}}\left(\int_0^{\frac{\bar{v}}{2}-\delta} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1 + \int_{\frac{\bar{v}}{2}-\delta}^{\bar{v}} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1\right)\end{aligned} \tag{2.14}$$

We evaluate the two integrals separately. The first integral amounts to

$$\begin{aligned}&\int_0^{\frac{\bar{v}}{2}-\delta} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1 = \int_0^{\frac{\bar{v}}{2}-\delta} \frac{v_1^2}{\bar{v}(\frac{\bar{v}}{2}-\delta)}\left(1 - \frac{v_1^2}{\bar{v}(\frac{\bar{v}}{2}-\delta)}\right)\mathrm{d}v_1 \\ &= \frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}-\delta\right)^2}\int_0^{\frac{\bar{v}}{2}-\delta} v_1^2\left(\frac{\bar{v}^2}{2} - \bar{v}\delta - v_1^2\right)\mathrm{d}v_1 = \frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}-\delta\right)^2}\left(\frac{v_1^3}{3}\left(\frac{\bar{v}^2}{2} - \bar{v}\delta\right) - \frac{v_1^5}{5}\right)\Bigg|_0^{\frac{\bar{v}}{2}-\delta} \\ &= \frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}-\delta\right)^2}\left[\frac{\bar{v}}{3}\left(\frac{\bar{v}}{2}-\delta\right)^4 - \frac{1}{5}\left(\frac{\bar{v}}{2}-\delta\right)^5\right],\end{aligned}$$

and the second integral amounts to

$$\int_{\frac{\bar{v}}{2}-\delta}^{\bar{v}} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1 = \int_{\frac{\bar{v}}{2}-\delta}^{\bar{v}} \left(1-\frac{(\bar{v}-v_1)^2}{\bar{v}(\frac{\bar{v}}{2}+\delta)}\right)\left(\frac{(\bar{v}-v_1)^2}{\bar{v}(\frac{\bar{v}}{2}+\delta)}\right)\mathrm{d}v_1$$
$$= \frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}+\delta\right)^2}\int_{\frac{\bar{v}}{2}-\delta}^{\bar{v}}\left(\frac{\bar{v}^2}{2}+\bar{v}\delta-(\bar{v}-v_1)^2\right)(\bar{v}-v_1)^2\,\mathrm{d}v_1$$
$$= \frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}+\delta\right)^2}\left(\left(\frac{\bar{v}^2}{2}+\bar{v}\delta\right)\left(\frac{v_1^3}{3}-\bar{v}v_1^2+\bar{v}^2v_1\right)-\frac{1}{5}(v_1-\bar{v})^5\right)\Big|_{\frac{\bar{v}}{2}-\delta}^{\bar{v}}$$
$$= \frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}+\delta\right)^2}\left[\left(\frac{\bar{v}^2}{2}+\bar{v}\delta\right)\left(\frac{\bar{v}^3}{3}-\frac{1}{3}\left(\frac{\bar{v}}{2}-\delta\right)^3+\bar{v}\left(\frac{\bar{v}}{2}-\delta\right)^2-\bar{v}^2\left(\frac{\bar{v}}{2}-\delta\right)\right)-\frac{1}{5}\left(\frac{\bar{v}}{2}+\delta\right)^5\right].$$

Considering the above equation, we also observe that

$$\left(\frac{\bar{v}^3}{3}-\frac{1}{3}\left(\frac{\bar{v}}{2}-\delta\right)^3+\bar{v}\left(\frac{\bar{v}}{2}-\delta\right)^2-\bar{v}^2\left(\frac{\bar{v}}{2}-\delta\right)\right)=\frac{1}{3}\left(\bar{v}-\left(\frac{\bar{v}}{2}-\delta\right)\right)^3=\frac{1}{3}\left(\frac{\bar{v}}{2}+\delta\right)^3,$$

where the first equality is obtained using the identity $a^3-b^3+3ab^2-3a^2b=(a-b)^3$ with $a=\bar{v}$ and $b=\left(\frac{\bar{v}}{2}-\delta\right)$. Upon substituting the obtained expressions for the two integrals back into (2.14), we obtain

$$\frac{1}{\bar{v}}\int_0^{\bar{v}} \mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1$$
$$= \frac{1}{\bar{v}}\left[\frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}-\delta\right)^2}\left(\frac{\bar{v}}{3}\left(\frac{\bar{v}}{2}-\delta\right)^4-\frac{1}{5}\left(\frac{\bar{v}}{2}-\delta\right)^5\right)+\frac{1}{\bar{v}^2\left(\frac{\bar{v}}{2}+\delta\right)^2}\left(\frac{\bar{v}}{3}\left(\frac{\bar{v}}{2}+\delta\right)^4-\frac{1}{5}\left(\frac{\bar{v}}{2}+\delta\right)^5\right)\right]$$
$$= \left[\frac{1}{3\bar{v}^2}\left(\left(\frac{\bar{v}}{2}-\delta\right)^2+\left(\frac{\bar{v}}{2}+\delta\right)^2\right)-\frac{1}{5\bar{v}^3}\left(\left(\frac{\bar{v}}{2}-\delta\right)^3+\left(\frac{\bar{v}}{2}+\delta\right)^3\right)\right]$$
$$= \frac{2}{3}\left(\frac{1}{4}+\left(\frac{\delta}{\bar{v}}\right)^2\right)-\frac{2}{5}\left(\frac{1}{8}+\frac{3}{2}\left(\frac{\delta}{\bar{v}}\right)^2\right)=\frac{7}{60}+\frac{1}{15}\left(\frac{\delta}{\bar{v}}\right)^2.$$

Condition (2.3) can now be expressed as

$$\begin{aligned}
&\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)+\frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1]\le 1+\frac{1}{\bar{v}}\int_0^{\bar{v}}\mathbb{F}_1(v_1)(1-\mathbb{F}_1(v_1))\mathrm{d}v_1\\
&\iff \frac{\bar{v}+4\delta}{2\bar{v}+4\delta}+\frac{3\bar{v}-2\delta}{6\bar{v}}\le 1+\frac{7}{60}+\frac{1}{15}\left(\frac{\delta}{\bar{v}}\right)^2\\
&\iff 0\le\frac{67}{60}+\frac{1}{15}\rho^2-\frac{1+4\rho}{2+4\rho}-\frac{3-2\rho}{6},
\end{aligned}\tag{2.15}$$

where $\rho = \frac{\delta}{\bar{v}} \in [0, \frac{1}{2}]$. By multiplying both sides of (2.15) by $60(2+4\rho)$, we obtain

$$0 \leq 60(2+4p)\left(\frac{67}{60} + \frac{1}{15}\rho^2 - \frac{1+4\rho}{2+4\rho} - \frac{3-2\rho}{6}\right) \iff 0 \leq 16\rho^3 + 88\rho^2 - 52\rho + 14$$
$$\iff 0 \leq 8\rho^3 + 44\rho^2 - 26\rho + 7.$$

It is now sufficient to show that $8\rho^3 + 44\rho^2 - 26\rho + 7$ is non-negative for any $\rho \in [0, \frac{1}{2}]$ to complete the proof for Step 2. To prove this, it is sufficient to show that the minimum value of polynomial is non-negative for $\rho \in [0, \frac{1}{2}]$. Noting that $8\rho^3 + 44\rho^2 - 26\rho + 7$ is strictly convex on the interval $[0, \frac{1}{2}]$ and differentiable, we set its derivative equal to zero to find the minimizer:

$$\frac{\mathrm{d}}{\mathrm{d}\rho}(8\rho^3 + 44\rho^2 - 26\rho + 7) = 24\rho^2 + 88\rho - 26 = 0.$$

The above equality is solved by $\rho = \frac{-88+32\sqrt{10}}{48} \approx 0.2748$, which is the minimizer of the polynomial function across $[0, \frac{1}{2}]$. One can verify that the polynomial evaluated at this minimizer amounts to a non-negative value. This concludes the proof for Step 2.

Therefore, $\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \leq 1 + \frac{1}{\bar{v}}\int_0^{\bar{v}} \mathbb{F}_1(v_1)(1 - \mathbb{F}_1(v_1))\mathrm{d}v_1$ holds for any triangular distribution. This concludes the proof of (iii).

(iv) Let $\mathbb{F}_1$ be a regular scaled beta distribution on the interval $\mathscr{V} = [0, \bar{v}]$ with $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \geq \frac{\bar{v}}{2}$. Note that $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \geq \frac{\bar{v}}{2}$ is true if and only if $\alpha \geq \beta$ (see Remark (2.18)).

We show that when $\alpha \geq \beta$ (and equivalently when $\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \geq \frac{\bar{v}}{2}$), $\mathbb{F}_1$ satisfies the condition in Remark 2.14, which is a sufficient condition for (2.3). Fix any such α and β, and denote by $\mathbb{F}_1$ and f_1 its cdf and pdf, respectively:

$$\mathbb{F}_1(v_1) = \frac{1}{B(\alpha,\beta)} \int_0^{v_1} \left(\frac{t}{\bar{v}}\right)^{\alpha-1} \left(1 - \frac{t}{\bar{v}}\right)^{\beta-1} \mathrm{d}t \quad \text{and} \quad f_1(v_1) = \frac{1}{B(\alpha,\beta)} \left(\frac{v_1}{\bar{v}}\right)^{\alpha-1} \left(1 - \frac{v_1}{\bar{v}}\right)^{\beta-1},$$

where the domain of $\mathbb{F}_1$ and f_1 is $[0, \bar{v}]$, and $B(\alpha, \beta)$ is the beta function defined as

$$B(\alpha,\beta) = \int_0^{\bar{v}} \left(\frac{t}{\bar{v}}\right)^{\alpha-1} \left(1 - \frac{t}{\bar{v}}\right)^{\beta-1} \mathrm{d}t.$$

We next demonstrate that sufficient condition $\mathbb{F}_1(\frac{\bar{v}}{2}) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \leq 1$ for (2.3) holds if and only if $\alpha \geq \beta$. To this end, we will start from the inequality $\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \leq 1$ and show that it can equivalently be expressed as $\alpha \geq \beta$. This inequality can be rewritten as

$$\mathbb{F}_1\left(\frac{\bar{v}}{2}\right) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \leq \frac{\int_0^{\bar{v}} t^{\alpha-1}(\bar{v}-t)^{\beta-1}\mathrm{d}t}{\int_0^{\bar{v}} t^{\alpha-1}(\bar{v}-t)^{\beta-1}\mathrm{d}t}. \tag{2.16}$$

Note that

$$\mathbb{F}_1\left(\frac{\overline{v}}{2}\right) = \frac{\int_0^{\frac{\overline{v}}{2}} t^{\alpha-1}\left(\overline{v}-t\right)^{\beta-1}\mathrm{d}t}{\int_0^{\overline{v}} t^{\alpha-1}\left(\overline{v}-t\right)^{\beta-1}\mathrm{d}t} \quad \text{and} \quad \mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] = \frac{\int_0^{\overline{v}} t t^{\alpha-1}\left(\overline{v}-t\right)^{\beta-1}\mathrm{d}t}{\int_0^{\overline{v}} t^{\alpha-1}\left(\overline{v}-t\right)^{\beta-1}\mathrm{d}t}.$$

Substituting these definitions into (2.16), we have

$$\frac{\int_0^{\frac{\overline{v}}{2}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t}{\int_0^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t} + \frac{1}{\overline{v}}\frac{\int_0^{\overline{v}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t}{\int_0^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t} \leq \frac{\int_0^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t}{\int_0^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t}.$$

We eliminate common denominator to obtain

$$\begin{aligned}
&\int_0^{\frac{\overline{v}}{2}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t + \frac{1}{\overline{v}}\int_0^{\overline{v}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \int_0^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t.\\
&\iff \frac{1}{\overline{v}}\int_0^{\overline{v}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \int_{\frac{\overline{v}}{2}}^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t\\
&\iff \int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t + \int_{\frac{\overline{v}}{2}}^{\overline{v}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \overline{v}\int_{\frac{\overline{v}}{2}}^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t\\
&\iff \int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \int_{\frac{\overline{v}}{2}}^{\overline{v}} \overline{v}t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t - \int_{\frac{\overline{v}}{2}}^{\overline{v}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t\\
&\iff \int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \int_{\frac{\overline{v}}{2}}^{\overline{v}} t^{\alpha-1}(\overline{v}-t)^{\beta}\mathrm{d}t.
\end{aligned}$$

Next, we apply a change of variables. To this end, we let $y = \overline{v} - t$ and $\mathrm{d}y = -\mathrm{d}t$ and substitute the new variable y into the right-hand side of the above inequality. Note that this substitution also commands a change in the limits of the integration on the right hand side (the lower bound, $\frac{\overline{v}}{2}$, upon substituting in the definition of y, yields $\overline{v} - \frac{\overline{v}}{2} = \frac{\overline{v}}{2}$, and similarly, the upper bound, $\overline{v}$, yields $\overline{v} - \overline{v} = 0$).

$$\int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq -\int_{\frac{\overline{v}}{2}}^{0} (\overline{v}-y)^{\alpha-1}y^{\beta}\mathrm{d}y \iff \int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \int_0^{\frac{\overline{v}}{2}} (\overline{v}-y)^{\alpha-1}y^{\beta}\mathrm{d}y.$$

Changing y back to t at the right-hand side, we obtain

$$\begin{aligned}
&\int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t \leq \int_0^{\frac{\overline{v}}{2}} (\overline{v}-t)^{\alpha-1}t^{\beta}\mathrm{d}t\\
&\iff 0 \leq \int_0^{\frac{\overline{v}}{2}} (\overline{v}-t)^{\alpha-1}t^{\beta}\mathrm{d}t - \int_0^{\frac{\overline{v}}{2}} t\ t^{\alpha-1}(\overline{v}-t)^{\beta-1}\mathrm{d}t\\
&\iff 0 \leq \int_0^{\frac{\overline{v}}{2}} t^{\beta}(\overline{v}-t)^{\beta-1}\left[(\overline{v}-t)^{\alpha-\beta} - t^{\alpha-\beta}\right]\mathrm{d}t.
\end{aligned}$$

Note that for any $t \in [0, \frac{\bar{v}}{2}]$, the expression $t^{\beta}(\bar{v}-t)^{\beta-1}$ is non-negative and $(\bar{v}-t)^{\alpha-\beta} - t^{\alpha-\beta}$ is also non-negative because $(\bar{v}-t) \geq t$ and $\alpha \geq \beta$. Therefore, $\mathbb{F}_1(\frac{\bar{v}}{2}) + \frac{1}{\bar{v}}\mathbb{E}_{\mathbb{F}_1}[\tilde{v}_1] \leq 1$ holds for a scaled beta distribution when $\alpha \geq \beta$. Hence $\mathbb{F}_1 \in \mathscr{G}(\mathscr{V})$. This concludes the proof of (iv). □

Proofs for competitive ratios in Table 2.1. The competitive ratios are characterized for the single-bidder case, where the bidder's value $\tilde{v}_1 \sim \mathbb{F}_1$. All studied distributions in Table 1 are regular (Bagnoli and Bergstrom, 2005; Ewerhart, 2013). Recall that we can express $\text{Rev}\left(M^{\text{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}_1\right)$ as $\frac{\bar{v}}{2}\left(1-\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)$ and $\text{Rev}^{\star}(\mathbb{F}_1)$ as $r^{\star}(1-\mathbb{F}_1(r^{\star}))$, where $r^{\star}$ represents the optimal Myerson price obtained by solving the equation $r^{\star} = \frac{1-\mathbb{F}_1(r^{\star})}{f_1(r^{\star})}$. Therefore, we can write the competitive ratio as

$$\frac{\text{Rev}(M^{\text{sp}}(\frac{\bar{v}}{2}), \mathbb{F}_1)}{\text{Rev}^{\star}(\mathbb{F}_1)} = \frac{\frac{\bar{v}}{2}\left(1-\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)}{r^{\star}(1-\mathbb{F}_1(r^{\star}))}.$$

In the following, we find a closed form for the competitive ratio in case of (i) triangular distributions, (ii) power distributions, and (iii) exponential distributions. For beta distributions, our analysis is limited to computing the competitive ratio numerically as displayed in Figure 2.2.

(i) **Triangular distribution:** Let $\mathbb{F}_1$ be a triangular distribution on the interval $[0, \bar{v}]$. Let c be the mode. The pdf and cdf of this distribution are defined in equation (2.13). By subsituting the definitions of pdf and cdf into the definition of the virtual value, we have

$$\varphi_1(v_1) = \begin{cases} \frac{3v_1}{2} - \frac{\bar{v}c}{2v_1} & \text{for } 0 < v_1 \leq c \\ \frac{3v_1}{2} - \frac{\bar{v}}{2} & \text{for } c < v_1 < \bar{v}. \end{cases}$$

By definition of the virtual value, the optimal Myerson price is given by $r^{\star}$ such that $\varphi_1(r^{\star}) = 0$. We will characterize $r^{\star}$ in closed form for two cases separately: Case 1 ($0 \leq c \leq \frac{\bar{v}}{3}$) and Case 2 ($\frac{\bar{v}}{3} < c \leq \bar{v}$).

Case 1 ($0 \leq c \leq \frac{\bar{v}}{3}$): Fix any c such that $0 \leq c \leq \frac{\bar{v}}{3}$. Observe that when $0 < v_1 \leq c$, virtual value is non-positive. Therefore, the optimal reserve price $r^{\star}$ is determined by $\frac{3r}{2} - \frac{\bar{v}}{2} = 0$, which yields $r^{\star} = \frac{\bar{v}}{3}$.

Case 2 ($\frac{\bar{v}}{3} < c \leq \bar{v}$): Fix any c such that $\frac{\bar{v}}{3} < c \leq \bar{v}$. Observe that when $c \leq v_1 < \bar{v}$, virtual value is positive. Therefore, the optimal reserve price $r^{\star}$ is determined by $\frac{3r}{2} - \frac{\bar{v}c}{2r} = 0$, implying $r^{\star} = \sqrt{\frac{\bar{v}c}{3}}$.

We therefore have

$$r^{\star} = \begin{cases} \frac{\bar{v}}{3} & \text{for } 0 \leq c \leq \frac{\bar{v}}{3} \\ \sqrt{\frac{\bar{v}c}{3}} & \text{for } \frac{\bar{v}}{3} < c \leq \bar{v} \end{cases} \quad \text{and} \quad 1-\mathbb{F}_1(r^{\star}) = \begin{cases} \frac{(\bar{v}-\frac{\bar{v}}{3})^2}{\bar{v}(\bar{v}-c)} & \text{for } 0 \leq c \leq \frac{\bar{v}}{3} \\ \frac{2}{3} & \text{for } \frac{\bar{v}}{3} < c \leq \bar{v}. \end{cases}$$

Next, we derive the competitive ratio in closed form in three steps each covering a distinct case: Step 1 ($0 \leq c \leq \frac{\bar{v}}{3}$), Step 2 ($\frac{\bar{v}}{3} < c \leq \frac{\bar{v}}{2}$) and Step 3 ($\frac{\bar{v}}{2} < c \leq \bar{v}$).

Step 1 ($0 \leq c \leq \frac{\bar{v}}{3}$): In this case, we know that $r^\star = \frac{\bar{v}}{3}$ and $1 - \mathbb{F}_1(r^\star) = \frac{(\bar{v}-\frac{\bar{v}}{3})^2}{\bar{v}(\bar{v}-c)}$. As $\frac{\bar{v}}{2} > \frac{\bar{v}}{3} \geq c$, we have $1 - \mathbb{F}_1(\frac{\bar{v}}{2}) = \frac{(\bar{v}-\frac{\bar{v}}{2})^2}{\bar{v}(\bar{v}-c)}$. By substituting these values into the competitive ratio, we find

$$\frac{\mathrm{Rev}\left(M^{\mathrm{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}_1\right)}{\mathrm{Rev}^\star(\mathbb{F}_1)} = \frac{\frac{\bar{v}}{2}\left(1 - \mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)}{r^\star\left(1 - \mathbb{F}_1(r^\star)\right)} = \frac{27}{32}.$$

Step 2 ($\frac{\bar{v}}{3} < c \leq \frac{\bar{v}}{2}$): In this case, we know that $r^\star = \sqrt{\frac{\bar{v}c}{3}}$ and $1 - \mathbb{F}_1(r^\star) = \frac{2}{3}$. As $\frac{\bar{v}}{2} \geq c$ by assumption, we have $1 - \mathbb{F}_1(\frac{\bar{v}}{2}) = \frac{(\bar{v}-\frac{\bar{v}}{2})^2}{\bar{v}(\bar{v}-c)}$. By substituting these values into the competitive ratio, we find

$$\frac{\mathrm{Rev}\left(M^{\mathrm{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}_1\right)}{\mathrm{Rev}^\star(\mathbb{F}_1)} = \frac{\frac{\bar{v}}{2}\left(1 - \mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)}{r^\star\left(1 - \mathbb{F}_1(r^\star)\right)} = \frac{3\sqrt{3}\bar{v}^2}{16\sqrt{\bar{v}c}(\bar{v}-c)}.$$

Step 3 ($\frac{\bar{v}}{2} < c \leq \bar{v}$): As $\frac{\bar{v}}{3} < \frac{\bar{v}}{2} < c$, we have $r^\star = \sqrt{\frac{\bar{v}c}{3}}$ and $1 - \mathbb{F}_1(r^\star) = \frac{2}{3}$. As $\frac{\bar{v}}{2} < c$, we also have $1 - \mathbb{F}_1(\frac{\bar{v}}{2}) = 1 - \frac{\bar{v}^2}{4\bar{v}c}$. By substituting these values into the competitive ratio, we find

$$\frac{\mathrm{Rev}\left(M^{\mathrm{sp}}\left(\frac{\bar{v}}{2}\right), \mathbb{F}_1\right)}{\mathrm{Rev}^\star(\mathbb{F}_1)} = \frac{\frac{\bar{v}}{2}\left(1 - \mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)}{r^\star\left(1 - \mathbb{F}_1(r^\star)\right)} = \frac{3\sqrt{3}\bar{v}(4c-\bar{v})}{16\sqrt{\bar{v}}c^{\frac{3}{2}}}.$$

(ii) **Power distribution:** Let $\mathbb{F}_1$ be a power distribution. The pdf and cdf of this distribution are defined as $f_1(v_1) = kv_1^{k-1}$ and $\mathbb{F}_1(v_1) = v_1^k$ for some $k \geq 1$. By substituting these definitions into $r^\star = \frac{1-\mathbb{F}_1(r^\star)}{f_1(r^\star)}$, we have

$$r^\star = \frac{1-(r^\star)^k}{k(r^\star)^{k-1}} \iff k(r^\star)^k - 1 + (r^\star)^k = 0 \iff r^\star = \left(\frac{1}{k+1}\right)^{\frac{1}{k}}.$$

Substituting $\frac{\bar{v}}{2} = \frac{1}{2}$ and $r^\star = \left(\frac{1}{k+1}\right)^{\frac{1}{k}}$ into the competitive ratio, we find

$$\frac{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\bar{v}}{2}), \mathbb{F}_1)}{\mathrm{Rev}^\star(\mathbb{F}_1)} = \frac{\frac{\bar{v}}{2}\left(1 - \mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)}{r^\star\left(1 - \mathbb{F}_1(r^\star)\right)} = \frac{\frac{1}{2}\left(1 - \left(\frac{1}{2}\right)^k\right)}{\left(\frac{1}{k+1}\right)^{\frac{1}{k}}\left(1 - \left(\frac{1}{k+1}\right)\right)} = \frac{\left(\frac{1}{2}\right)\left(\frac{2^k-1}{2^k}\right)}{\left(\frac{1}{k+1}\right)^{1/k}\left(\frac{k}{k+1}\right)}.$$

(iii) **Exponential distribution:**

Let $\mathbb{F}_1$ be a truncated exponential distribution on the interval $[0, \bar{v}]$. The pdf and cdf of this distribution are defined as $f_1(v_1) = \frac{\lambda e^{-\lambda v_1}}{1-e^{-\lambda \bar{v}}}$ and $\mathbb{F}_1(v_1) = \frac{1-e^{-\lambda v_1}}{1-e^{-\lambda \bar{v}}}$. By substituting these definitions into $r^\star = \frac{1-\mathbb{F}_1(r^\star)}{f_1(r^\star)}$, we have

$$r^\star = \frac{1-\frac{1-e^{-\lambda r^\star}}{1-e^{-\lambda \bar{v}}}}{\frac{\lambda e^{-\lambda r^\star}}{1-e^{-\lambda \bar{v}}}} \iff \lambda r^\star e^{-\lambda r^\star} - e^{-\lambda r^\star} = -e^{-\lambda \bar{v}} \iff 1-\lambda r^\star = W(e^{1-\lambda \bar{v}}),$$

where W is the Lambert W function. Substituting $r^\star = \frac{1-W(e^{1-\lambda \bar{v}})}{\lambda}$ into the competitive ratio, we find

$$\begin{aligned}\frac{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\bar{v}}{2}),\mathbb{F}_1)}{\mathrm{Rev}^\star(\mathbb{F}_1)} = \frac{\frac{\bar{v}}{2}\left(1-\mathbb{F}_1\left(\frac{\bar{v}}{2}\right)\right)}{r^\star\left(1-\mathbb{F}_1(r^\star)\right)} &= \frac{\frac{\bar{v}}{2}\left(1-\frac{1-e^{-\frac{\lambda\bar{v}}{2}}}{1-e^{-\lambda\bar{v}}}\right)}{\frac{1-W(e^{1-\lambda\bar{v}})}{\lambda}\left(1-\frac{1-e^{-\lambda\frac{1-W(e^{1-\lambda\bar{v}})}{\lambda}}}{1-e^{-\lambda\bar{v}}}\right)} \\ &= \frac{\lambda\bar{v}}{2\left(1-W(e^{1-\lambda\bar{v}})\right)}\left(\frac{e^{\frac{-\lambda\bar{v}}{2}}-e^{-\lambda\bar{v}}}{e^{\left(W(e^{1-\lambda\bar{v}})-1\right)}-e^{-\lambda\bar{v}}}\right).\end{aligned}$$

This completes the proof. □

Proof of Proposition 2.19. We use the fact that for $\mathbb{F} \in \mathscr{G}(\mathscr{V}^N)$, $M^{\mathrm{sp}}(\frac{\bar{v}}{2})$ guarantees 2-approximation for $N=1$, and Theorems 2.11 and 2.12 to find a lower bound on competitive ratio for $M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2})$. Consider any $\bar{v} \in [\hat{v}-\varepsilon, \hat{v}+\varepsilon]$, any $i \in \mathscr{N}$ and any distribution $\mathbb{F}_i \in \mathscr{G}(\mathscr{V})$. We have the following inequality:

$$\frac{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2}),\mathbb{F}_i)}{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\bar{v}}{2}),\mathbb{F}_i)} = \frac{(\frac{\hat{v}-\varepsilon}{2})(1-\mathbb{F}_i(\frac{\hat{v}-\varepsilon}{2}))}{\frac{\bar{v}}{2}(1-\mathbb{F}_i(\frac{\bar{v}}{2}))} \geq \frac{(\frac{\hat{v}-\varepsilon}{2})(1-\mathbb{F}_i(\frac{\hat{v}-\varepsilon}{2}))}{\frac{\bar{v}}{2}(1-\mathbb{F}_i(\frac{\hat{v}-\varepsilon}{2}))} = \frac{\hat{v}-\varepsilon}{\bar{v}} \geq \frac{\hat{v}-\varepsilon}{\hat{v}+\varepsilon}.$$

The first expression compares the seller's revenue obtained through two mechanisms: $M^{\mathrm{sp}}\left(\frac{\hat{v}-\varepsilon}{2}\right)$ and $M^{\mathrm{sp}}\left(\frac{\bar{v}}{2}\right)$, where $\bar{v}$ represents the true upper bound. Thanks to the non-increasing nature of the function $1-\mathbb{F}_i(v_i)$, we have the inequality $\left(1-\mathbb{F}_i\left(\frac{\bar{v}}{2}\right)\right) \leq \left(1-\mathbb{F}_i\left(\frac{\hat{v}-\varepsilon}{2}\right)\right)$. The first inequality follows from this observation. The last inequality follows from the fact that $\bar{v} \leq \hat{v}+\varepsilon$.

We express the competitive ratio as shown to derive the following bound:

$$\gamma\left(M^{\mathrm{sp}}\left(\frac{\hat{v}-\varepsilon}{2}\right),\mathbb{F}_i\right) = \frac{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2}),\mathbb{F}_i)}{\mathrm{Rev}^\star(\mathbb{F}_i)} = \frac{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2}),\mathbb{F}_i)}{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\bar{v}}{2}),\mathbb{F}_i)} \cdot \frac{\mathrm{Rev}(M^{\mathrm{sp}}(\frac{\bar{v}}{2}),\mathbb{F}_i)}{\mathrm{Rev}^\star(\mathbb{F}_i)} \geq \frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)},$$

where the inequality follows from that the first fraction at the left-hands side of this inequality is bounded from below by $\frac{\hat{v}-\varepsilon}{\hat{v}+\varepsilon}$, as shown earlier, and the second fraction is bounded from below by $\frac{1}{2}$ by Theorem 2.8. This inequality establishes a lower bound on the competitive ratio, indicating the fraction of revenue that can be achieved by the seller using mechanism $M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2})$, compared to the revenue achieved by the optimal revenue under distribution $\mathbb{F}_i$.

By the derivations above, we have $\gamma\big(M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2}),\mathbb{F}_i\big) \geq \frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)}$ for all $i \in \mathcal{N}$. From Theorem 2.11, $\gamma\big(\hat{M}^{\mathrm{sp}}(\boldsymbol{r}^\star),\mathbb{F}\big) \geq \frac{1}{2}$ when bidders' values are independent and not necessarily identically distributed, and $\gamma\big(\hat{M}^{\mathrm{sp}}(\boldsymbol{r}^\star),\mathbb{F}\big) = 1$ when bidders' values are i.i.d. By Theorem 2.12, we can now conclude that under $\mathbb{F} \in \mathscr{G}(\mathscr{V}^N)$, we have $\gamma\big(M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2}),\mathbb{F}\big) \geq \frac{1}{2}\frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)} = \frac{\hat{v}-\varepsilon}{4(\hat{v}+\varepsilon)}$ for independent but not necessarily identical bidders' values, meaning $M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2})$ guarantees at least a $\frac{4(\hat{v}+\varepsilon)}{\hat{v}-\varepsilon}$-approximation. For the case of i.i.d. bidders' values, similarly we have $\gamma\big(M^{\mathrm{sp}}(\frac{\hat{v}-\varepsilon}{2}),\mathbb{F}\big) \geq 1\frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)} = \frac{\hat{v}-\varepsilon}{2(\hat{v}+\varepsilon)}$. The claim thus follows. □

Chapter 3

Capacity Allocation: Balancing Contract Commitments and Spot Market Opportunities

3.1 Introduction

Before modern contract practices in freight shipping and logistics, capacity (amount of freight a carrier can carry) was allocated through informal agreements and long-term partnerships. As supply chains became more structured and competitive, firms began using allotment contracts more formally, specifying both the price and the capacity offered. Today, effective capacity allocation is a key challenge in supply chain management. Firms decide how to allocate their limited capacity (typically a long-term decision that cannot be changed in a short term) across different customers, products, or time periods, often before knowing the exact demand. Allotment contracts guarantee a certain amount of capacity at a predetermined price for the contracted customers. These contracts ensure steady income and planned usage, but they make it harder for the company to adjust if better opportunities arise later. Improper capacity allocation can have serious consequences. On one hand, allocating too much capacity to low-value contracts can lead to missed opportunities in the spot market, where prices may be higher. On the other hand, allocating too little can damage customer relationships or result in unused capacity and reduced revenue.

This work is carried out in collaboration with a major European ferry operator. We study a capacity allocation problem faced by this maritime asset provider operating short-sea liner services (e.g., ferries, Roll-on/Roll-off (RoRo) vessels, and feeders) that transport cargo carried in trucks, trailers, and containers. Capacity is measured in lane meters, and sailings are scheduled weekly (the company has a fixed number of sailings each week with a known capacity), typically ranging from 1,800 to 4,600 lane meters. Customers are freight forwarders or cargo owners who use contracts or book in the spot market. The allocation process unfolds in two phases: capacity is first offered through long-term contracts, and remaining capacity is later sold on the spot market. Contracted customers receive agreed-upon capacity but pay only for capacity used at the time of departure. Demand from these customers is uncertain and only realized at departure, exposing the asset provider to lost sales as unused capacity generates no revenue. If a contracted customer's demand exceeds their allocated capacity, they may purchase additional capacity at the contract price, provided that there is available capacity in the spot market and their contract price is above the spot market price. The spot market price is dynamic and updated based on factors such as remaining capacity and time until departure.

Deciding how much capacity to allocate among contracted customers and how much to reserve for the spot market is a tactical decision that directly impacts total revenue, which is the sum of revenue from contracts and the spot market. While revenue from contracts is determined by individual contracted prices, customer demands, and capacity allotments, spot market revenue is relatively more difficult to understand (and formulate) because it depends on several dynamic factors, including demand, remaining capacity, and the dynamics of price updates. Treating these drivers separately may result in non-implementable allotments or even an intractable model that cannot generate allocation decisions. A promising direction, which we choose to follow, is to develop a model that represents spot market revenue at an aggregate level as a function of the remaining capacity allocated to the spot market. This formulation is motivated by the commonly used bid-price policy (see Talluri and Van Ryzin (1998) for details of the policy) in managing spot market sales, where firms accept spot demand only if the offered price exceeds a threshold reflecting the opportunity cost of capacity. Importantly, this bid price is inversely related to the available spot capacity: as capacity for the spot market decreases, the bid price rises, and as capacity increases, the bid price falls. This inverse relationship naturally induces diminishing marginal revenue with respect to spot capacity, since each additional unit allocated to the spot market is either sold at a lower price or left unsold. A similar structure of diminishing marginal revenue for the scatter market (which corresponds to the spot market in our case) is described in Araman and Popescu (2010).

In this chapter, we study a capacity allocation problem in which a firm must decide how to allocate its limited capacity to the contracted customers and the spot market in the presence of demand uncertainty. The objective is to maximize expected revenue. It is important to note that customer contracts are negotiated annually. The key components of these contracts—price and minimum capacity allotments—are determined based on the previous year's performance and are market-specific. These strategic commercial negotiations are beyond the scope of this study. While our focus is on liner shipping, the proposed solution approach can be extended to a broader class of revenue management problems that arise in various industries. These include cruise lines, hotels, airlines (when upgrades are restricted), and any setting where capacity can be partitioned among customer segments without substitution across segments.

The main contributions of this chapter can be summarized as follows.

- We formulate the maritime asset provider's capacity allocation problem as an integer program with integer-valued capacity allocation decisions, aiming to maximize a concave revenue function subject to affine constraints on the capacity decisions.

- To gain insights into the optimal allocations, we relax the integrality constraints and derive the optimality conditions for the relaxed problem. Specifically, we show that the optimal allocations equalize the marginal revenue (roughly, the rate of change that would result from slightly perturbing the allocated capacity) from each contracted customer, all of which must be greater than or equal to the marginal revenue from the spot market.

- Leveraging the insights from the optimality conditions, we propose a computationally efficient algorithm, Expected Marginal Revenue with a Spot market (EMR-S), to solve the original integer program. The EMR-S algorithm iteratively allocates capacity, one unit at a time, to the contracted customer with the highest current marginal revenue, as long as it exceeds that of the spot market. We prove that EMR-S is exact in the sense that it outputs an optimal solution. We also conduct numerical experiments to compare the computational efficiency of EMR-S with directly solving the integer program using Gurobi, which reveal that EMR-S is significantly faster.

- Through computational experiments on real data, we demonstrate that our method is suitable for the partner company's operations and results in revenue improvements ranging from 13 to 55 percent across different sailings.

3.1.1 Structure of the chapter

The remainder of the chapter is organized as follows. Section 3.2 reviews related work and positions our contribution within the literature streams of capacity allocation and revenue management. Section 3.3 defines the problem and introduces the necessary preliminaries. Section 3.4 derives structural insights, studying a convex relaxation of the original problem. Section 3.5 presents the EMR-S algorithm, establishes its optimality, and assesses its computational efficiency. Section 3.6 presents the experiments with real data. Section 3.7 concludes the chapter.

3.1.2 Notation

Random vectors are designated by tilde signs (e.g., $\tilde{\boldsymbol{d}}$), and their realizations are denoted by the same symbols without tildes (e.g., $\boldsymbol{d}$). We write $\mathbf{1}(E)$ to denote the indicator function of a logical expression E, where $\mathbf{1}(E)$ equals 1 if E is true, and 0 otherwise.

3.2 Literature Review

Capacity management with contracts is an active area of research in freight transportation. Most carriers sell a significant portion of their capacity through long-term contracts to agencies (commonly referred to as forwarders). This contracted capacity, known as allotments, was first studied by Kasilingam (1997), who investigated the differences between air cargo and passenger revenue management across various aspects such as forecasting, overbooking, and seat allotment. There is a substantial body of literature on capacity allocation problems in revenue management. For comprehensive reviews, we refer the reader to Talluri and Van Ryzin (2006), and Klein et al. (2020). Most studies in this area have primarily focused on seat allocation and inventory control (Belobaba, 1987, 1992; Lee and Hersh, 1993; Tavana and Weatherford, 2017), pricing and overbooking control (Kunnumkal and Topaloglu, 2011), and booking control (Amaruchkul et al., 2007; Huang and Chang, 2010; Levin et al., 2012; Moussawi-Haidar, 2014).

This chapter is mainly related to capacity inventory control where the basic problem is to allocate a capacity for each fare class (or each customer with a distinct price) to maximize the total expected revenue. Littlewood (1972) took the first step by analyzing this problem in the context of a single-segment transport system. They proposed a decision rule for a two-fare-bucket scenario: accept a request for a discount-fare booking if its revenue exceeds the expected revenue from potential future full-fare bookings. This decision rule, now known as Littlewood's Rule, yields an optimal allocation when there are only two fare buckets, the demand for each bucket is independent, and arrival occurs in order from low to high fare classes. Building on this foundation, Belobaba (1987) introduced a more general heuristic for setting seat protection levels and booking limits across more than two nested fare classes. This approach, known as the Expected Marginal Seat Revenue (EMSR) model, became a cornerstone of revenue management in transportation. Belobaba (1992) introduced a modified EMSR method incorporating a protection model, which better approximates the optimality conditions for nested booking classes and after this work, the new version was named EMSRb, while the original version was renamed as EMSRa. In the context of risk aversion in capacity control, a heuristic approach suggested by Weatherford (2004) extends EMSR to a concept called EMSU (Expected Marginal Seat Utility). In this extension, the price of a seat is replaced by the utility (derived as a function of price).

However, these classical models assume fare buckets are nested, products are undifferentiated, and customers always book the lowest available fare regardless of their true willingness to pay, thus allowing for substitution between fare classes. In contrast, when fare buckets are fully differentiated and substitution is not permitted (e.g., in scenarios where products are partitioned and customers cannot be freely upgraded), the standard EMSR models are no longer adequate. To address this, Tavana (2004) develop a version of the EMSR model, called EMSRc, tailored for capacity allocation under fully differentiated fare environment. Tavana and Weatherford (2017) reintroduce the method in nested fare structures and show that it performs better incomparison to EMSRa and EMSRb in an unrestricted fare environment. A possible extension to this setting is the incorporation of spot market. This chapter addresses the management of capacity across both contracts and the spot market. The allocation policy we adopt for distributing capacity between contract commitments and spot market is conceptually similar to the EMSRc framework. In particular, we compare the marginal revenue associated with different customer segments distinguished by price (or fare buckets) and the marginal revenue of allocating capacity to the spot market.

There are studies that explore the coordination between allotment contracts and spot market booking decisions. Levin et al. (2012) and Moussawi-Haidar (2014) study booking control policies for the spot market aimed at maximizing the combined revenue from allotments and the expected profit from spot market sales and propose a dynamic control model that makes real-time decisions on whether to accept or reject spot booking requests. Various machine learning enhanced algorithms have been developed for airline seat inventory control problems (e.g., Gosavi et al. (2002); Lawhead and Gosavi (2019)). However, these contexts exhibit characteristics that differ from freight booking control problems. Dumouchelle et al. (2024), using a fast approximation of the value of a cargo container, developed a reinforcement learning approach to make accept/reject decisions on booking offers in order to maximize profit. These papers address capacity control problems with only acceptance or rejection decisions whereas in our problem the decision variables directly represent the allocation in each fare bucket.

Besides freight shipping companies, there are several industries in which capacity is sold partly in a primary (advance purchase) market and partly in a secondary (spot) market, such as manufacturing, food, and broadcasting media. The model used in these cases is a variation of the newsvendor model. Gillai and Lee (2009) study how secondary markets affect retailers' strategies for disposing of excess inventory. In a two-period model with a fixed primary market price and an endogenous secondary market price, retailers choose order quantities strategically. Marginal revenue is assumed to be a decreasing linear function of excess inventory. Yang and Yu (2025) consider the newsvendor setting for the perishable food. By modeling consumers' ordering decisions, they study a novel clearance scheme in which the store sells surprise bags containing an uncertain quantity of product units. Araman and Popescu (2010) study the ad allocation problem for more traditional media, specifically broadcasting. Their model is concerned with how to allocate limited advertising space between up-front contracts and the so-called scatter market (i.e., a spot market). One feature that distinguishes our freight capacity allocation from the above literature is the integer programming formulation under demand uncertainty.

In our chapter, we consider a spot market revenue function that is concave in the capacity reserved for the spot market. The concavity of the spot revenue function captures the diminishing marginal revenue of capacity. There has been little work that considers diminishing marginal revenue for capacity allocated outside of allotments and sold at a market-clearing price. Cachon and Kök (2007) endogenize the clearance price and demonstrate that several intuitive estimation procedures of salvage value lead to significantly poor performance. Araman and Popescu (2010) considers a similar setup with multiple clients under contracts, where the aggregate revenue from the spot market is concave in the remaining capacity. Their model focuses on a linear function for spot market revenue, whereas we consider a general concave function. Although the spot market in our work is modeled at an aggregate level, the setups are different, and our results are new in the context of freight contracts.

To summarize, to the best of our knowledge, none of the existing studies integrates optimal capacity allocation in the presence of both differentiated fare buckets and a spot market, and this chapter aims to fill that gap. It is important to note that we view the allocation problem as a tactical issue where customer prices are fixed and known. Unlike the traditional revenue management setting, operational-level decisions in maritime freight shipping—such as freight rate pricing in liner container services—are highly sensitive. Contracted customers hold strong bargaining power, which is addressed through strategic decisions, such as offering lower prices to customers with higher demand, in a competitive environment to maintain long-term relationships. For a review of dynamic pricing, which involves adjusting product prices over time, we refer readers to Du et al. (2025).

3.3 Problem Formulation and Preliminaries

An asset provider who is endowed with c units of capacity faces the challenge of how to allocate these units through allotment contracts or on the spot market. All potential allotment contracts are the result of a negotiation process. As a result of this process, each potential allotment customer and the asset provider agree on a contracted price per unit of lane meters. The contracted price depends on the specific practices of the asset providers and the type of customer. For example, customers with high demand typically negotiate a lower rate in their allotment contracts. These contracted prices are exogenous in our model. The capacity allocated to contracted customers is determined at the beginning of the capacity allocation phase, whereas booking requests on the spot market occur continuously until the time of departure. We assume that there are N contracted customers indexed by $i \in \mathcal{N} = \{1, \ldots, N\}$ and denote by $p_i \in \mathbb{R}_+$ the contracted price of customer $i \in \mathcal{N}$. For each customer $i \in \mathcal{N}$, we assume that their demand $\tilde{d}_i$ is a random variable and that the random vector $\tilde{\boldsymbol{d}} = (\tilde{d}_1, \tilde{d}_2, \ldots, \tilde{d}_N)$ of the demands of all contracted customers follows a joint probability $\mathbb{P}$. Since demand is an integer, $\tilde{d}_i$ can only take non-negative integer values. We denote by $\mathbb{F}_i$ the marginal cumulative distribution function of $\tilde{d}_i$, that is, $\mathbb{F}_i(z) := \mathbb{P}(\tilde{d}_i \leq z)$ for all $z \in \mathbb{Z}_+$.

The asset provider should determine the units b_i to allocate to each contracted customer $i \in \mathcal{N}$ without knowing their exact demand, relying only on their distribution and given contracted prices. We denote by $\boldsymbol{b} = (b_1, b_2, \ldots, b_N)$ the vector of allocations to contracted customers. All units not allocated to contracted customers become available on the spot market. Any excess capacity $\max\{b_i - d_i, 0\}$ allocated to customer $i \in \mathcal{N}$ cannot be allocated to the spot market, as it was reserved for the contracted customer until after spot market sales. We reduce the complex spot market dynamics to a single function and model the expected revenue from the spot market using a concave, non-decreasing function $\mathrm{R_s}\left(c - \sum_{i=1}^{N} b_i\right)$ of the remaining capacity[1]. This reflects the practical situation that marginal revenue from the spot market decreases as more capacity is allocated to it. In practice, increasing spot market capacity often requires price reductions to stimulate demand and attract lower-paying, non-contracted customers.

At the time of departure, the demands are realized. If the demand d_i of customer i does not exceed the allocation b_i, the customer ships all demand d_i at price p_i. Otherwise, the customer ships b_i at the same price. We formulate the asset provider's capacity allocation problem with the objective of maximizing expected revenue as follows:

$$\max_{\boldsymbol{b} \in \mathscr{B}} \quad \mathbb{E}\left[\sum_{i=1}^{N} p_i \min\{\tilde{d}_i, b_i\}\right] + \mathrm{R_s}\left(c - \sum_{i=1}^{N} b_i\right), \tag{$\mathscr{P}$}$$

where

$$\mathscr{B} = \left\{\boldsymbol{b} \in \mathbb{Z}_+^N : \sum_{i=1}^{N} b_i \leq c, \ \underline{b}_i \leq b_i \ \ \forall i \in \mathcal{N}\right\},$$

and $\underline{b}_i \geq 0$ denotes the minimum capacity to be allocated to customer $i \in \mathcal{N}$, which may be agreed upon in the contract or set strictly greater than zero to strengthen customer relationships by balancing profitability with customer retention. In the reminder, we assume that $\sum_{i=1}^{N} \underline{b}_i \leq c$ because otherwise the problem is infeasible.

[1] Our choice to model the expected revenue from the spot market as a function of the remaining capacity is also motivated by the real data provided by our partner company. In particular, because the data is limited, making reliable inferences about the spot market based on multiple factors is challenging and risks overfitting. Our real-data experiments (to be presented in Section 3.6) indicate that remaining capacity alone provides a good enough explanation of spot market revenue. Furthermore, this approach enables us to propose a computationally efficient exact algorithm that generates an optimal output, making it advantageous in terms of tractability.

3.4 Structural Properties of Optimal Allocations

For the purposes of the analysis in this section, we temporarily relax the integrality requirement in $\mathscr{B}$ and assume that b_i , $i \in \mathscr{N}$, can take continuous values. Additionally, we temporarily assume that $\mathbb{F}_i$ has a density function (and therefore $\tilde{d}_i$ can take continuous values) for all $i \in \mathscr{N}$ and that R_s is differentiable. We characterize the structural properties of the optimal allocations for the corresponding convex relaxation of problem ($\mathscr{P}$), denoted by ($\mathscr{P}_{\text{conv}}$), in order to gain insights. In Section 3.5, we will leverage these insights in developing an exact algorithm to solve the original problem ($\mathscr{P}$). We note that the temporary assumptions of this section will not be required in Section 3.5.

Theorem 3.1. *Assume, without loss of generality, that the contracted customers are reindexed such that* $p_1\,(1-\mathbb{F}_1(\underline{b}_1)) \geq p_2\,(1-\mathbb{F}_2(\underline{b}_2)) \geq \cdots \geq p_N\,(1-\mathbb{F}_N(\underline{b}_N))$. *For any optimal solution* $\boldsymbol{b}^\star$ *to* ($\mathscr{P}_{conv}$)*, there exists* $K \in \{0,1,\ldots,N\}$ *such that* $\boldsymbol{b}^\star$ *satisfies*

$$\begin{aligned}
&p_1\,(1-\mathbb{F}_1(b_1^\star)) = p_2\,(1-\mathbb{F}_2(b_2^\star)) = \cdots = p_K\,(1-\mathbb{F}_K(b_K^\star)) \geq \mathrm{R}_s'\left(c-\sum_{i=1}^{N} b_i^\star\right), \\
&b_i^\star = \underline{b}_i \;\; \forall i \in \{K+1,\ldots,N\}.
\end{aligned} \tag{3.1}$$

The equality condition across $p_i(1-F_i(b_i^\star))$ holds only for the first K customers who receive a strictly positive allocation beyond their minimum. When $K=0$, it means that no customer receives a strictly positive allocation beyond their minimum.

Remark 3.2. To build intuition about the structure of the optimal allocations, note that $p_i\,(1-\mathbb{F}_i(b_i))$ can be interpreted as the rate of change in expected revenue from a small perturbation in the allocation to customer i, given that customer i has already been allocated b_i units. This quantity is non-increasing in b_i. Similarly, $\mathrm{R}_\mathrm{s}'\big(c-\sum_{i=1}^{N} b_i\big)$ represents the rate of change in expected revenue from a small perturbation in the spot-market allocation, given that its current allocation is $c-\sum_{i=1}^{N} b_i$. Since the spot revenue function R_s is concave, its derivative R_s' is non-increasing. These monotonicity properties suggest a natural way to think about the allocation decisions: allocate capacity one by one, each time assigning a unit capacity to the option that yields the highest rate of revenue increase, that is, either to the contracted customer with the highest current rate, or to the spot market if its rate is higher than that of any customer. We formalize this intuition through an algorithm presented in Section 3.5, and we prove its optimality in the same section. From now on, we refer to the rates of change outlined in this remark as marginal expected revenues.

3.5 Exact Algorithm

Even though the findings of Section 3.4 were derived under the assumption that allocations can take continuous values as well as temporary assumptions on $\mathbb{F}_i$, $i \in \mathscr{N}$, and R_s, they offer valuable insights for developing an algorithm to solve the original problem ($\mathscr{P}$). Building on the ideas discussed in Remark 3.2, we now develop the Expected Marginal Revenue with Spot (EMR-S) algorithm to solve problem ($\mathscr{P}$). Our proposed algorithm first assigns the minimum required allocations to the contracted customers. It then assigns the remaining units one by one to the option with the highest current marginal expected revenue. Once the marginal expected gain from the spot market exceeds that of any contracted customer, the algorithm allocates all remaining capacity to the spot market and terminates. The complete details of EMR-S are presented in Algorithm 1.

Algorithm 1 Expected Marginal Revenue with Spot (EMR-S) Algorithm

Inputs: c, $\underline{b}_i$, $\mathbb{F}_i$ for all $i \in \mathscr{N}$
Initialize $b_i \leftarrow \underline{b}_i$ for all $i \in \mathscr{N}$ and $c_{\text{rem}} \leftarrow \left(c - \sum_{i=1}^{N} \underline{b}_i\right)$
while $c_{\text{rem}} > 0$ **do**
    for $i \in \mathscr{N}$ **do**
        Compute marginal revenue $\text{MR}_i \leftarrow p_i(1 - \mathbb{F}_i(b_i))$
    end for
    Compute marginal spot revenue $\Delta \mathrm{R}_s \leftarrow \mathrm{R}_s(c_{\text{rem}}) - \mathrm{R}_s(c_{\text{rem}} - 1)$
    Find customer $i^\star \leftarrow \min \arg\max_{i \in \mathscr{N}} \text{MR}_i$
    if $\text{MR}_{i^\star} \geq \Delta \mathrm{R}_s$ **then**
        Allocate a unit to customer $i^\star$: $b_{i^\star} \leftarrow b_{i^\star} + 1$
        $c_{\text{rem}} \leftarrow c_{\text{rem}} - 1$
    else
        Allocate c_{rem} to spot market
        $c_{\text{rem}} \leftarrow 0$
    end if
end while
Outputs: b_i for all $i \in \mathscr{N}$

The algorithm terminates in at most $\left(c - \sum_{i=1}^{N} \underline{b}_i\right) \leq c$ iterations. In each iteration, it computes marginal revenues for all N customers and selects the customer with highest marginal revenue, both of which require $O(N)$ operations. Therefore, the time complexity of the algorithm is $O(cN)$.

Next, we show that Algorithm 1 is exact in the sense that it gives an optimal solution to the original capacity allocation problem ($\mathscr{P}$). The following lemma will be useful in proving this exactness result.

Lemma 3.3. *Problem ($\mathscr{P}$) is equivalent to*

$$\max_{\boldsymbol{b}\in\mathscr{B}} \sum_{i=1}^{N} p_i \sum_{t=1}^{b_i} \mathbb{P}(\tilde{d}_i \geq t) + \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i\right).$$

Using the equivalent formulation of ($\mathscr{P}$) in Lemma 3.3, we prove that no other feasible allocation can offer more revenue than the allocations produced by EMR-S.

Theorem 3.4. *EMR-S outputs an optimal solution to problem ($\mathscr{P}$).*

3.5.1 Computational Efficiency of Algorithm

We study the computational efficiency of EMR-S compared to solving the integer programming problem ($\mathscr{P}$) directly using Gurobi. To this end, we consider a problem instance with $N = 100$ customers. For each customer, we construct a demand distribution by sampling 1000 values from a scaled Beta distribution with probability density function

$$f^{\alpha,\beta}(x) = \frac{1}{1000}\frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)} \quad x \in [0,1000],$$

where $B(\alpha,\beta)$ is the Beta function with shape parameters α and β, and the support lies between 0 and 1000. After sampling from this continuous distribution, we round the realizations to the nearest integer to have integer-valued demand. For each customer, the parameters α and β are integers drawn independently from the uniform distribution on $\{1,2,\ldots,9\}$. Similarly, the contract price for each customer is independently sampled from the uniform distribution on $[1,10]$. The expected revenue from the spot market is modeled using a logarithmic function:

$$\mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i\right) = 10^5 \log\left(1 + c - \sum_{i=1}^{N} b_i\right).$$

We assume $\underline{b}_i = 0$ for all $i \in \mathscr{N}$.

We evaluate computation time for EMR-S and Gurobi solver for different capacity values ranging from 10,000 to 90,000. The results, reported in Figure 3.1, show that EMR-S is significantly faster than the exact solution via Gurobi, particularly at large capacity levels.

All experiments were conducted on a high-performance computing server running Ubuntu Linux (kernel version `6.8.0-62-generic`, `x86_64` architecture). The machine is equipped with two Intel® Xeon® Gold 6330 CPUs, each with 28 cores and 2 threads per core, totaling 112 logical processors, and has 251 GB of RAM and 255 GB of swap memory. All code was executed in a 64-bit environment using Python 3.12.3. The integer programming problem ($\mathscr{P}$) was solved via

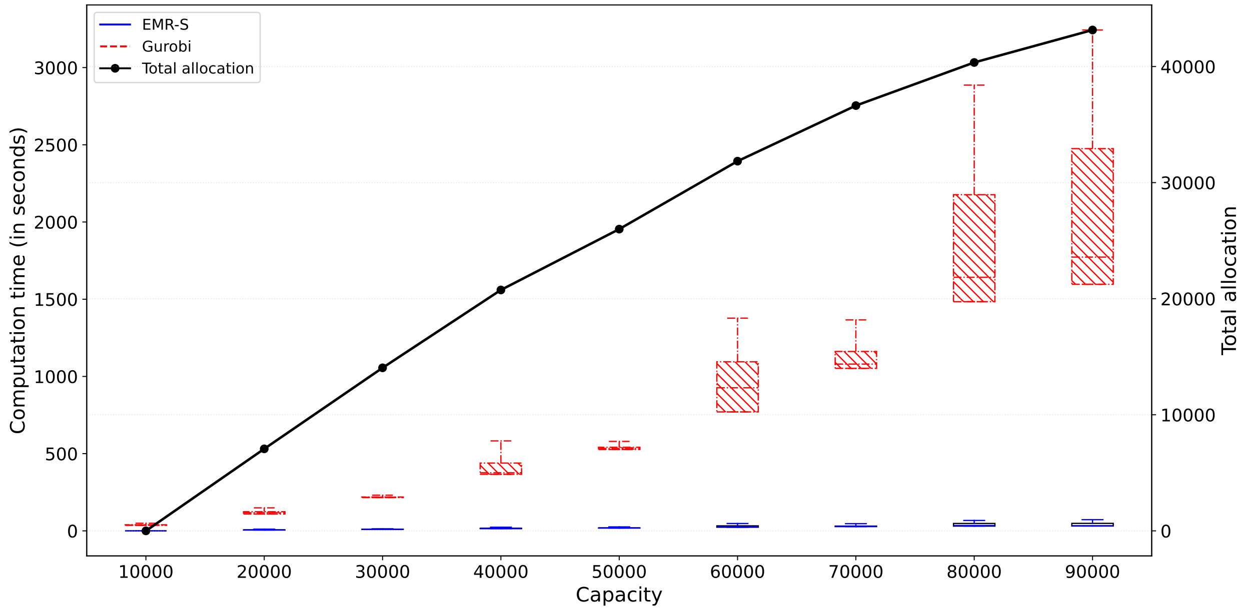

Figure 3.1: Comparison of the computation time statistics: EMR-S vs Gurobi. Total allocation is the sum of optimal allocations across all contracted customers

the Gurobi 12.0.0 optimizer using its Python interface, while EMR-S was implemented in native Python.

As the total capacity c increases, the optimal allocations naturally increase, since more units can be assigned to contracted customers, which may generate additional revenue. Optimal allocations also ensure that the capacity reserved for the spot market generates sufficiently high marginal revenue; otherwise, allocating too much capacity to customers may leave the spot market with very little or no opportunity to generate revenue. To solve this problem with the integer allocation variables b_i, Gurobi explores the search space of feasible values[2] that expands in each b_i as c grows, making the problem increasingly difficult to prune. In contrast, the EMR-S algorithm proceeds by allocating units one by one and stops once the marginal revenue from the spot market exceeds that of all customers, thereby achieving faster computation time.

3.6 Practical Implementation of Allocation Policy

We next showcase the performance of our proposed algorithm on real-world data for allocating freight capacity to contracted customers of a maritime RORO (Roll-on/Roll-off) company. The company operates in both freight and travel segments; however, our study focuses exclusively on the freight operations. Freight capacity is sold to customers either through contracts or on the spot

[2]For example, by branching on decisions such as $b_i \leq k$ versus $b_i \geq k+1$, thereby forming a branch-and-bound tree.

market. The customer base primarily consists of logistics service providers, freight forwarders, and shippers. The dataset contains demand information from anonymized customers, including 405 contracted and 1920 non-contracted customers, over a three-year period from 1st January 2022 to 31st December 2024. The data spans inbound and outbound departures on two key routes (four markets in total), comprising 9724 sailings across all markets. Capacity is measured in lane meters, with freight typically transported in the form of trucks that occupy either 14 or 17 lane meters.

In practice, the company operates with standardized sailings, but it divides (under its direct control) the total sailing capacity between travel and freight. The decision on how much capacity to assign to freight depends on several factors, including the time of year, holiday seasons, and the day and time of the week. For example, during the summer holidays in July and August, only a small share of the sailing capacity is assigned to freight. Thus, sailings often have varying capacities for freight. In the remainder of this chapter we use the term sailing capacity to denote the capacity allocated to freight.

Contracted and non-contracted customers may purchase capacity on both routes and across different markets. We treat the demand on each market of the routes independently and refer to each leg as a separate market. This is motivated by the fact that demand is unidirectional and does not depend on the flow in the opposite direction. For each contracted customer, contracts are annual and specify the allocated capacity on a per-sailing basis throughout the contract year. The company monitors actual usage to make allocation decisions for the next contract year.

We use all customer data from the 1st January 2022 to 31st December 2023 as our training set to learn the distributions of customers and to estimate spot market marginal revenue function. Each entry in the dataset represents a unique demand request identified by a booking number, the route and market on which demand was requested, the associated demand per vehicle (in lane meters), price (in Euros per lane meters), type of vehicle (from small enclosed vans to large commercial trucks), whether the request was part of an allocated capacity or excess capacity, and the end state of the booking. The end state indicates whether the booking was shipped or left unused.[3] We consider shipped bookings as actual demand, for which customers pay the carrier.

Customers can create multiple booking requests, which are classified either as falling within the customer's allocated (contracted) capacity or as additional bookings beyond the allocated capacity. The company allows contracted customers to satisfy excess demand by drawing from capacity originally reserved for the spot market. While the company has the option to reject requests that exceed the allocated capacity, such rejections are not observed in the data. Because past allocation decisions have often resulted in underutilized sailing capacity, the company has often accepted

[3] As a customer may adjust a booking over time (e.g., by assigning a vehicle with a different booking number or requesting a different departure date), the database contains additional entries to capture these changes. The final entry of each booking reflects the true demand from the customer and is therefore the focus of our analysis.

additional requests. When the additional requests are accepted, customers pay the contracted price. The acceptance decision follows a bid price mechanism: a request is accepted only if the contracted price exceeds the bid price.

To motivate our modeling approach, we first examine the sailing capacities across markets. Sailing capacities vary considerably in size across different routes, driven by the heterogeneous nature of demand and operational constraints. To illustrate these differences, Table 3.1 presents key summary statistics for all sailings, as well as for the train and test subsets.

Entire Data (2022-2024)								
Market	Total	Mean	Std Dev	Min	P25	P50	P75	Max
R1L1	2172	4227.24	372.75	1800.0	4125.0	4222.5	4485.0	4785.0
R1L2	2169	4341.39	397.88	1800.0	4175.0	4485.0	4485.0	4835.0
R2L1	2693	2007.99	719.77	414.0	1419.0	1804.0	2955.0	3057.0
R2L2	2690	2022.78	700.51	205.0	1500.0	1846.0	2955.0	3057.0

Train Data (2022-2023)								
Market	Total	Mean	Std Dev	Min	P25	P50	P75	Max
R1L1	1448	4288.92	238.81	3200.0	4125.0	4325.0	4485.0	4785.0
R1L2	1445	4424.02	239.55	3200.0	4320.0	4485.0	4485.0	4835.0
R2L1	1842	1880.86	703.31	414.0	1419.0	1657.0	2649.0	2957.0
R2L2	1839	1898.60	691.87	205.0	1449.0	1759.0	2649.0	2955.0

Test Data (2024)								
Market	Total	Mean	Std Dev	Min	P25	P50	P75	Max
R1L1	724	4103.9	529.3	1800.0	4125.0	4125.0	4325.0	4730.0
R1L2	724	4176.49	564.99	1800.0	4125.0	4125.0	4485.0	4800.0
R2L1	851	2283.18	677.19	494.0	1640.0	2411.0	2955.0	3057.0
R2L2	851	2290.41	642.07	334.0	1725.0	2207.0	2955.0	3057.0

Table 3.1: Summary statistics of sailing capacities per route for all data, training data, and test data. 'Total' represents the total number of sailings operated in the corresponding market, 'Mean' represents the average capacity available, and 'Std Dev' represents the standard deviation. 'Min' and 'Max' represent the minimum and maximum recorded sailing capacities, respectively, on the market. 'P25', 'P50', and 'P75' represent the 25th, 50th (Median), and 75th percentiles, respectively.

Since the marginal spot revenue function depends on the capacity allocated to the spot market, it also systematically depends on sailing capacity. Estimating a single revenue function across all sailings could therefore introduce bias. To counter this effect of sailing capacity, we group sailings of comparable sizes. We analyze the groupings of the various markets separately. The objective is to establish groups that contain a sufficient number of observations that are also similar in their characteristics, thereby limiting within-group variation. In R1L1 and R1L2, the training data are relatively tightly clustered, with 50% of the observations falling within the ranges of 4125–4485 and 4320–4485, respectively, suggesting that further grouping may offer limited benefits. In contrast, R2L1 and R2L2 exhibit considerably greater dispersion, with 50% of the observations spread over a much wider range (1419–2649 and 1449–2649). Therefore, we adopt a size-based division that ensures that each group contains a sufficient number of observations. Considering ranges of 800 units yields three potential intervals: 600–1400, 1400–2200, and 2200–3000. This division produces the observations reported in Table 3.2. As there are 11 observations in R2L1 and 13 in R2L2 below the first threshold, we include them in the first group. A finer partitioning would create additional groups but result in insufficient data for estimation. For R1L1 and R1L2, sailings are treated as single homogeneous groups (3000–5000), since vessel sizes in these routes are relatively stable. This yields a total of eight market groups: R1L1, R1L2, and three subgroups each from R2L1 and R2L2. We denote the set of market groups as

$$\mathscr{M} = \begin{Bmatrix} \text{R1L1},\ \text{R1L2},\ \text{R2L1}_{0\text{–}1400},\ \text{R2L1}_{1400\text{–}2200},\ \text{R2L1}_{2200\text{–}3000}, \\ \text{R2L2}_{0\text{–}1400},\ \text{R2L2}_{1400\text{–}2200},\ \text{R2L2}_{2200\text{–}3000} \end{Bmatrix},$$

where the subscripts denote the vessel capacity range corresponding to each subgroup of R2L1 and R2L2.[4] The detailed summary of dataset can be found in Table 3.2.

For each of group $m \in \mathscr{M}$ we separately estimate demand of contracted customers and the marginal spot market revenue.

Estimation of Demand: The sailings are often underutilized, with an average utilization ratio—defined as average ratio of the occupied capacity at departure to the corresponding sailing capacity—of approximately 63%. Out of 9,724 sailings, only 983 achieved full utilization. Consequently, demand censoring (i.e., rejections based on bid prices) occurs in the lower tail of the customer distribution. Thus, the observed booking data provide a reasonable proxy for realized demand. For future practice, we recommend recording rejected requests to enable more accurate estimation of demand distributions. In the present study, demand distributions are inferred from historical bookings using a sample-average approximation approach.

[4] The test data for R1L1 and R1L2 include 33 sailings in each route with a size of less than 3000 lane meters. In the test data for R2L1 and R2L2, the only sailing exceeding 3000 lane meters has 3057 lane meters; hence, we include it in $\text{R2L1}_{2200\text{–}3000}$ and $\text{R2L2}_{2200\text{–}3000}$.

Market	Total sailings	Capacity groups				Contracted customers	Non-contracted customers
		0–1400	1400–2200	2200–3000	3000–5000		
R1L1	2172	– –	– –	– –	Train: 1448 Test: 724	274	883
R1L2	2169	– –	– –	– –	Train: 1445 Test: 724	271	853
R2L1	2693	Train: 462 Test: 143	Train: 856 Test: 195	Train: 534 Test: 503	– –	197	984
R2L2	2690	Train: 394 Test: 62	Train: 937 Test: 282	Train: 518 Test: 497	– –	195	864

Table 3.2: Summary of data

For the i^{th} customer, the empirical cumulative distribution function in group m, $\hat{F}_{i,m}(.)$ is given by:

$$\hat{F}_{i,m}(a) = \frac{\text{Number of demand observations of customer } i \text{ in group } m \text{ with value} \leq a}{\text{Total number of demand observations of customer } i \text{ in group } m}.$$

The total revenue consists of revenue from allocated capacity through contracts and revenue from non-allocated capacity. The EMR-S algorithm uses demand distributions to determine the initial allocation via contracts. The revenue generated from contracts for each sailing is calculated as the sum of the product of the contracted price and the minimum of the allocation and demand.

For revenue generated from the spot market, we consider: (i) excess demand fulfilled for contracted customers beyond their initial allocation, and (ii) accepted demand requests from non-contracted customers. Non-contracted customers are typically smaller clients—who are not allocated capacity in advance but have an associated price, i.e., a promised rate they agree to pay if their booking request is accepted. A non-contracted customer's request is accepted only if their associated price exceeds the current bid price. If accepted, the customer pays their associated price.

Since counterfactual outcomes are not observable, we cannot directly measure the revenue that the EMR-S algorithm would have generated from the capacity allocated to the spot market. In particular, the data does not reveal the spot revenue under the EMR-S allocation. Therefore, we evaluate the algorithm's performance using a semi-synthetic approach that is standard in observational data studies: we first estimate a spot revenue function on test data and then use this function to compute the expected spot market revenue corresponding to the capacity allocated to the spot market by EMR-S.

Estimation of Marginal Spot Revenue Function: Estimating the marginal spot revenue involves calculating the additional revenue generated per unit of spot capacity on a given sailing j within market group $m \in \mathscr{M}$. Our estimation approach considers the total number of customers in the spot market. Let there be K spot customers for sailing j in group m. Each spot customer $k \in \{1, \ldots, K\}$ has a fulfilled demand $d_{k,j,m}$ and an associated price p_k. The total spot revenue for sailing j in group m is then

$$\text{Spot Revenue}_{j,m} = \sum_{k=1}^{K} p_k \, d_{k,j,m}.$$

The spot capacity for sailing j in group m is the difference between the vessel's total capacity $c_{j,m}$ and the total contracted allocation:

$$\text{Spot Capacity}_{j,m} = c_{j,m} - \sum_{i=1}^{N} b_{i,j,m}.$$

Finally, the marginal spot revenue (denoted $MR_{j,m}$) for sailing j in group m is estimated as

$$MR_{j,m} = \frac{\text{Spot Revenue}_{j,m}}{\text{Spot Capacity}_{j,m}}.$$

Using the computed data points $\left(\text{SpotCapacity}_{j,m}, MR_{j,m}\right)$ for each sailing j in market m, we estimate a functional relationship that describes how marginal spot revenue varies with the amount of available spot capacity. Formally, for a given spot capacity x, we define the estimated marginal spot revenue function as

$$\Delta\hat{\text{R}}_s(x) : \mathbb{R}_+ \to \mathbb{R}_+,$$

which is estimated separately for each market m.

We compare different regression models, from simple linear, quadratic and logistic fit to nonlinear models such as random forests, and XGBoost. We apply 5-fold cross-validation on the training data.The dataset was partitioned into five equal-sized subsets; for each fold, four subsets were used for training and the remaining one for testing. Linear and quadratic fits were done using ordinary least squares and we estimate the parameters for the following functions:

$$\Delta\hat{\text{R}}_{\text{linear}}(x \mid \beta_0^L, \beta_1^L) = \beta_0^L + \beta_1^L x, \quad \text{and} \quad \Delta\hat{\text{R}}_{\text{quadratic}}(x \mid \beta_0^Q, \beta_1^Q, \beta_2^Q) = \beta_0^Q + \beta_1^Q x + \beta_2^Q x^2,$$

where x is the capacity allocated to the spot market, and β_0^L, β_1^L and $\beta_0^Q, \beta_1^Q, \beta_2^Q$ are parameters for the linear and quadratic functions. For logistic fit, we fitted a sigmoid curve using nonlinear curve fitting. In particular, we estimate the parameters L, k, u using the data for the following function:

$$\Delta\hat{\mathrm{R}}_{\mathrm{logistic}}(x \mid L, k, u) = \frac{L}{1 + e^{-k(x-u)}}.$$

Random Forest and XGBoost models were fine-tuned via grid search nested within each cross-validation fold. For Random Forest, we tuned the number of trees (`n_estimators`), the maximum tree depth (`max_depth`), and the minimum number of samples required to split a node (`min_samples_split`). The grid search explored `n_estimators` = {100, 200}, `max_depth` = {1, 2, 3, 5}, and `min_samples_split` = {2, 3, 5}. XGBoost was tuned similarly over the number of estimators, maximum depth, and learning rate (`learning_rate` = {0.01, 0.05, 0.1}); lower learning rates typically improve accuracy but require more trees. For each fold and model, the coefficient of determination (R^2) was computed on the test set. The R^2 values and the average R^2 score across the five folds for each model and each market group are reported in Table 3.3.

3.6.1 Numerical Results

Once we have both the empirical demand distributions and the marginal revenue function for the spot market, we move to evaluate the performance of the EMR-S allocation policy using the test data on each market group in $\mathscr{M}$. From Table 3.3, we find that logistic regression is a suitable choice since its R^2 values are comparable to those of random forest and XGBoost, while offering a simpler and more interpretable structure that captures diminishing marginal effects (see Figure 3.2). The estimated functional form is given by $\frac{L_{\mathrm{train}}}{1+e^{-k_{\mathrm{train}}(x-u_{\mathrm{train}})}}$, where x denotes the spot capacity, and $L_{\mathrm{train}}, k_{\mathrm{train}}, u_{\mathrm{train}}$ are parameters estimated from the training set. For each customer i, vessel j in group m in the test data, the EMR-S allocation $b_{i,j,m}^{\mathrm{EMR}}$ is determined based on the estimated marginal spot revenue function and the distribution of customers.

The total revenue under EMR-S consists of (i) the contracted revenue and (ii) the expected spot revenue. Since the realized demand $d_{i,j,m}^{\mathrm{test}}$ of contracted customer i on vessel j in group m is observable in the test data, the contracted revenue for a sailing j in test data is computed as $\sum_{i=1}^{N} p_i \min\{d_{i,j,m}^{\mathrm{test}}, b_{i,j,m}^{\mathrm{EMR}}\}$, i.e., the contract price multiplied by the minimum of the allocated contract capacity and the realized customer demand on that sailing. The expected spot revenue is computed by first estimating the spot revenue function from the test data, at the remaining capacity after contracts are allocated. This function is $\frac{L_{\mathrm{test}}}{1+e^{-k_{\mathrm{test}}(x-u_{\mathrm{test}})}}$, where $L_{\mathrm{test}}, k_{\mathrm{test}}, u_{\mathrm{test}}$ are parameters estimated from the test set. This marginal spot revenue is then multiplied to the total spot capacity $\left(c - \sum_{i=1}^{N} b_{i,j,m}^{\mathrm{EMR}}\right)$. Thus, the total estimated revenue is

$$\mathrm{Rev}_{j,m}^{\mathrm{EMR}} = \sum_{i=1}^{N} p_i \min\{d_{i,j,m}^{\mathrm{test}}, b_{i,j,m}^{\mathrm{EMR}}\} + \left(c - \sum_{i=1}^{N} b_{i,j,m}^{\mathrm{EMR}}\right) \frac{L_{\mathrm{test}}}{1 + e^{-k_{\mathrm{test}}\left(c - \sum_{i=1}^{N} b_{i,j,m}^{\mathrm{EMR}} - u_{\mathrm{test}}\right)}}.$$

Market Group		**Linear Reg.**	**Quadratic**	**Logistic Reg.**	**Random Forest**	**XGBoost**
R1L1	Set 1	0.6741	0.7620	0.7787	0.8062	0.8072
	Set 2	0.7014	0.7944	0.8102	0.8452	0.8467
	Set 3	0.7034	0.7673	0.7795	0.8306	0.8307
	Set 4	0.6866	0.7677	0.7803	0.8185	0.8164
	Set 5	0.6807	0.7786	0.7848	0.8299	0.8297
	Avg.	**0.6892**	**0.7740**	**0.7867**	**0.8261**	**0.8261**
R1L2	Set 1	0.6771	0.6809	0.6943	0.6954	0.6957
	Set 2	0.6018	0.6087	0.6184	0.6215	0.6294
	Set 3	0.6268	0.6278	0.6462	0.6547	0.6528
	Set 4	0.6050	0.6030	0.6174	0.6229	0.6241
	Set 5	0.5804	0.5874	0.6021	0.6181	0.6220
	Avg.	**0.6182**	**0.6216**	**0.6357**	**0.6425**	**0.6448**
$R2L1_{0-1400}$	Set 1	0.6700	0.6704	0.6680	0.6953	0.7094
	Set 2	0.5613	0.5615	0.5713	0.6218	0.6369
	Set 3	0.5807	0.5811	0.5733	0.5820	0.6153
	Set 4	0.6001	0.6004	0.5983	0.6078	0.6081
	Set 5	0.5995	0.5931	0.5986	0.6221	0.6421
	Avg.	**0.6023**	**0.6013**	**0.6019**	**0.6258**	**0.6423**
$R2L1_{1400-2200}$	Set 1	0.6059	0.6049	0.6230	0.6130	0.6152
	Set 2	0.5778	0.5772	0.5818	0.6102	0.6194
	Set 3	0.5818	0.5882	0.6146	0.6554	0.6588
	Set 4	0.6179	0.6232	0.6318	0.6519	0.6523
	Set 5	0.6232	0.6239	0.6472	0.6806	0.6774
	Avg.	**0.6013**	**0.6035**	**0.6197**	**0.6422**	**0.6446**
$R2L1_{2200-3000}$	Set 1	0.5496	0.5415	0.5549	0.5814	0.5816
	Set 2	0.7117	0.7127	0.7311	0.7429	0.7443
	Set 3	0.6478	0.6488	0.6755	0.6730	0.6874
	Set 4	0.7015	0.7019	0.7197	0.7279	0.7370
	Set 5	0.6393	0.6401	0.6690	0.6998	0.7083
	Avg.	**0.6500**	**0.6490**	**0.6700**	**0.6850**	**0.6917**
$R2L2_{0-1400}$	Set 1	0.6023	0.6040	0.6013	0.6095	0.6172
	Set 2	0.6970	0.7256	0.7302	0.6981	0.7088
	Set 3	0.6407	0.6781	0.6807	0.6759	0.6748
	Set 4	0.5781	0.5803	0.5723	0.5757	0.5707
	Set 5	0.6034	0.5985	0.6065	0.5743	0.5959
	Avg.	**0.6243**	**0.6373**	**0.6382**	**0.6267**	**0.6335**
$R2L2_{1400-2200}$	Set 1	0.5549	0.5429	0.5563	0.5712	0.6307
	Set 2	0.5555	0.5416	0.5489	0.5718	0.6305
	Set 3	0.5541	0.5561	0.5783	0.5822	0.6481
	Set 4	0.5471	0.5654	0.6078	0.6067	0.6423
	Set 5	0.5539	0.5632	0.6066	0.6059	0.6578
	Avg.	**0.5712**	**0.5696**	**0.5838**	**0.5939**	**0.5975**
$R2L2_{2200-3000}$	Set 1	0.6978	0.6985	0.7474	0.7615	0.7717
	Set 2	0.7412	0.7413	0.8110	0.8643	0.8778
	Set 3	0.6846	0.6811	0.6953	0.7100	0.7047
	Set 4	0.7558	0.7565	0.8170	0.8446	0.8620
	Set 5	0.6885	0.6813	0.7194	0.7654	0.7647
	Avg.	**0.7136**	**0.7117**	**0.7580**	**0.7892**	**0.7962**

Table 3.3: R^2 scores from cross-validation on training set

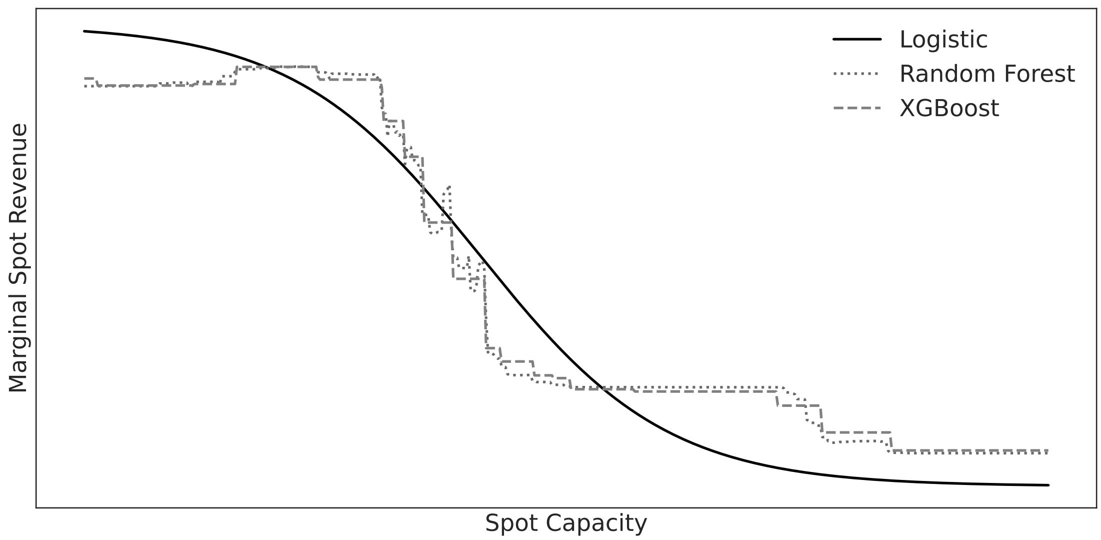

Figure 3.2: Structural comparison of estimated marginal revenue function from Logistic fit, Random Forest, and XGBoost on R1L2.

Note: Real data has been removed due to the non-disclosure agreement.

This procedure is repeated for all vessels in the test set. From the data, the observed revenue for sailing j in group m is denoted by $\text{Rev}_{j,m}^{\text{Observed}}$. We then calculate the aggregate revenue by summing over all sailings, and compute the ratio of the estimated revenue obtained under EMR-S to the actual revenue from historical allocations, given by:

$$\text{Revenue Ratio}_m = \frac{\sum_j \text{Rev}_{j,m}^{\text{EMR}}}{\sum_j \text{Rev}_{j,m}^{\text{Observed}}}.$$

This ratio for each market group is reported in Table 3.4.

Market Group	**Revenue Ratio**
R1L1	1.216
R1L2	1.133
R2L1$_{0-1400}$	1.557
R2L1$_{1400-2200}$	1.347
R2L1$_{2200-3000}$	1.383
R2L2$_{0-1400}$	1.165
R2L2$_{1400-2200}$	1.326
R2L2$_{2200-3000}$	1.437

Table 3.4: Ratio of estimated revenue under EMR-S to revenue under historical allocations across market groups

We compare the allocations made to contracted customers by EMR-S against the historical allocations. Figures 3.3 and 3.4 show the allocation by EMR-S often exceeds in total allocation to contracted customers across routes. In fact, the number of sailings where the ratio exceeds 1 is consistently higher than those at or below 1, R1L1 records 243 cases at or below 1 compared to 481 above; R1L2 has 321 versus 403; R2L1 shows 370 against 471; and R2L2 reports 394 compared to 447. More spcifically For R2L1 in the 0–1400 capacity group, the majority of instances fall within ratios exceeding 1, with 108 out of 143 cases where EMR-S allocates more than the company's heuristics. In the 1400–2200 and 2200–3000 groups, EMR-S provides higher allocations in 127 out of 195 and 258 out of 503 instances, respectively. For R2L2 in the 0–1400 group, the histogram shows a more balanced distribution, with 31 out of 62 instances exceeding a ratio of 1, indicating that EMR-S allocates more than the company's heuristics. In the 1400–2200 and 2200–3000 groups, EMR-S yields higher allocations in 108 out of 282 and 298 out of 497 cases, respectively.

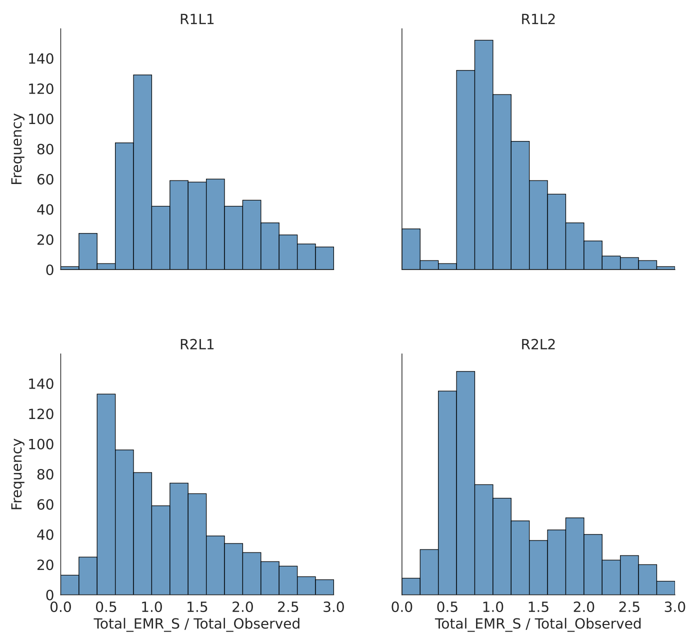

Figure 3.3: Comparison of total allocations to contracted customers by EMR-S (*Total_EMR_S*) and historical allocations (*Total_Observed*) for each sailing on the test data across routes

3.7 Conclusion

This chapter focuses on capacity planning decisions across contracts and spot market when customer demand is subject to uncertainty. Motivated by the industry perspective, we study the trade-off between spot and contract allocations to multiple customers and propose a stylized model and an exact algorithm to solve it that provide insights and outputs for operational decisions.

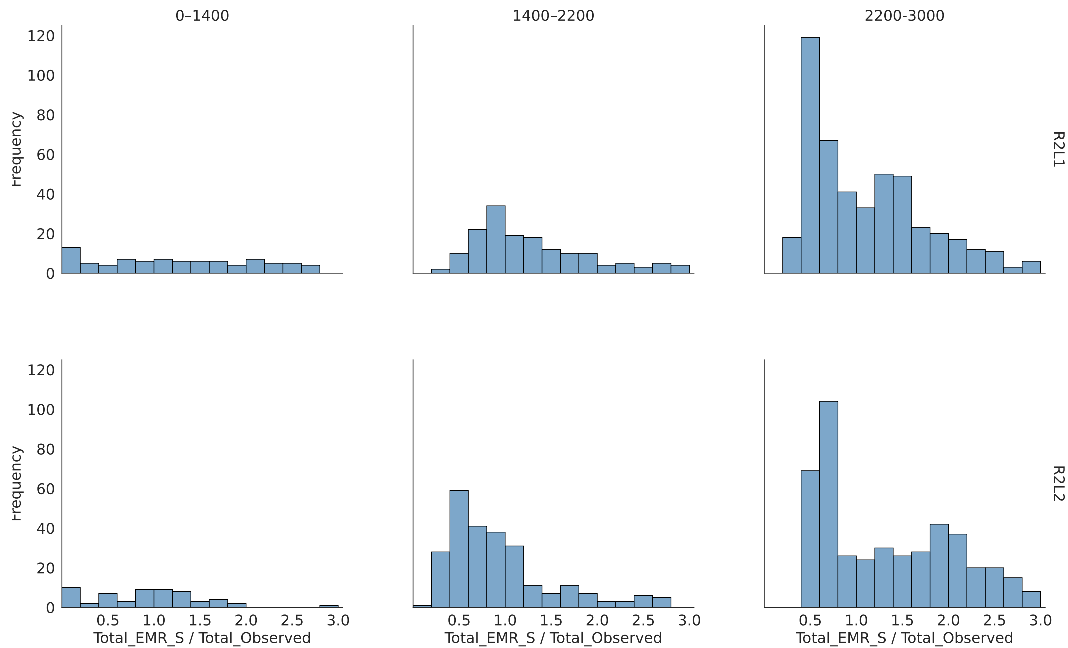

Figure 3.4: Comparison of total allocations by EMR-S ($Total\_EMR\_S$) and historical allocations ($Total\_Observed$) for each sailing in market groups on R2L1 and R2L2

In future work, it would be interesting to consider the problem of allotment and spot market coordination across a network of vessels. It would also be useful to study a richer model that incorporates strategic decisions, such as including overbooking as a decision variable. Aydın et al. (2013) develop a static model that determines the optimal overbooking limit. This overbooking limit can be used as an input to our algorithm when deciding the optimal allocation between contracts and the spot market. Another interesting avenue for future research is to characterize the optimal allocation while accounting for differing risk preferences. Traditional techniques typically assume that decision-makers are risk-neutral; however, this assumption often does not hold in practice, as individuals may exhibit varying degrees of risk aversion, as emphasized in many recent studies (Lancaster, 2003; Weatherford, 2004; Barz and Waldmann, 2007; Gönsch and Hassler, 2014). It could be valuable to explore new formulations for allocating capacity in order to maximize revenue while simultaneously meeting desired utilization targets as a measure of risk.

3.8 Appendix

Proof of Theorem 3.1. If $\sum_{i=1}^N \underline{b}_i = c$, the only feasible solution is given by $b_i^\star = \underline{b}_i$ for all $i \in \mathcal{N}$. This solution satisfies condition (3.1) for $K = 0$. Assume now that $\sum_{i=1}^N \underline{b}_i < c$. As problem $(\mathscr{P}_{\text{conv}})$ is a convex optimization problem with differentiable objective and satisfies the Slater's condition, the Karush-Kuhn-Tucker (KKT) conditions are necessary. Next, we establish the KKT conditions.

The Lagrangian function associated with $(\mathscr{P}_{\text{conv}})$ is given by

$$\mathscr{L}(\boldsymbol{b}, \lambda, \boldsymbol{\mu}) = \mathbb{E}\left[\sum_{i=1}^N p_i \min\{\tilde{d}_i, b_i\}\right] + \mathrm{R_s}\left(c - \sum_{i=1}^N b_i\right) + \lambda\left(c - \sum_{i=1}^N b_i\right) + \sum_{i=1}^N \mu_i(b_i - \underline{b}_i),$$

where $\lambda \geq 0$ is the dual variable of the capacity constraint $\sum_{i=1}^N b_i \leq c$, and $\mu_i \geq 0$, $i \in \mathcal{N}$, is dual variable of the constraint $\underline{b}_i \leq b_i$. The partial derivative of $\mathscr{L}$ with respect to each b_i is given by

$$\frac{\partial}{\partial b_i}\mathscr{L}(\boldsymbol{b}, \lambda, \boldsymbol{\mu}) = p_i\,(1 - \mathbb{F}_i(b_i)) - \mathrm{R}_\mathrm{s}'\left(c - \sum_{i=1}^N b_i\right) - \lambda + \mu_i. \tag{3.2}$$

For any optimal solution $\boldsymbol{b}^\star$ to $(\mathscr{P}_{\text{conv}})$, there exist $\lambda^\star, \boldsymbol{\mu}^\star$ such that the KKT conditions given below hold.

$$\begin{aligned}
&\sum_{i=1}^N b_i^\star \leq c,\ b_i^\star \geq \underline{b}_i\ \ \forall i \in \mathcal{N} && \text{(Primal feasibility)}\\
&\lambda^\star \geq 0,\ \mu_i^\star \geq 0\ \ \forall i \in \mathcal{N} && \text{(Dual feasibility)}\\
&\lambda^\star\left(c - \sum_{i=1}^N b_i^\star\right) = 0,\ \mu_i^\star(b_i^\star - \underline{b}_i) = 0\ \ \forall i \in \mathcal{N} && \text{(Complementary slackness)}\\
&p_i\,(1 - \mathbb{F}_i(b_i^\star)) - \mathrm{R}_\mathrm{s}'\left(c - \sum_{i=1}^N b_i^\star\right) - \lambda^\star + \mu_i^\star = 0\ \ \forall i \in \mathcal{N} && \text{(Stationarity)}
\end{aligned}$$

Next, we prove the claim by leveraging the necessary KKT conditions above. To this end, note that the stationarity condition for any $i \in \mathcal{N}$ is equivalent to

$$p_i\,(1 - \mathbb{F}_i(b_i^\star)) + \mu_i^\star = \mathrm{R}_\mathrm{s}'\left(c - \sum_{i=1}^N b_i^\star\right) + \lambda^\star$$

and that the right-hand side of this equation is the same for every $i \in \mathcal{N}$. Therefore, the stationarity condition holds if and only if

$$\begin{aligned} & p_1\left(1-\mathbb{F}_1(b_1^\star)\right)+\mu_1^\star = p_2\left(1-\mathbb{F}_2(b_2^\star)\right)+\mu_2^\star = \dots \\ & = p_N\left(1-\mathbb{F}_N(b_N^\star)\right)+\mu_N^\star = \mathrm{R}_\mathrm{s}'\left(c-\sum_{i=1}^N b_i^\star\right)+\lambda^\star . \end{aligned} \tag{3.3}$$

Recall that we assumed, without loss of generality, that $p_1\left(1-\mathbb{F}_1(\underline{b}_1)\right) \geq p_2\left(1-\mathbb{F}_2(\underline{b}_2)\right) \geq \cdots \geq p_N\left(1-\mathbb{F}_N(\underline{b}_N)\right)$. Therefore, for (3.3), and hence the stationarity condition, to be satisfied, the dual variables $\mu_i^\star$ must be non-decreasing with i, i.e., $\mu_i^\star \leq \mu_j^\star$ if $i<j$. Let $K=\max\{i \in \mathscr{N} : \mu_i^\star = 0\}$ and note that $\mu_i^\star > 0$ for all $i \in \{K+1,\dots,N\}$. By complementary slackness, this means that $b_i^\star = \underline{b}_i$ for all $i \in \{K+1,\dots,N\}$. We thus proved that for $K=\max\{i \in \mathscr{N} : \mu_i^\star = 0\}$, $\boldsymbol{b}^\star$ satisfies $b_i^\star = \underline{b}_i$ for all $i \in \{K+1,\dots,N\}$. As $\mu_i^\star = 0$ for every $i \in \{1,\dots,K\}$, by (3.3), we have

$$p_1\left(1-\mathbb{F}_1(b_1^\star)\right) = \cdots = p_K\left(1-\mathbb{F}_K(b_K^\star)\right) = \mathrm{R}_\mathrm{s}'\left(c-\sum_{i=1}^N b_i^\star\right)+\lambda^\star .$$

The observation that $\lambda^\star \geq 0$ by dual feasibility completes the proof. □

Proof of Lemma 3.3. The decision variables in ($\mathscr{P}$) as well as the demand $\tilde{d}_i$, $i \in \mathscr{N}$, can only take non-negative integer values. Thus, we have

$$\begin{aligned} \mathbb{E}[\min\{\tilde{d}_i, b_i\}] &= \sum_{k=0}^{\infty} \min\{k, b_i\} \cdot \mathbb{P}(\tilde{d}_i = k) = \sum_{k=0}^{\infty}\left(\sum_{t=1}^{b_i} \mathbf{1}\{k \geq t\}\right)\mathbb{P}(\tilde{d}_i = k) \\ &= \sum_{t=1}^{b_i}\sum_{k=0}^{\infty} \mathbf{1}\{k \geq t\}\mathbb{P}(\tilde{d}_i = k) = \sum_{t=1}^{b_i}\sum_{k=t}^{\infty} \mathbb{P}(\tilde{d}_i = k) = \sum_{t=1}^{b_i} \mathbb{P}(\tilde{d}_i \geq t). \end{aligned}$$

This completes the proof. □

Proof of Theorem 3.4. In the proof, we assume without loss of generality that $\underline{b}_i = 0$ for all $i \in \mathscr{N}$.[5] Let $\boldsymbol{b}^\mathrm{E}$ denote the allocation returned by the EMR-S, and let $\boldsymbol{b}^\star$ be an optimal allocation in problem ($\mathscr{P}$). Define the index sets

$$B_\ell = \left\{i \in \mathscr{N} \mid b_i^E < b_i^\star\right\}, \quad B_r = \left\{i \in \mathscr{N} \mid b_i^E > b_i^\star\right\}.$$

Let $\Delta B = \sum_{i \in \mathscr{N}}\left(b_i^\star - b_i^E\right)$ and note that we can write it equivalently as

$$\Delta B = \sum_{i \in B_\ell}\left(b_i^\star - b_i^E\right) - \sum_{j \in B_r}\left(b_j^E - b_j^\star\right).$$

[5] One can perform a change of variables by setting $b_i' = b_i - \underline{b}_i$ and transform the problem into an equivalent one with redefined variables and constants, preserving the structure of the optimization problem.

We denote by ΔV the difference between the objective values of $\boldsymbol{b}^\star$ and $\boldsymbol{b}^{\mathrm{E}}$, specifically, the objective value of $\boldsymbol{b}^\star$ minus that of $\boldsymbol{b}^{\mathrm{E}}$. Leveraging Lemma 3.3 and the definitions of B_ℓ and B_r, we can express ΔV as

$$\begin{aligned}
\Delta V &= \left(\sum_{i=1}^{N} p_i \sum_{t=1}^{b_i^\star} \mathbb{P}(\tilde{d}_i \geq t) + \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^\star\right)\right) - \left(\sum_{i=1}^{N} p_i \sum_{t=1}^{b_i^{\mathrm{E}}} \mathbb{P}(\tilde{d}_i \geq t) + \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^{\mathrm{E}}\right)\right) \\
&= \sum_{i=1}^{N} p_i \left(\sum_{t=1}^{b_i^\star} \mathbb{P}(\tilde{d}_i \geq t) - \sum_{t=1}^{b_i^{\mathrm{E}}} \mathbb{P}(\tilde{d}_i \geq t)\right) + \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^\star\right) - \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^{\mathrm{E}}\right) \\
&= \left(\sum_{i \in B_\ell} p_i \sum_{t=b_i^E+1}^{b_i^\star} \mathbb{P}(\tilde{d}_i \geq t)\right) - \left(\sum_{j \in B_r} p_j \sum_{t=b_j^\star+1}^{b_j^E} \mathbb{P}(\tilde{d}_j \geq t)\right) + \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^\star\right) - \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^{\mathrm{E}}\right).
\end{aligned}$$

By definition, ΔV is non-negative because $\boldsymbol{b}^\star$ represents an optimal solution. In the reminder of the proof, we will show that ΔV is also non-positive (and therefore equals zero), which implies that $\boldsymbol{b}^{\mathrm{E}}$ is optimal.

Slightly abusing the notation, we write $\Delta \mathrm{R}_s(x) = \mathrm{R}_s(x) - \mathrm{R}_s(x-1)$ for all $x \geq 0$. The concavity of R_s implies that the spot-market marginal revenue is non-increasing, that is,

$$\Delta \mathrm{R}_s(x-1) \;\geq\; \Delta \mathrm{R}_s(x) \quad \forall x \geq 0. \tag{3.4}$$

Leveraging this observation, we will next generate upper bounds on each of

$$\begin{aligned}
\Delta V_1 &= \left(\sum_{i \in B_\ell} p_i \sum_{t=b_i^E+1}^{b_i^\star} \mathbb{P}(\tilde{d}_i \geq t)\right) - \left(\sum_{j \in B_r} p_j \sum_{t=b_j^\star+1}^{b_j^E} \mathbb{P}(\tilde{d}_j \geq t)\right), \\
\Delta V_2 &= \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^\star\right) - \mathrm{R}_s\left(c - \sum_{i=1}^{N} b_i^{\mathrm{E}}\right),
\end{aligned}$$

which will translate into a bound on $\Delta V = \Delta V_1 + \Delta V_2$.

For ΔV_1, we have

$$\begin{aligned}
\Delta V_1 &\le \sum_{i\in B_\ell} p_i(b_i^\star - b_i^E)\mathbb{P}(\tilde{d}_i \ge b_i^E + 1) - \sum_{j\in B_r} p_j(b_j^E - b_j^\star)\mathbb{P}(\tilde{d}_j \ge b_j^E) \\
&\le \left(\sum_{i\in B_\ell}(b_i^\star - b_i^E)\right)\max_{i\in B_\ell}\left\{p_i\mathbb{P}(\tilde{d}_i \ge b_i^E + 1)\right\} - \left(\sum_{j\in B_r}(b_j^E - b_j^\star)\right)\min_{j\in B_r}\left\{p_j\mathbb{P}(\tilde{d}_j \ge b_j^E)\right\} \\
&\le \left(\sum_{i\in B_\ell}(b_i^\star - b_i^E)\right)\Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right) - \left(\sum_{j\in B_r}(b_j^E - b_j^\star)\right)\Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right) \\
&= (\Delta B)\Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right),
\end{aligned}$$

where the first inequality holds because $\mathbb{P}(\tilde{d}_i \ge b_i^E + 1) \ge \mathbb{P}(\tilde{d}_i \ge t)$ for all $t \ge b_i^E + 1$ and $p_i \in \mathbb{R}_+$ for all $i \in B_\ell$, and $\mathbb{P}(\tilde{d}_j \ge b_j^E) \le \mathbb{P}(\tilde{d}_j \ge t)$ for all $t \le b_j^{\mathrm{E}}$ and $p_j \in \mathbb{R}_+$ for all $j \in B_r$. The third inequality holds because by the procedure of EMR-S, we should have

$$\Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right) > \max_{i\in\mathcal{N}} p_i(1 - \mathbb{F}_i(b_i^E)) \ge \max_{i\in B_\ell} p_i\mathbb{P}(\tilde{d}_i \ge b_i^E + 1) \tag{3.5}$$

and

$$\min_{j\in B_r} p_j\mathbb{P}(\tilde{d}_j \ge b_j^E) = \min_{j\in B_r} p_j(1 - \mathbb{F}_j(b_j^E - 1)) \ge \Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right). \tag{3.6}$$

In fact, with regard to (3.5), note that if $\Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right) \le \max_{i\in\mathcal{N}} p_i(1 - \mathbb{F}_i(b_i^E))$, then $\boldsymbol{b}^E$ cannot be the output of EMR-S because in this case the algorithm would allocate less than $c - \sum_{i=1}^N b_i^E$ to the spot market and more to contracted customers. Similarly, with regard to (3.6), note that if $\min_{j\in B_r} p_j(1 - \mathbb{F}_j(b_j^E - 1)) < \Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E\right)$, then in view of (3.4), we have $\min_{j\in B_r} p_j(1 - \mathbb{F}_j(b_j^E - 1)) < \Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^E - t\right)$ for all $t \in \{0, 1, \ldots, c - \sum_{i=1}^N b_i^E\}$. This means that irrespective of the order of the allocations in the iterations of EMR-S, allocations $\boldsymbol{b}^E$ cannot be the output because EMR-S would allocate less than b_j^E to at least one $j \in B_r$.

Noting that $\Delta V_2 \le 0$ when $\Delta B \ge 0$ and $\Delta V_2 \ge 0$ when $\Delta B < 0$, we will bound ΔV_2 above by $-(\Delta B)\Delta \mathrm{R}_s\left(c - \sum_{i=1}^N b_i^{\mathrm{E}}\right)$ in two cases ($\Delta B \ge 0$ and $\Delta B < 0$) separately. This will mean that ΔV is bounded above by zero in both cases.

If $\Delta B \geq 0$, we have

$$\begin{aligned}\Delta V_2 &= \mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}} - \Delta B\right) - \mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}}\right)\\ &= \sum_{t=1}^{\Delta B}\left[\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E} - t\right) - \mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E} - t + 1\right)\right]\\ &= -\sum_{t=1}^{\Delta B}\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E} - t + 1\right) \leq -(\Delta B)\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E}\right),\end{aligned}$$

where the second equality follows from the definition of $\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E} - t + 1\right)$, and the inequality follows from (3.4) and that $\Delta B \geq 0$.

If $\Delta B < 0$, then we have

$$\begin{aligned}\Delta V_2 &= \mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}} + |\Delta B|\right) - \mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}}\right)\\ &= \sum_{t=1}^{|\Delta B|}\left[\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}} + t\right) - \mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}} + t - 1\right)\right]\\ &= \sum_{t=1}^{|\Delta B|}\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}} + t\right) \leq |\Delta B|\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}}\right) = -(\Delta B)\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{\mathrm{E}}\right),\end{aligned}$$

where the (in)equalities follow from similar arguments as before.

Therefore, for all possible values of ΔB, we obtain

$$\Delta V = \Delta V_1 + \Delta V_2 \leq (\Delta B)\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E}\right) - (\Delta B)\Delta\mathrm{R}_s\left(c-\sum_{i=1}^{N} b_i^{E}\right) = 0.$$

This completes the proof. □

Chapter 4

Conclusion

This thesis investigates the problem of resource allocation in environments where agents' actions induce uncertainty in the realized utility of resources. The primary objective is to develop resource-allocation algorithms that are both simple and computationally efficient and can achieve good performance that aligns with the objective. The algorithms discussed in the thesis exploit the inherent structure of the problem to obtain good revenue.

The thesis is structured into two principal chapters. Chapter 2 focuses on an allocation problem in an auction environment to sell a single item where the bidders' values are uncertain. The seller knows only an upper bound on these values and does not know their distribution. The objective is to design a simple and deterministic mechanism that guarantees good revenue performance across a broad class of distributions. We propose a second-price auction with a reserve price set at half of the upper bound. We show that our mechanism achieves at least $\frac{1}{4}$ ($\frac{1}{2}$ under i.i.d. values) of the maximum achievable expected revenue for a broad range of distributions of practical importance. Chapter 3 studies the allocation of limited capacity between multiple contracts and the spot market under demand uncertainty, with the objective of maximizing expected revenue. We formulate this problem as an integer program. Considering a convex relaxation of the formulated problem, we show that at optimality, the marginal revenue must be equalized across all customers. To solve the original integer program, we propose a computationally efficient algorithm that leverages the structural properties of the problem to determine the revenue-maximizing capacity allocation.

We demonstrate that, by considering simple algorithms, it is possible to achieve performance comparable to that of more complex algorithms—those that are difficult to analyze, implement, or scale in practice. This chapter concludes with a summary of the main contributions of the thesis, a discussion and an overview of future research directions and open problems.

4.1 Chapter 2 : Future Directions

In Chapter 2, we recognize that no deterministic mechanism can provide a constant-factor approximation guarantee across all distributions when only upper-bound information is available. Thus we introduce a new class of distributions that captures cases of practical relevance and establish revenue guarantees for this class under independently distributed bidders' values. These results provide a foundation for discussing the potential conservatism of considering overly broad distribution classes, which may include distributions that are less relevant in practice.

This chapter focuses on second price auction with deterministic reserve prices, which are straightforward to implement and explain. Numerical experiments in Chapter 2 indicate that our proposed mechanism performs well for more general distributions than those we rigorously studied. Thus, it would be interesting to further explore the performance of proposed deterministic mechanism in settings beyond the assumptions of the studied ambiguity set. Since the proof technique for the revenue guarantees of our mechanism relies on comparison to the single sample mechanism (Dhangwatnotai et al., 2015), these guarantees could potentially be extended by analyzing the single sample mechanism directly. Therefore, an interesting avenue is to investigate whether the single sample mechanism could yield revenue guarantees when expanded to an even broader set of distributions (beyond regular distributions for which the current revenue guarantees of the single sample mechanism holds for).

There have been attempts to study alternative classes of distributions that capture forms of non-regularity. For example, Sivan and Syrgkanis (2013) in the Bayesian setting (where bidder values follow known distributions) address the case of population heterogeneity, in which bidders are drawn from distinct groups (e.g., budget-conscious households vs. high-income professionals). Their analysis highlights how such heterogeneity naturally leads to irregular distributions. Building on this perspective, one could explore whether the single sample mechanism can provide meaningful revenue guarantees when applied to uncertainty sets that explicitly incorporate heterogeneous populations. One idea to assess the value of single sample in an heterogeneous environment is to capture information through the squared coefficient of variation—a measure defined as the ratio of variance (σ^2) to the squared mean (μ^2), which characterizes dispersion. Studying distribution classes defined by the coefficient of variation enables us to move beyond regularity assumptions while still maintaining analytical tractability. These classes are broad enough to capture multimodal, heavy-tailed, and discontinuous distributions. Recent works by Giannakopoulos et al. (2023), Wang et al. (2024), and others studies robust auctions under the assumption that both the mean and variance of the value distribution are known. Unlike them, one can seek to bound the coefficient of variation for describing ambiguity sets of practical importance.

Figure 4.1 illustrates a potential uncertainty set constructed using the nesting relationship among

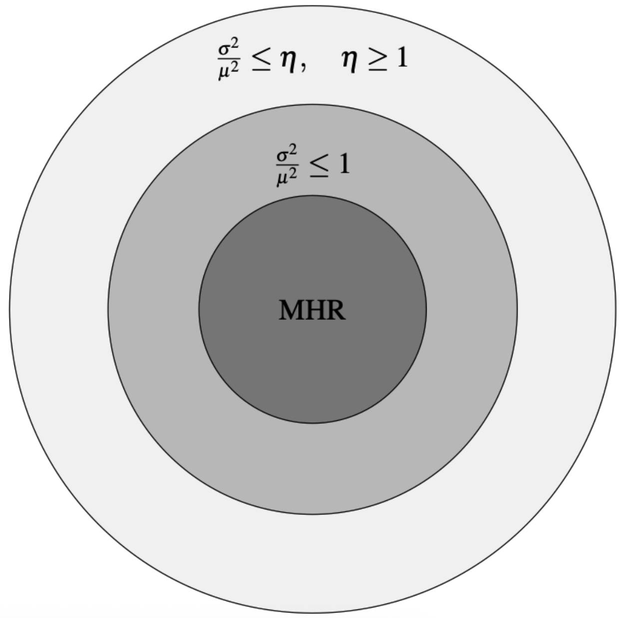

Figure 4.1: Nested distribution classes

distribution classes. The innermost circle represents distributions with monotone hazard rate (MHR). These are contained within the broader class of distributions with a coefficient of variation at most 1, which itself lies within the class of distributions with a coefficient of variation bounded by some $\eta \geq 1$. As η grows to infinity, the set expands to include all distributions (extending η to infinity would be of limited practical interest). The uncertainty set of distributions whose coefficient of variation is bounded by $\eta = 1$ therefore encompasses a wide range of structural families—such as MHR distributions—as well as multi-modal, discontinuous, and even heavy-tailed distributions. An open question is how to identify a meaningful η for real-world applications.

4.2 Chapter 3 : Future Directions

Chapter 3 focuses on a class of resource allocation problems where capacity must be allocated well in advance of actual consumption. In such settings, the actions of agents who ultimately receive the allocation are uncertain. Our EMR-S mechanism can be employed to generate an optimal policy whenever the marginal utility of each agent can be computed. The algorithm significantly reduces computation time when compared to state-of-the-art solvers like Gurobi. The algorithm we presented can potentially be used for different revenue management problems, which include cruise lines, hotels, airlines (when upgrades are restricted), and any setting where capacity can be partitioned among customer segments without substitution across segments.

The standard revenue management problems often rely on common-knowledge assumptions, which provide a useful and tractable foundation for analysis. Extending these baseline assumptions to include behavioral variation observed in practice allows for a richer understanding of decision-making. In particular, much of the literature—including our work in this chapter—assumes that decision-makers are risk-neutral. This benchmark is both insightful and practically relevant, yet it is also natural to ask what happens when decision-makers exhibit different risk attitudes. In many real-world contexts, decision-makers who exibits preferences and biases over objectives—for instance may be risk-averse, valuing stability and consistent utilization, or risk-seeking, willing to tolerate variability in pursuit of higher returns or market share. Recent studies highlight the importance of accounting for such heterogeneous risk preferences (Lancaster, 2003; Weatherford, 2004; Barz and Waldmann, 2007; Gönsch and Hassler, 2014). A natural direction for future research is therefore to incorporate risk attitudes explicitly into revenue management systems, such as pricing and capacity allocation. For example, in the freight shipping setting, the partner company from Chapter 3 faces the challenge of allocating limited capacity between long-term loyal customers and new ones. Loyal customers reliably use capacity and are often offered lower per-unit prices to preserve long-term relationships, while new customers offer potential for higher revenue but introduce uncertainty about utilization. In this context, utilization targets can be interpreted as reflecting the firm's risk attitude: a risk-averse company may prioritize loyal customers to ensure a stable utilization even at the cost of lower revenue, whereas a risk-seeking company may allocate more to uncertain but higher-yielding new customers. One way to capture this trade-off is to formulate the objective as a convex combination of revenue and utilization.

Bibliography

Allouah, A., Bahamou, A., and Besbes, O. (2022). Pricing with samples. *Operations Research*, 70(2):1088–1104.

Allouah, A., Bahamou, A., and Besbes, O. (2023). Optimal pricing with a single point. *Management Science*, 69(10):5866–5882.

Allouah, A. and Besbes, O. (2020). Prior-independent optimal auctions. *Management Science*, 66(10):4417–4432.

Amaruchkul, K., Cooper, W. L., and Gupta, D. (2007). Single-leg air-cargo revenue management. *Transportation Science*, 41(4):457–469.

Anunrojwong, J., Balseiro, S. R., and Besbes, O. (2023). Robust auction design with support information. *arXiv preprint*, page arXiv:2305.09065.

Araman, V. F. and Popescu, I. (2010). Media revenue management with audience uncertainty: Balancing upfront and spot market sales. *Manufacturing & Service Operations Management*, 12(2):190–212.

Aydın, N., Birbil, Ş. İ., Frenk, J., and Noyan, N. (2013). Single-leg airline revenue management with overbooking. *Transportation Science*, 47(4):560–583.

Azar, P. and Micali, S. (2012). Optimal parametric auctions. Technical report available at https://dspace.mit.edu/bitstream/handle/1721.1/70556/MIT-CSAIL-TR-2012-011.pdf.

Azar, P., Micali, S., Daskalakis, C., and Weinberg, S. M. (2013). Optimal and efficient parametric auctions. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 596–604.

Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory*, 26(2):445.

Bandi, C. and Bertsimas, D. (2014). Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research*, 39(4):1012–1038.

Barz, C. and Waldmann, K.-H. (2007). Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research*, 65(3):565–579.

Belobaba, P. P. (1987). Survey paper—airline yield management an overview of seat inventory control. *Transportation Science*, 21(2):63–73.

Belobaba, P. P. (1992). Optimal vs. heuristic methods for nested seat allocation. In *Presentation at ORSA/TIMS Joint National Meeting*, volume 1.

Bergemann, D. and Schlag, K. (2011). Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543.

Bergemann, D. and Schlag, K. H. (2008). Pricing without priors. *Journal of the European Economic Association*, 6(2-3):560–569.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

Cachon, G. P. and Kök, A. G. (2007). Implementation of the newsvendor model with clearance pricing: How to (and how not to) estimate a salvage value. *Manufacturing & Service Operations Management*, 9(3):276–290.

Carrasco, V., Luz, V. F., Kos, N., Messner, M., Monteiro, P., and Moreira, H. (2018). Optimal selling mechanisms under moment conditions. *Journal of Economic Theory*, 177:245–279.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488.

Chen, H., Hu, M., and Perakis, G. (2022). Distribution-free pricing. *Manufacturing & Service Operations Management*, 24(4):1939–1958.

Chen, Z., Hu, Z., and Wang, R. (2024). Screening with limited information: A dual perspective. *Operations Research*, 72(4):1487–1504.

Cohen, M. C., Perakis, G., and Pindyck, R. S. (2021). A simple rule for pricing with limited knowledge of demand. *Management Science*, 67(3):1608–1621.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2015). Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333.

Du, C., He, F., and Lin, X. (2025). Dynamic pricing for air cargo revenue management. *Transportation Research Part E: Logistics and Transportation Review*, 197:104088.

Dumouchelle, J., Frejinger, E., and Lodi, A. (2024). Reinforcement learning for freight booking control problems. *Journal of Revenue and Pricing Management*, 23(4):318–345.

Ewerhart, C. (2013). Regular type distributions in mechanism design and concavity. *Economic Theory*, 53(3):591–603.

Giannakopoulos, Y., Poças, D., and Tsigonias-Dimitriadis, A. (2023). Robust revenue maximization under minimal statistical information. *ACM Transactions on Economics and Computation*, 10(3):1–34.

Gillai, B. and Lee, H. L. (2009). Internet-based distribution channel for product diversion with potential manufacturer's intervention. *Consumer-Driven Demand and Operations Management Models: A Systematic Study of Information-Technology-Enabled Sales Mechanisms*, pages 115–154.

Gönsch, J. and Hassler, M. (2014). Optimizing the conditional value-at-risk in revenue management. *Review of Managerial Science*, 8(4):495–521.

Gosavi, A., Bandla, N., and Das, T. K. (2002). A reinforcement learning approach to a single leg airline revenue management problem with multiple fare classes and overbooking. *IIE transactions*, 34(9):729–742.

Hart, S. and Nisan, N. (2017). Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347.

Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234.

Huang, K. and Chang, K.-c. (2010). An approximate algorithm for the two-dimensional air cargo revenue management problem. *Transportation Research Part E: Logistics and Transportation Review*, 46(3):426–435.

Kasilingam, R. G. (1997). Air cargo revenue management: Characteristics and complexities. *European Journal of Operational Research*, 96(1):36–44.

Klein, R., Koch, S., Steinhardt, C., and Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2):397–412.

Koçyiğit, Ç., Iyengar, G., Kuhn, D., and Wiesemann, W. (2020). Distributionally robust mechanism design. *Management Science*, 66(1):159–189.

Koçyiğit, Ç., Kuhn, D., and Rujeerapaiboon, N. (2024). Regret minimization and separation in multi-bidder, multi-item auctions. *INFORMS Journal on Computing*, 36(6):1543–1561.

Koçyiğit, Ç., Rujeerapaiboon, N., and Kuhn, D. (2021). Robust multidimensional pricing: Separation without regret. *Mathematical Programming*, pages 1–34.

Kunnumkal, S. and Topaloglu, H. (2011). A stochastic approximation algorithm to compute bid prices for joint capacity allocation and overbooking over an airline network. *Naval Research Logistics (NRL)*, 58(4):323–343.

Lancaster, J. (2003). The financial risk of airline revenue management. *Journal of Revenue and Pricing Management*, 2(2):158–165.

Lawhead, R. J. and Gosavi, A. (2019). A bounded actor–critic reinforcement learning algorithm applied to airline revenue management. *Engineering Applications of Artificial Intelligence*, 82:252–262.

Lee, T. C. and Hersh, M. (1993). A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, 27(3):252–265.

Levin, Y., Nediak, M., and Topaloglu, H. (2012). Cargo capacity management with allotments and spot market demand. *Operations Research*, 60(2):351–365.

Littlewood, K. (1972). Forecasting and control of passenger bookings. *The Airline Group of the International Federation of Operational Research Societies*, 12:95–117.

Moussawi-Haidar, L. (2014). Optimal solution for a cargo revenue management problem with allotment and spot arrivals. *Transportation Research Part E: Logistics and Transportation Review*, 72:173–191.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Rujeerapaiboon, N., Wei, Y., and Xue, Y. (2023). Target-oriented regret minimization for satisficing monopolists. In *International Conference on Web and Internet Economics*, pages 563–581. Springer.

Sivan, B. and Syrgkanis, V. (2013). Vickrey auctions for irregular distributions. In *Web and Internet Economics: 9th International Conference, WINE 2013, Cambridge, MA, USA, December 11-14, 2013, Proceedings 9*, pages 422–435. Springer.

Suzdaltsev, A. (2022). Distributionally robust pricing in independent private value auctions. *Journal of Economic Theory*, 206:105555.

Talluri, K. and Van Ryzin, G. (1998). An analysis of bid-price controls for network revenue management. *Management Science*, 44(11-part-1):1577–1593.

Talluri, K. T. and Van Ryzin, G. J. (2006). *The theory and practice of revenue management*, volume 68. Springer Science & Business Media.

Tavana, H. (2004). Optimal airline seat allocation in a fenceless environment using expected revenue maximization. *AGIFORS Reservations and Yield Management Study Group (Auckland, New Zealand)*.

Tavana, H. and Weatherford, L. (2017). Application of an alternative expected marginal seat revenue method (emsrc) in unrestricted fare environments. *Journal of Air Transport Management*, 62:65–77.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.

Wang, S., Liu, S., and Zhang, J. (2024). Minimax regret robust screening with moment information. *Manufacturing & Service Operations Management*, 26(3):992–1012.

Weatherford, L. R. (2004). Additional paper: Emsr versus emsu: revenue or utility? *Journal of Revenue and Pricing Management*, 3(3):277–284.

Wilson, R. (1987). Game-theoretic analysis of trading processes. *Bewley TF, ed. Adv. Econom. Theory: Fifth World Congress*, (Cambridge University Press, Cambridge, UK):33–70.

Yang, L. and Yu, M. (2025). Too good to go: Combating food waste with surprise clearance. *Management Science (forthcoming)*.