MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Dynamic Pricing and Capacity Optimization in Railways

Manchiraju, Chandrasekhar; Dawande, Milind; Janakiraman, Ganesh; Raghunathan, Arvind

TR2025-007 January 08, 2025

Abstract

Problem definition: Revenue management in railways distinguishes itself from that in traditional sectors such as airline, hotel, and fashion retail, in several important ways: (i) Capacity is substantially more flexible, in the sense that changes to the capacity of a train can often be made throughout the sales horizon. Consequently, the joint optimization of prices and capacity assumes genuine importance. (ii) Capacity can only be added in discrete "chunks", i.e., coaches. (iii) Passengers with unreserved tickets can travel in any of the multiple trains available during the day. Further, passengers in unreserved coaches are allowed to travel by standing, thus giving rise to the need to manage congestion. Motivated by our work with a major railway company in Japan, we analyze the problem of jointly optimizing pricing and capacity − this problem is a more-general version of the canonical multiproduct dynamic-pricing problem. Methodology/Results: Our analysis yields four asymptotically optimal policies. From the viewpoint of the pricing decisions, our policies can be classified into two types − static and dynamic. With respect to the timing of the capacity decisions, our policies are again of two types − fixed capacity and flexible capacity. We establish the convergence rates of these policies: when demand and supply are scaled by a factor k P N, the optimality gaps of the static policies scale proportional to ?k, and those of the dynamic policies scale proportional to log k. We illustrate the attractive performance of our policies on a test-suite of instances based on real-world operations of the high-speed "Shinkansen" trains in Japan, and develop associated insights. Managerial implications: Our work provides railway administrators with simple and effective policies for pricing, capacity, and congestion management. Our policies cater to different contingencies that decision- makers may face in practice: the need for static or dynamic prices, and for fixed or flexible capacity.

Manufacturing and Service Operations Management 2024

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Dynamic Pricing and Capacity Optimization in Railways

Chandrasekhar Manchiraju
Eli Broad College of Business, Michigan State University, chandumanchi@gmail.com

Milind Dawande, Ganesh Janakiraman
Naveen Jindal School of Management, The University of Texas at Dallas, {milind, ganesh}@utdallas.edu

Arvind Raghunathan
Mitsubishi Electric Research Laboratories Inc., raghunathan@merl.com

Problem definition: Revenue management in railways distinguishes itself from that in traditional sectors such as airline, hotel, and fashion retail, in several important ways: (i) Capacity is substantially more flexible, in the sense that changes to the capacity of a train can often be made throughout the sales horizon. Consequently, the *joint optimization* of prices and capacity assumes genuine importance. (ii) Capacity can only be added in discrete "chunks", i.e., coaches. (iii) Passengers with unreserved tickets can travel in any of the multiple trains available during the day. Further, passengers in unreserved coaches are allowed to travel by standing, thus giving rise to the need to manage congestion. Motivated by our work with a major railway company in Japan, we analyze the problem of jointly optimizing pricing and capacity – this problem is a more-general version of the canonical multiproduct dynamic-pricing problem.
Methodology/Results: Our analysis yields four asymptotically optimal policies. From the viewpoint of the pricing decisions, our policies can be classified into two types – static and dynamic. With respect to the timing of the capacity decisions, our policies are again of two types – fixed capacity and flexible capacity. We establish the convergence rates of these policies: when demand and supply are scaled by a factor $\kappa \in \mathbb{N}$, the optimality gaps of the static policies scale proportional to $\sqrt{\kappa}$, and those of the dynamic policies scale proportional to $\log \kappa$. We illustrate the attractive performance of our policies on a test-suite of instances based on real-world operations of the high-speed "Shinkansen" trains in Japan, and develop associated insights.
Managerial implications: Our work provides railway administrators with simple and effective policies for pricing, capacity, and congestion management. Our policies cater to different contingencies that decision-makers may face in practice: the need for static or dynamic prices, and for fixed or flexible capacity.

1. Introduction

Revenue management techniques can be fundamentally classified into two types – price-based and quantity-based. While price-based techniques manage demand through the pricing of the products offered through time, quantity-based techniques directly control the consumption of products through time. Reveneue management practices originated in the airline industry and have now been adopted in several other major industries, including hotels, restaurants, retail, and broadcast television; see e.g., Talluri and Van Ryzin (2006).

Researchers have also recognized the potential of using revenue management techniques in railway operations; e.g., Konno and Raghunathan (2020), Kamandanipour et al. (2020), Wang et al. (2016), and Hetrakul and Cirillo (2014). In most developing countries, railways continue to fulfill a strong social obligation as an affordable means of transportation for the common people. However, as technology has progressed and competition from other means of transportation has increased, the demand from railway passengers for the latest services, amenities, and comforts has also grown. Accordingly, to compete effectively, railways (both government-owned and privately-owned) are increasingly placing emphasis on improving profits, so that better facilities and services can be provided. The increasing adoption of revenue management principles in railways worldwide is in line with this goal. Below, we discuss railway practice in the context of pricing (static or dynamic pricing) and the timing of capacity decisions (fixed or flexible capacity), whether or not unreserved travel is allowed, and whether or not standing travel is allowed.

Railway practice on pricing: In the United States, Amtrak uses dynamic pricing for all passenger trains.

> "Similar to airlines, Amtrak uses dynamic pricing, meaning fares will rise and fall based on consumer demand. With some careful attention, and luck, you can get a fare at the minimum fare price even days before your departure." [The Urbanist 2019]

In Europe, most of the international trains use dynamic pricing. However, some countries use fixed fares for domestic travel.

> "International train fares within Europe all seem to have this dynamic pricing where the fare goes up as the date approaches and more tickets are sold. However, some countries including Switzerland and the Scandinavian countries have fixed domestic fares that can be reasonable if bought on travel day. However, they also have "super saver" fares on some routes where they offer very cheap tickets if you buy far enough in advance." [Price of Travel 2022]

The Indian railway network, the second busiest in the world after China, uses a mix of fixed fares and dynamic pricing. The exclusive "premium" trains (fast, direct routes) use dynamic pricing, while most non-premium trains use fixed pricing.

> "Premium trains are popular for their speed, direct routes, and covering long distance journeys in less time duration . . . Currently, there are more than 140 trains catering in this segment operated by the Indian Railways . . . What goes common with these trains is the "Dynamic Fare Pricing" of the journey tickets. Meaning, the ticket prices change according to the availability of the berths in these particular trains." [Rail Mitra 2019]

In China, dynamic pricing is being explored. Currently, except for a few high-speed trains on the busiest routes, most of the trains have fixed fares. Dynamic pricing was introduced on the Beijing-Shanghai high-speed rail in December 2020.

> "Rail line operator Beijing-Shanghai High Speed Railway Co Ltd announced in October 2020 that it would move away from flat fares to a variable pricing model." [Business Traveller 2020]

Railway practice on capacity: While most networks allow for additional coaches to be added to trains to cater to high demand, there does not seem to be precise information in the public domain on the deadline for adding such coaches. Our discussions with practitioners indicate that the entire spectrum of deadlines ranging from the *start* of the selling season (fixed capacity) to the *end* of the selling season (flexible capacity) is possible.

Railway practice on unreserved travel and standing travel: Most major railway networks allow unreserved travel and standing travel. In the United States, Amtrak allows unreserved travel only on select routes, but does not allow standing travel [Amtrak 2022, Trains.com 2019].

Railway operations distinguish themselves from those in airlines in many important ways; we discuss a few below.

(a) *Joint Optimization of Pricing and Capacity:* Studies in the mainstream revenue management literature typically analyze the pricing (or inventory control) of multiple products over a finite sales horizon, under exogenously-specified capacities of the resources needed to offer these products. That is, the capacity of each resource is assumed to be fixed at the beginning of the sales horizon, and dynamic pricing is used to maximize the revenue generated over the horizon; examples include the sale of airline seats (Bumpensanti and Wang 2020, Kunnumkal and Talluri 2016), hotel rooms (Zhang and Weatherford 2017), and fashion products (Ferreira et al. 2018, Ferreira et al. 2016). In such applications, the assumption of fixed capacity is reasonable since there is little flexibility to dynamically change capacities, or the lead times required for such changes are significant. For instance, in the case of airline tickets, the limited availability of aircrafts and the simultaneous need for other scarce resources such as pilots and crews leave little room to meaningfully make dynamic updates to capacity.

In significant contrast, capacity is substantially more flexible in railway operations. First, the ability to add individual coaches to a train allows for a more granular control over capacity. Second, several countries have extensive rail networks with critical resources (coaches, drivers, maintenance staff and equipment) spread over the entire network. Third, these resources are typically cheaper than those in the airline industry and hence their availability is not as tightly constrained. Together, these characteristics often make it feasible to quickly support changes in the capacity of a train throughout the sales horizon. Consequently, the joint optimization of pricing and capacity gains genuine importance in railways.

(b) *Discrete Capacity:* Capacity can only be added in discrete sizes, namely coaches. This discreteness leads to technical challenges – in particular, the need to solve non-convex problems, which typically requires solving their continuous relaxations followed by rounding procedures.

(c) *Unreserved Capacity and Congestion Management:* In many countries around the world, railway travel is the most affordable, and hence popular, option. Consequently, passenger demand usually exceeds the supply of seats, and railway firms are often mandated to provide unreserved coaches where passengers are allowed to stand (if all the seats are taken). Often, each train is required to offer a certain minimum amount of unreserved capacity. The presence of unreserved capacity leads to several important features. First, unreserved tickets are typically valid for a particular day and are not train-specific; thus, passengers with such tickets can travel in any of the multiple trains available on the same day. Second is the substitutability of demand between reserved and unreserved tickets. For example, an increase in the price of reserved tickets for an itinerary can lead to an increase in the demand for unreserved tickets. Third is the need to manage congestion (number of standing passengers) in the unreserved coaches: allowing more standing passengers can increase revenue, but may deteriorate travel experience.

In this paper, motivated by our work with a major railway company in Japan, we study the following joint pricing and capacity optimization problem:

Consider a railway firm that operates trains at different times of the day on a single railway line that serves multiple stations. Let $L+1$ be the total number of stations on the line, including the start and end stations. Thus, the railway line comprises of L consecutive *legs* $(i, i+1)$, $i = 1, 2, \ldots, L$. Let M be the number of distinct trains that operate in a day from start to end on this line. Let $N = \binom{L+1}{2}$ be the number of available *itineraries* for any train m, $1 \leqslant m \leqslant M$, where an itinerary is defined by a starting station i and an ending station j, $i < j$. The firm offers two types of tickets for any itinerary – a *reserved* ticket and an *unreserved* ticket. A reserved ticket specifies the itinerary, date, and the specific train on which the customer must travel; the customer cannot travel with this ticket on any train other than the one specified. The customer is also guaranteed a seat for the entire journey. An unreserved ticket, on the other hand, specifies only the itinerary and the date of travel, and does not specify the train on which the customer must travel. The customer can travel with this ticket on any train that operates on the date of travel; however, she is not guaranteed a seat. Corresponding to the two types of tickets, the firm operates two types of *coaches* on each train – reserved and unreserved; passengers who buy a reserved (resp., unreserved) ticket can travel only on a reserved (resp., unreserved) coach.

We focus on a single day of operations on this railway line. Tickets for travel on this focal day are sold during an advance *sales horizon* consisting of T periods: 1,2,..., T, which immediately precede that day; for brevity, we denote the sales horizon by $[1, T]$. For this single day of operations, the firm sells both reserved and unreserved tickets for all the itineraries and can vary the price of any ticket, on any train, during the sales horizon. The demand arrivals for the tickets are stochastic, with arrival rates that depend on ticket prices. For every coach used on a train, the firm incurs a fixed operational cost. There are also limits on the number of coaches used on each train; specifically, there is an upper bound on the total number of coaches and a lower bound on the number of unreserved coaches used on each train. Furthermore, the firm needs to guarantee a seat for every reserved ticket sold, on all the itineraries, on all the trains. On the unreserved coaches, since standing is allowed, the firm needs to manage the congestion due to potential standing. To model this congestion, we include a per-leg penalty cost on each leg for each standing passenger. Further details on the demand and supply models are provided in Section 2. For this setting, the firm's goal is to maximize its expected net profit (namely, revenue less the operational cost and the congestion penalty) by making the following decisions: (i) The prices of reserved and unreserved tickets on each itinerary, on each train, at every time instant during the sales horizon and (ii) The capacity, i.e., the number of reserved and unreserved coaches, of each train. This problem, which we refer to as $\mathbb{P}_S$, is formally defined in Section 2.3.

1.1. Our Results and Contributions

We offer four asymptotically optimal policies for our problem of joint pricing and capacity decisions:

- **Static- and Dynamic-Pricing Policies:** From the viewpoint of the pricing decisions, our four policies can be classified into two types – static and dynamic; see Figure 1. While the dynamic-pricing policies are more sophisticated and offer a relatively better performance, the static-pricing policies are also attractive due to their simplicity and a surprisingly strong overall performance. As discussed earlier, both types of policies are of interest in practice.

- **Fixed- and Flexible-Capacity Policies:** With respect to the timing of the capacity decisions, our policies can be classified into two types – fixed capacity and flexible capacity. In the latter, capacity is "flexible" throughout the sales horizon, in the sense that the capacity decisions are made only at the *end* of the sales horizon. While the firm would naturally prefer the luxury of delaying capacity decisions to the extent possible, operating conditions in busy seasons (when the demand for resources is high) may sometimes necessitate that the firm make these decisions in advance. This motivates policies in which the capacity decisions are made and fixed at the *beginning* of the sales horizon and then remain unchanged; accordingly, we refer to these policies as "fixed" capacity policies.

For all four policies, we establish performance guarantees on the rates at which they converge to optimality. When supply and demand are scaled by a factor $\kappa \in \mathbb{N}$, the optimality gaps of both our static policies scale proportional to $\sqrt{\kappa}$ (Theorems 2 and 3) whereas the optimal profit of the upper-bound problem scales proportional to κ, thereby establishing the asymptotic optimality of these policies (Corollaries 1 and 2). For both of our dynamic policies, the optimality gaps scale proportional to $\log \kappa$ (Theorems 4 and 5) and hence establish their asymptotic optimality (Corollaries 3 and 4). Figure 1 also provides a convenient view of our main results.

We also illustrate the impressive performance of all our policies on a test bed with problem instances generated from the real-world operations of the high-speed "Shinkansen" trains in Japan. The average relative optimality gap (over all the instances) of each of our policies is at most 10%. Our numerical analysis also shows that dynamic-pricing policies significantly outperform static-pricing policies when the demand arrival rates are high, while the performance of flexible-capacity policies is significantly better than that of fixed-capacity policies under low demand arrival rates.

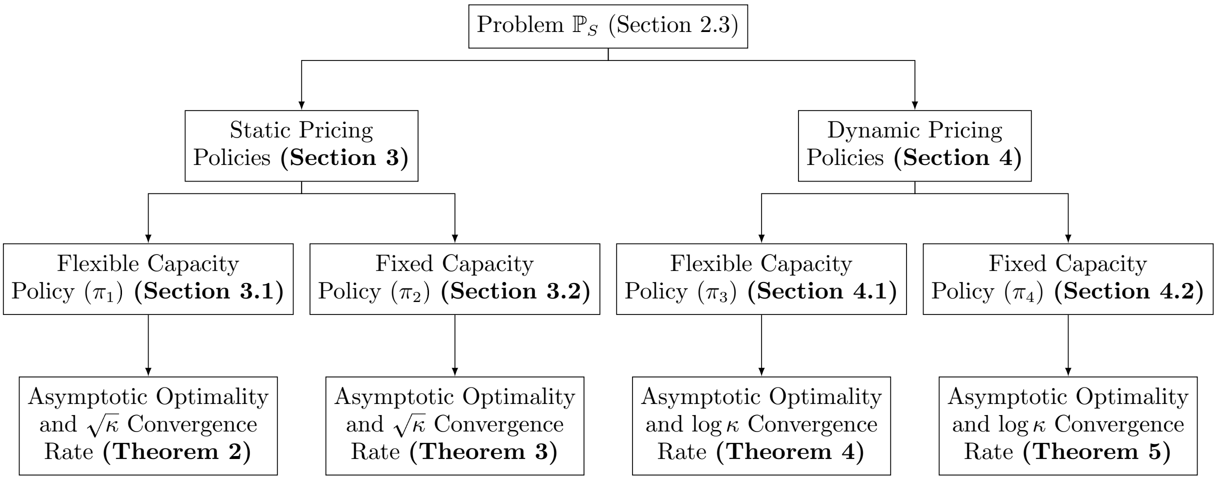

Figure 1 Organization of our analysis and summary of results.

To the best of our knowledge, ours is the first comprehensive study of revenue management in railways – as discussed earlier, this setting has several distinguishing features of railway operations that are new to the literature, including the joint optimization of pricing and capacity decisions, discrete "jumps" in capacity, and unreserved capacity along with the need for congestion management. While it is clear that the capacity decisions in our problem are new relative to the existing revenue management literature, it is important to note that these decisions determine the cost incurred by the firm. It is the analysis of the total cost in our problem that necessitates a substantial amount of new and challenging analysis. Although there are other challenges as well, e.g., capacity can only be acquired in discrete steps and there are bounds to be satisfied on the number of reserved and unreserved coaches used on each train, our main technical innovation is in the capacity analysis. Below, we explain why the analysis of the cost is challenging and how we address it.

Main Challenge: Bounding the Total Cost

Bounding the total cost incurred requires bounding *both* the reserved and unreserved costs. In what follows, we first highlight and then elaborate on the important challenges in bounding the total cost incurred. Subsequently, we also briefly discuss other challenges in our analysis, including the restriction that capacity can only be added in multiples of a discrete size.

The objective of our joint optimization problem $\mathbb{P}_S$, namely the maximization of the firm's expected net profit over the sales horizon, consists of three components: the revenue earned from ticket sales, the cost incurred in operating (reserved and unreserved) coaches, and the total penalty cost incurred as a result of passengers standing in unreserved coaches. The most significant challenges arise due to the analysis of the cost portion. In particular, to establish the asymptotic optimality of any of our policies and the order of our optimality loss, we need to derive an upper bound on the difference between the cost incurred by the policy and the benchmark (optimal cost incurred in the deterministic problem $\mathbb{P}_D$, defined in Section 2.3). This particular component of our problem necessitates a significant amount of new analysis. Broadly, the need for new analysis arises because the computation of the cost incurred by any policy in our problem, over any time interval, is *non-separable*; i.e., the cost incurred over a time interval cannot be easily computed by decomposing it as the sum of costs incurred in each period of the interval. Further, computing the total cost incurred by a policy requires characterizing the demand arrivals under that policy until the end of the sales horizon; this characterization is also non-trivial. In what follows, we first explain the difficulties in computing the total cost incurred by our policy due to the following two reasons: (a) the non-separability of cost over time and (b) the difficulty in characterizing demand arrivals. Then, in (c), we elaborate on how we handle these challenges in cost computation discussed in (a) and (b). We now proceed with the details.

(a) In problem $\mathbb{P}_S$, the cost incurred by a policy in a given time period, say t, cannot be computed in a straightforward manner. The fundamental reason behind this difficulty is the fact that the sale of a ticket on a specific leg does not necessarily increase the cost incurred! Indeed, the cost incurred at any time is determined by the number of coaches needed, which, in turn, is determined by the leg on which the *maximum* number of tickets are sold. Consequently, a ticket sold on a leg that has fewer ticket sales than the maximum does not lead to an increase in cost. Thus, the cost incurred by a policy in period t depends on the *state of the system* in that period; that is, the number of tickets sold on each leg of each train until period t. In turn, this state determines the number of coaches required to satisfy the demand generated until period t, which ultimately determines the cost incurred until period t. Due to this dependency of cost on the state of the system, computing the expected cost incurred until period t as the sum of the individual expected costs incurred from periods 1 through t is intractable. Hence, the exisiting analyses in the literature cannot be extended in a straightforward manner to handle the analysis of the cost function. Below, in (c), we explain how we handle this challenge in the context of our dynamic-price flexible-capacity policy.

(b) We now discuss the difficulty in characterizing demand arrivals until the end of the sales horizon. Consider any policy (π) for problem $\mathbb{P}_S$. To derive a lower bound on the profit obtained under this policy, we separately analyze the revenue and cost portions of the profit. Let $R^\pi[t_1, t_2]$ and $C^\pi[t_1, t_2]$ denote the revenue earned and cost incurred, respectively, under policy π, from period t_1 until period t_2. Thus, we need to obtain a lower bound on $R^\pi[1, T]$ and an upper bound on $C^\pi[1, T]$. Instead of bounding $R^\pi[1, T]$, we obtain our desired lower bound on $R^\pi[1, \tau]$, where τ is the first period in which all the available seats on one of the legs of one of the trains are sold out. Thus, the pricing decisions after τ do not affect the bound on the revenue. However, the analysis of cost has to be different. This is because the pricing decisions after period τ affect the total cost incurred and therefore, bounding the cost incurred only until period τ does not result in a valid upper bound. Hence, relative to pricing policies for revenue maximization problems which do not involve costs, our policy has to be developed more carefully. Further, after period τ, it is difficult to characterize the demand arrivals in problem $\mathbb{P}_S$ since the demand on the leg on which capacity is exhausted is turned off after period τ. This leads to difficulties in bounding $C^{\pi}[1,T]$.

We now explain how we handle the above challenges in the context of our analysis of the dynamic-price flexible-capacity policy (policy π_3 in Section 4.1).

(c) Our dynamic-price flexible-capacity policy is a *linear rate-control* policy (see, e.g., Jasin 2014, Besbes and Maglaras 2012, and Atar and Reiman 2012). As discussed above, relative to the revenue management literature, the analysis of such a policy for problem $\mathbb{P}_S$ requires us to address the cost component of the objective function. To handle the challenge in bounding $C^{\pi_3}[1,T]$ (the total cost under policy π_3), we consider problem $\mathbb{P}_U$, an unconstrained version of our problem $\mathbb{P}_S$ in which there is no capacity restriction on any train, i.e., no bound on the number of coaches used on any train. For this unconstrained problem, we define a modified version of our policy, denoted by $\tilde{\pi}_3$, in which ticket sales are allowed until the end of the sales horizon, assuming unlimited capacity. Using the constructions of our original policy π_3 and its modified version $\tilde{\pi}_3$, we show that the cost incurred by the modified policy $\tilde{\pi}_3$ in the unconstrained problem $\mathbb{P}_U$ is higher than the cost incurred by policy π_3 in problem $\mathbb{P}_S$. Since capacity is never exhausted in problem $\mathbb{P}_U$, it is easier to characterize the demand arrivals under policy $\tilde{\pi}_3$ until the end of the sales horizon. Therefore, instead of bounding $C^{\pi_3}[1,T]$, we focus on bounding $C^{\tilde{\pi}_3}[1,T]$, the total cost incurred by the modified policy $\tilde{\pi}_3$ in problem $\mathbb{P}_U$.

As discussed earlier in (a), bounding $C^{\tilde{\pi}_3}[1,T]$ is also challenging since it is not easy to compute the total expected cost incurred over the sales horizon as the sum of expected costs in each period of the horizon. To handle this challenge, we use the following approach: Recall that the number of reserved coaches used on a train is determined by the leg with the highest number of tickets sold. Therefore, the total expected cost of reserved capacity for a train is determined by the expectation of the maximum over the number of tickets sold on each leg. Based on this observation, we first focus on bounding the (stochastic) cost incurred on each leg of a train along every sample path. Using these bounds, we bound the *maximum* cost incurred over all the legs for each sample path. Finally, using these bounds on the maximum cost, we obtain our desired upper bound on the total *expected* cost of reserved capacity incurred on a train. A similar approach is also used to obtain our bound for the expected cost from unreserved capacity.

Other Challenges: Discrete Capacity and Bounds on the Number of Coaches

Below, we summarize some of the other challenges encountered in our analysis and the approach we adopt to handle them.

- The fact that capacity can only be added in multiples of a discrete size, namely the size of a coach, also poses a challenge in our analysis. Our policies use the solution to the deterministic relaxation of problem $\mathbb{P}_S$; this relaxation is a non-convex optimization problem due to the discreteness in capacity. To address this difficulty, we examine the continuous counterpart of this relaxation by further relaxing capacity to be continuous and solving the resulting convex program.

This solution, which uses a fractional number of coaches, is then modified by careful rounding to ensure that our policies are feasible and also offer the desired performance guarantees.

- In our setting, capacity is not exogenously specified but there are restrictions on the total number of coaches used on each train and a lower bound on the number of unreserved coaches used on each train. Therefore, to guarantee feasibility and to prove the asymptotic convergence rate for any of our policies, we need to devise specific limits to impose on the number of reserved and unreserved seats to sell on each train, under that policy.

1.2. Literature Review

Broadly, our joint pricing and capacity optimization problem belongs to the stream of work on multi-product dynamic-pricing problems in the revenue management literature. An early study in this stream is Gallego and Van Ryzin (1997), which examines the dynamic pricing of multiple products over a finite sales horizon and under fixed capacities of the resources that are required to offer these products. For this problem, the authors develop two asymptotically optimal static policies, and establish the rates at which they converge to optimality. Several subsequent studies revisit the multi-product dynamic-pricing problem and develop policies with stronger performances; we discuss a few here. Maglaras and Meissner (2006) develop asymptotically optimal, resolving-based dynamic policies and use an extensive numerical study to demonstrate their superior performance over static policies. Jasin (2014) develops two dynamic policies – re-optimized static control and linear rate correction – and shows that the optimality gaps of these policies scale proportional to $\log \kappa$ when the demands and capacities are scaled by $\kappa \in \mathbb{N}$. Ke et al. (2019) adopt an approximate dynamic programming approach to develop a pricing policy and conduct an extensive numerical study to demonstrate its strong performance. Wang and Wang (2022) establish an $\mathcal{O}(1)$ optimality gap for the re-optimized static control policy in Jasin (2014). In addition to these studies, several multi-product dynamic-pricing problems that incorporate a variety of additional constraints, motivated by real-world applications, have also been addressed in the literature; examples include Lei et al. (2018) [e-commerce retail], Ferreira et al. (2018) [fashion retail], Zhang and Weatherford (2017) [hotels] and Chen et al. (2016) [airlines, hotels].

The above studies analyze the pricing of products under fixed and exogenously specified capacities of the resources. In contrast, capacities are decisions in our analysis. Moreover, as discussed earlier, our setting of railway operations has other distinguishing features such as capacity acquisition in multiples of a discrete size, unreserved coaches, standing travel resulting in congestion, and bounds on the total number of coaches and the number of unreserved coaches. Put succinctly, the problem we analyze in this paper can be viewed as a more-general version of the canonical multi-product dynamic-pricing problem.

A majority of revenue management studies in the context of railways use quantity-based techniques. Ciancimino et al. (1999), one of the earliest such work, examines the problem of determining seat allocations to maximize revenue on a single-train line with a single fare-class and multiple legs. A subsequent study by You (2008) analyzes the extension to two fare-classes – the author formulates this problem as a mixed integer nonlinear program and proposes a heuristic. In their overview of revenue management models and approaches in the context of railways, Armstrong and Meissner (2010) extend the two fare-class model in You (2008) to multiple fare-classes. Wang et al. (2016) study a seat-allocation problem with stochastic (discrete random) demand. They construct a variety of seat-allocation policies and numerically compare these policies. Zhai et al. (2018) study the problem of maximizing passenger throughput for multiple lines that serve a set of origin-destination pairs, by optimizing seat allocation. Zhu et al. (2021) study a dynamic capacity-control problem that makes instantaneous and non-changeable seat assignments for tickets sold, in addition to deciding the itineraries open for sale at any time during the sales horizon. They propose a resolving-based policy that achieves uniformly bounded regret under mild assumptions. Konno and Raghunathan (2020) study a data-driven pricing and seat-allocation problem by formulating it as an integer program. There are several studies that use mathematical programming models (in particular, linear integer programming) to formulate dynamic pricing problems and meta-heuristics to generate effective feasible solutions; examples include Hohberger (2020) and Qin et al. (2019). Kamandanipour et al. (2020) use simulation together with a simulated annealing meta-heuristic for a dynamic pricing and capacity allocation model. There are also empirical studies of revenue management in railways that develop models to predict customer choice; for example, Hetrakul and Cirillo (2014) use a latent-class choice model.

2. Model

For the setting described in Section 1, we discuss the demand and supply models in Sections 2.1 and 2.2. Then, Section 2.3 formulates our problem as well as a related deterministic variant.

2.1. Demand Model

Recall from Section 1 that we consider a single day of operations of a railway line, and a preceding sales horizon $[1,T]$ during which tickets for the N itineraries on each of the M trains are sold for travel on the focal day. Also, recall that a passenger with an unreserved ticket for an itinerary can travel on any of the M trains. Our model of the demand-price relationship of the itineraries is general, in the following sense:

- The price of a reserved ticket for itinerary n on any train affects the demand of a reserved ticket for itinerary n on all the trains, and also the demand of an unreserved ticket for itinerary n.

- The price of an unreserved ticket for itinerary n affects the demand of an unreserved ticket for itinerary n and also the demand of reserved tickets for itinerary n on all the trains.

We assume that the price of a ticket (reserved or unreserved) for itinerary n does not affect the demand of a ticket for any other itinerary. In period $t \in \{1,2,...,T\}$, let $p_{n,m}(t)$ be the price of a reserved ticket for itinerary n on train m and let $q_n(t)$ be the price of an unreserved ticket for itinerary n. Let $\vec{p}_n(t) = \big(p_{n,1}(t), ..., p_{n,M}(t)\big)$ denote the vector of prices of a reserved ticket for itinerary n on the M trains. Similar to most discrete-time models in the revenue management literature, we assume that the duration of a period is sufficiently small so that at most one customer arrives during each period. In period t, let $X_{n,m}(t) = 1$ if a customer arrives and purchases a reserved ticket for itinerary n on train m, and $X_{n,m}(t) = 0$ otherwise. Similarly, let $Y_n(t) = 1$ if a customer arrives in period t and purchases an unreserved ticket for itinerary n, and $Y_n(t) = 0$ otherwise $\big($note that $\sum_{m=1}^{M}\sum_{n=1}^{N} X_{n,m}(t) + \sum_{n=1}^{N} Y_n(t) \leqslant 1$, since at most one customer arrives during each period$\big)$. Let $x_{n,m}(t)$ and $y_n(t)$ denote the expectations of $X_{n,m}(t)$ and $Y_n(t)$, respectively; we refer to $x_{n,m}(t)$ and $y_n(t)$ as the *demand rates*. Let $\vec{x}_n(t) = \big(x_{n,1}(t), ..., x_{n,M}(t)\big)$ denote the vector of demand rates of a reserved ticket on itinerary n for the M trains. We note that both $\vec{x}_n(t)$ and $y_n(t)$ depend on $\vec{p}_n(t)$ as well as $q_n(t)$. For better exposition, we will also benefit from another succinct notation for the demand rates: In any period t, let $\vec{\lambda}_n(\vec{p}_n(t), q_n(t))$ denote the demand-rate vector of a ticket for itinerary n as a function of the current price vector of that itinerary; that is,

$$\begin{aligned}\big(\vec{x}_n(t), y_n(t)\big) &= \big(x_{n,1}(t), .., x_{n,M}(t), y_n(t)\big)\\ &= \vec{\lambda}_n\big(\vec{p}_n(t), q_n(t)\big) = \Big(\lambda_{n,1}\big(\vec{p}_n(t), q_n(t)\big), ..., \lambda_{n,M}\big(\vec{p}_n(t), q_n(t)\big), \lambda_n^u\big(\vec{p}_n(t), q_n(t)\big)\Big).\end{aligned}$$

Here, $\lambda_{n,m}$ is the demand function of a reserved ticket for itinerary n on train m for $1 \leqslant m \leqslant M$, and λ_n^u is the demand function of an unreserved ticket for itinerary n. We assume that all the demand functions are stationary[1] and known.

Similar to several other studies in the dynamic pricing literature (see e.g., Gallego and Van Ryzin 1997, Jasin 2014), we assume the following regularity conditions:

1. For itinerary n, let $\mathcal{P}_n^C$ and Λ_n^C denote the set of feasible prices and the set of feasible demand rates, respectively. For all $1 \leqslant n \leqslant N$, both $\mathcal{P}_n^C$ and Λ_n^C are convex, and the demand function $\vec{\lambda}_n(\cdot) : \mathcal{P}_n^C \to \Lambda_n^C$ is bounded and continuously differentiable.

2. For all $1 \leqslant n \leqslant N$, $\vec{\lambda}_n(\cdot)$ has an inverse function denoted by $\vec{\zeta}_n(\cdot) : \Lambda_n^C \to \mathcal{P}_n^C$; that is, $\vec{\lambda}_n\big(\vec{\zeta}_n(\vec{x}_n, y_n)\big) = (\vec{x}_n, y_n)$. Therefore, either the demand rates of the itineraries or, equivalently, their prices can be viewed as decision variables.

3. In any time period, there exists a "null" price vector $\widehat{\vec{p}}_{n,\infty} = (\vec{p}_{n,\infty}, q_{n,\infty})$ for each itinerary n such that $\vec{\lambda}_n(\vec{p}_{n,\infty}, q_{n,\infty}) = \vec{0}$ for all $1 \leqslant n \leqslant N$. Let $\widehat{\vec{p}}_\infty = (\widehat{\vec{p}}_{1,\infty}, ..., \widehat{\vec{p}}_{N,\infty})$ be the price vector that turns off the demand on all the itineraries of all the trains.

4. In any time period, the revenue rate of the firm as a function of the demand rates of the itineraries in that period is defined by:

$$r(\vec{x}_1, ..., \vec{x}_N, y_1, ..., y_N) = \sum_{n=1}^{N} r_n(\vec{x}_n, y_n), \text{ where } r_n(\vec{x}_n, y_n) = (\vec{x}_n, y_n) \cdot \vec{\zeta}_n(\vec{x}_n, y_n)$$

[1] The time-homogeneity assumption on the demand functions is purely for ease of exposition and is not necessary for our results to hold.

is the revenue rate from itinerary n. The function r is bounded, twice differentiable, and jointly concave in the demand rates. The joint-concavity assumption is standard in the literature and holds for many common demand models such as the multinomial logit and nested logit (Li and Huh 2011) models, and the linear demand model (Maglaras and Meissner 2006).

2.2. Supply Model

We now discuss the supply side. The reserved and unreserved coaches of the trains are identical, with each coach consisting of k seats. Let us first consider the supply for reserved tickets.

Reserved Supply: To compute the capacity needed to satisfy the demand for reserved tickets on each train, we need to examine the reserved-ticket demand arrivals on each leg of a train. Let S_ℓ denote the set of all itineraries that use leg ℓ, $1 \leqslant \ell \leqslant L$, and let $\bar{S}_n$ denote the set of all legs used by itinerary n, $1 \leqslant n \leqslant N$. In any period t, let $\overline{X}_{\ell,m}(t) = \sum_{n \in S_\ell} X_{n,m}(t)$. Since we assume that at most one customer arrives in each period, from the definition of $X_{n,m}(t)$, we know that $\overline{X}_{\ell,m}(t)$ takes the value 1 if a customer arrives in period t and purchases a reserved ticket that uses leg ℓ of train m, and takes the value 0, otherwise. The demand rate of $\overline{X}_{\ell,m}(t)$ is given by $\bar{x}_{\ell,m}(t) = \sum_{n \in S_\ell} x_{n,m}(t)$. Now, let $\overline{X}_{\ell,m}[1,t] = \sum_{\hat{t}=1}^{t} \overline{X}_{\ell,m}(\hat{t})$; i.e., $\overline{X}_{\ell,m}[1,t]$ denotes the total number of purchases, until the end of period t, of reserved tickets that use leg ℓ of train m. Therefore, $\max\left\{\overline{X}_{\ell,m}[1,T] : 1 \leqslant \ell \leqslant L\right\}$ denotes the total number of seats required to satisfy the reserved demand on train m. Since travel by standing is not allowed in the reserved coaches, the number of reserved coaches[2] needed on train m is $\left\lceil \frac{\max\left\{\overline{X}_{\ell,m}[1,T]:1 \leqslant \ell \leqslant L\right\}}{k} \right\rceil$, where k is the number of seats in a coach. Thus, the pricing policy of the firm indirectly determines the total number of reserved coaches required on each train. Next, we discuss the supply for unreserved tickets.

Unreserved Supply: Recall the random variable $Y_n(t)$, which equals 1 if a customer arrives in period t and purchases an unreserved ticket for itinerary n, and equals 0 otherwise. Also, recall that an unreserved ticket for an itinerary is not associated with a specific train – the customer can travel on any train on the day of her travel on that itinerary. We assume that a customer who buys an unreserved ticket for itinerary n chooses to travels on train m with a probability $\delta_{n,m}$ that is known to the firm; these probabilities can be estimated from historical sales.

Define the random variable $Y_{n,m}(t)$ as follows: $Y_{n,m}(t) = 1$, if a customer arrives in period t and purchases an unreserved ticket on itinerary n and eventually (on the day of travel) chooses to travel on train m; $Y_{n,m}(t) = 0$, otherwise. Note that the exact realization of $Y_{n,m}(t)$ is known only on the day of travel and not in period t. Recall that the demand rate of an unreserved ticket for itinerary n is $y_n(t)$; thus, the expectation of $Y_{n,m}(t)$ is $\delta_{n,m} \cdot y_n(t)$. Similar to $\overline{X}_{\ell,m}(t)$, we define $\overline{Y}_{\ell,m}(t) = \sum_{n \in S_\ell} Y_{n,m}(t)$. Since at most one customer arrives in any period, $\overline{Y}_{\ell,m}(t)$ equals 1 if a customer who purchases (in period t) an unreserved ticket that uses leg ℓ ends up traveling on

[2] For any $x \in \mathbb{R}_+$, the floor $\lfloor x \rfloor$ is the largest integer smaller than or equal to x. The ceiling $\lceil x \rceil = \lfloor x \rfloor + 1$.

train m, and equals 0 otherwise. The exact realization of $\overline{Y}_{\ell,m}(t)$ will also be known only on the day of travel and not in period t. Let $\overline{Y}_{\ell,m}[1,t] = \sum_{\hat{t}=1}^{t} \overline{Y}_{\ell,m}(\hat{t})$. Thus, among customers who (until period t) buy unreserved tickets that use leg ℓ, the number who choose to travel on train m is $\overline{Y}_{\ell,m}[1,t]$. On the unreserved coaches, passengers are allowed to stand and the firm can therefore sell more tickets than the seats available. Therefore, the number of unreserved coaches is not directly determined by the sale of unreserved tickets, and is a separate decision for the firm. Let $\vec{b}_u = (b_{u,1}, ..., b_{u,M})$, where $b_{u,m}$ denotes the decision variable for the number of unreserved coaches used on train m.

Costs and Constraints: Since standing is allowed on unreserved coaches, the firm's pricing policies, aimed solely at profit maximization, could lead to the number of passengers in a train exceeding the available unreserved capacity. To model the dissatisfaction of unreserved customers who do not find a seat, we include a penalty cost of c_s per leg, on each leg, for each standing passenger[3]. To operate one coach on the railway line, the firm incurs an operational cost of c_o, which includes costs such as per-mile electricity cost, cleaning fee, and maintenance cost. Due to a variety of reasons, including the hauling capacity of electric locomotives, the length of railway platforms, and the length of loop lines that provide bypassing routes, there is an upper bound on the number of coaches that can be used on a train. Let $\overline{b}_m$ denote the maximum number of coaches that can be used on train m. Finally, governmental regulations mandate a minimum number of unreserved coaches on each train. Accordingly, we impose a lower bound $\underline{b}_m$ on the number of unreserved coaches used on train m.

Table 1 summarizes the notation for all our variables and parameters. With all the assumptions and constraints discussed above, we now proceed to provide a precise formulation of our problem.

2.3. Formulation

The firm decides the prices of all the reserved and unreserved tickets in each time period, and the number of unreserved coaches to use on each train, subject to the capacity constraints discussed above. Let Π denote the class of all non-anticipatory policies. For the quantities that we have defined in this section, we will use the superscript π, wherever appropriate, to denote the corresponding quantities under policy $\pi \in \Pi$. The firm's profit-maximization problem can now be written as:

$$\begin{aligned} \mathcal{J}_S = \max_{\pi \in \Pi} \mathbb{E} \Bigg[& \sum_{t=1}^{T} \sum_{n=1}^{N} r_n(\vec{x}_n^{\pi}(t), y_n^{\pi}(t)) - \sum_{m=1}^{M} c_o \left\lceil \frac{\max\left\{\overline{X}_{\ell,m}^{\pi}[1,T] : 1 \leqslant \ell \leqslant L\right\}}{k} \right\rceil \\ & - \sum_{m=1}^{M} c_o b_{u,m}^{\pi} - \sum_{m=1}^{M} \sum_{\ell=1}^{L} c_s \left[\overline{Y}_{\ell,m}^{\pi}[1,T] - b_{u,m}^{\pi} k \right]^+ \Bigg] \end{aligned} \qquad (\mathbb{P}_S)$$

subject to:

[3] While this assumption makes the total penalty cost linear in the number of standing passengers, in Appendix K (supplementary appendix) we show that our analysis and results throughout the paper extend to the case where the total penalty is a convex function of the number of standing passengers.

Parameter	**Definition**
$t \in \{1,2,...,T\}$	Time Period.
$m \in \{1,2,...,M\}$	Train.
$\ell \in \{1,2,...,L\}$	Leg; i.e., a pair of successive stations where a train stops.
$n \in \{1,2,..,N\}$	Itinerary; i.e., a combination of consecutive legs.
S_ℓ	Set of itineraries that use leg ℓ.
$\overline{S}_n$	Set of legs used by itinerary n.
$X_{n,m}(t)$ $\big($resp., $Y_{n,m}(t)\big)$	Random demand of a reserved (resp., unreserved) ticket for itinerary n on train m in period t.
$\overline{X}_{\ell,m}(t)$ $\big($resp., $\overline{Y}_{\ell,m}(t)\big)$	Random demand of a reserved (resp., unreserved) ticket for leg ℓ on train m in period t.
$\overline{X}_{\ell,m}[1,t]$ $\big($resp., $\overline{Y}_{\ell,m}[1,t]\big)$	Total random demand of a reserved (resp., unreserved) ticket for leg ℓ on train m until the end of period t.
$x_{n,m}(t), y_{n,m}(t), \bar{x}_{\ell,m}(t), \bar{y}_{\ell,m}(t)$	Expected rates of $X_{n,m}(t)$, $Y_{n,m}(t)$, $\overline{X}_{\ell,m}(t)$, $\overline{Y}_{\ell,m}(t)$ in period t.
k	Number of seats in a coach of a train.
$\bar{b}_m$	Maximum number of coaches allowed on train m.
$\underline{b}_m$	Minimum number of unreserved coaches to be used on train m.
c_o	Cost of operating a coach.
c_s	Per-person penalty cost for travel by standing on a leg.
Variable	**Definition**
$p_{n,m}(t)$	Price of a reserved ticket for itinerary n on train m in period t.
$q_n(t)$	Price of an unreserved ticket for itinerary n in period t.
$b_{u,m}$	Number of unreserved coaches used on train m.

Table 1 Our main notation

$$\overline{X}^{\pi}_{\ell,m}[1,T] + b^{\pi}_{u,m} k \leqslant \bar{b}_m k \quad \text{(a.s.)} \quad \forall\ 1 \leqslant \ell \leqslant L,\ 1 \leqslant m \leqslant M, \tag{1}$$

$$b^{\pi}_{u,m} \geqslant \underline{b}_m, \quad b^{\pi}_{u,m} \in \mathbb{N} \quad \forall\ 1 \leqslant m \leqslant M, \tag{2}$$

$$(\vec{x}^{\pi}_n(t), y^{\pi}_n(t)) \in \Lambda^C_n \quad \forall\ 1 \leqslant n \leqslant N.$$

We refer to the above problem as $\mathbb{P}_S$ and its optimal expected profit as $\mathcal{J}_S$. The first term in the objective function of $\mathbb{P}_S$ represents the total expected revenue, while the second and third terms represent the expected operational costs from reserved and unreserved coaches, respectively. The last term represents the expected total penalty cost incurred due to standing passengers. Recall that there is an upper bound $\bar{b}_m$ on the number of coaches used on train m; mathematically,

$$\left\lceil \frac{\max\left\{\overline{X}^{\pi}_{\ell,m}[1,T] : 1 \leqslant \ell \leqslant L\right\}}{k} \right\rceil + b^{\pi}_{u,m} \leqslant \bar{b}_m \quad \text{(a.s.)} \quad \forall\ 1 \leqslant m \leqslant M. \tag{3}$$

It is easy to verify that constraints (3) are equivalent to constraints (1). Constraints (2) model the requirement that there should be least $\underline{b}_m$ unreserved coaches on train m.

Problem $\mathbb{P}_S$ is more general than the canonical multi-product dynamic-pricing problem in the revenue management literature; see e.g., Gallego and Van Ryzin (1997). It is therefore well-understood that obtaining an optimal solution to $\mathbb{P}_S$ is intractable, in general. Therefore, we aim to develop efficient policies with attractive performance guarantees. To this end, we first define the following deterministic optimization problem $\mathbb{P}_D$ that is obtained from problem $\mathbb{P}_S$ by dropping the ceiling function and the integrality restriction on $b^{\pi}_{u,m}$, replacing the stochastic quantities by their expectations, and making the demand-rate decision variables time-invariant.

$$\mathcal{J}_D = \max_{\vec{x},\vec{y},\vec{b}_u} \Bigg[\sum_{n=1}^{N} r_n(\vec{x}_n, y_n)\ T - \sum_{m=1}^{M} \frac{c_o T}{k} \max\{\bar{x}_{\ell,m} : 1 \leqslant \ell \leqslant L\} - \sum_{m=1}^{M} c_o b_{u,m}$$
$$- \sum_{m=1}^{M} \sum_{\ell=1}^{L} c_s \Big[\bar{y}_{\ell,m} T - b_{u,m} k \Big]^{+} \Bigg] \qquad (\mathbb{P}_D)$$

subject to:

$$\bar{x}_{\ell,m} T + b_{u,m}\, k \leqslant \bar{b}_m\, k \quad \forall\ 1 \leqslant \ell \leqslant L,\ 1 \leqslant m \leqslant M,$$
$$b_{u,m} \geqslant \underline{b}_m, \quad b_{u,m} \in \mathbb{R}_+\ \forall\ 1 \leqslant m \leqslant M,$$
$$\bar{x}_{\ell,m} = \sum_{n \in S_\ell} x_{n,m}, \quad \bar{y}_{\ell,m} = \sum_{n \in S_\ell} \delta_{n,m} \cdot y_n \quad \forall\ 1 \leqslant \ell \leqslant L,\ 1 \leqslant m \leqslant M,$$
$$(\vec{x}_n, y_n) \in \Lambda_n^C \quad \forall\ 1 \leqslant n \leqslant N. \qquad (4)$$

The following result establishes an upper bound on the optimal profit of problem $\mathbb{P}_S$. The proofs of all the technical results are in the appendix.

Theorem 1. *The optimal profit of $\mathbb{P}_S$ is at most that of $\mathbb{P}_D$, i.e., $\mathcal{J}_S \leqslant \mathcal{J}_D$.*

Next, we use the solution of $\mathbb{P}_D$ to define our policies for $\mathbb{P}_S$. Section 3 (resp., Section 4) develops our static-pricing (resp., dynamic-pricing) policies.

3. Static Pricing Policies

We first develop our static-price flexible-capacity policy and show that when demand-arrival rates and capacities are both scaled by a factor $\kappa \in \mathbb{N}$, the optimality gap of this policy scales proportional to $\sqrt{\kappa}$ whereas the profit of the upper-bound problem $\mathbb{P}_D$ scales proportional to κ, thereby proving that this policy is asymptotically optimal (Theorem 2). Subsequently, we develop our static-price fixed-capacity policy and establish the same results for that policy (Theorem 3). Later, in Section 5, we also numerically evaluate the two policies and demonstrate the superior performance of the flexible-capacity policy.

3.1. Static Pricing Policy with Flexible Capacity

We construct our static-price flexible-capacity policy using an optimal solution to the deterministic problem $\mathbb{P}_D$. Let $\left(\vec{x}^*, \vec{y}^*, \vec{b}^*_u\right)$ be an optimal solution to problem $\mathbb{P}_D$, where $\vec{x}^* = \left(\vec{x}^*_1, ..., \vec{x}^*_N\right)$, $\vec{x}^*_n = \left(x^*_{n,1}, ..., x^*_{n,M}\right)$ for $1 \leqslant n \leqslant N$, $\vec{y}^* = \left(y^*_1, ..., y^*_N\right)$, and $\vec{b}^*_u = \left(b^*_{u,1}, ..., b^*_{u,M}\right)$. For $1 \leqslant \ell \leqslant L$ and $1 \leqslant m \leqslant M$, let $\bar{x}^*_{\ell,m} = \sum_{n \in S_\ell} x^*_{n,m}$ and $\bar{y}^*_{\ell,m} = \sum_{n \in S_\ell} \delta_{n,m} \cdot y^*_n$. Also, let $\left(\vec{p}^*_n, q^*_n\right) = \vec{\zeta}_n\left(\vec{x}^*_n, y^*_n\right)$ be the corresponding optimal prices of a reserved ticket and an unreserved ticket for itinerary n. Below, we first describe the main idea behind our policy and then define the policy precisely.

For each n, $1 \leqslant n \leqslant N$, and for each m, $1 \leqslant m \leqslant M$, we maintain the price of a reserved ticket for itinerary n on train m at $p^*_{n,m}$ and the price of an unreserved ticket for itinerary n at q^*_n, until the earliest time period in which, for some $\hat{\ell}$ and $\hat{m}$, the number of reserved seats sold on leg $\hat{\ell}$ of train $\hat{m}$ equals a carefully-chosen upper bound on the reserved capacity of train $\hat{m}$; $1 \leqslant \hat{\ell} \leqslant L$ and $1 \leqslant \hat{m} \leqslant M$. After that time period, we stop the sales of all the reserved and unreserved tickets for any itinerary n that uses leg $\hat{\ell}$, i.e., for all $n \in S_{\hat{\ell}}$. For any itinerary n that does not use leg $\hat{\ell}$, i.e., $n \notin S_{\hat{\ell}}$, we continue to maintain the price of a reserved ticket on train m at $p^*_{n,m}$ and the price of an unreserved ticket at q^*_n, until a time period in which the number of reserved seats sold for some leg $\tilde{\ell} \neq \hat{\ell}$ of a train reaches the upper bound on the reserved capacity of that train; from this time period, we stop the sales of all the reserved and unreserved tickets for any itinerary n that uses leg $\tilde{\ell}$. We proceed in a similar manner until the sales for all the itineraries have stopped or the end of the sales horizon is reached.

At the end of the sales horizon, the capacity decisions are made as follows. On each train, we choose the minimum number of reserved coaches sufficient to provide the required number of reserved seats. The vector of the number of unreserved coaches on the trains is jointly chosen so that, over all the trains, the sum of operational costs and penalty cost from standing is minimized.

We now precisely define our policy, which we denote by π_1. We first describe the pricing decisions and then the capacity decisions.

Algorithm for Static-Price Flexible-Capacity Policy

- **Pricing decisions:**

1. In time period $t = 1$, for itinerary $n, 1 \leqslant n \leqslant N$, choose $\left(\vec{p}^{\pi_1}_n(1), q^{\pi_1}_n(1)\right) = (\vec{p}^*_n, q^*_n)$.

2. In time periods $t > 1$, for all itineraries $(1 \leqslant n \leqslant N)$, do the following: For itinerary n, if $\overline{X}^{\pi_1}_{\ell,m}[1, t-1] = k\left(\bar{b}_m - \lfloor b^*_{u,m} \rfloor\right)$ for some $\ell \in \bar{S}_n$ and $1 \leqslant m \leqslant M$, then choose the null price vector, i.e., $\left(\vec{p}^{\pi_1}_n(t), q^{\pi_1}_n(t)\right) = \left(\vec{p}_{n,\infty}, q_{n,\infty}\right)$; otherwise, choose $\left(\vec{p}^{\pi_1}_n(t), q^{\pi_1}_n(t)\right) = (\vec{p}^*_n, q^*_n)$.

- **Capacity decisions:**

3. The capacity to be used on each train is determined from the total demand generated at the end of the sales horizon. Define

$$b^{\pi_1}_{r,m} = \left\lceil \frac{\max\left\{\overline{X}^{\pi_1}_{\ell,m}[1,T] : 1 \leqslant \ell \leqslant L\right\}}{k} \right\rceil \ \forall 1 \leqslant m \leqslant M.$$

On train m $(1 \leqslant m \leqslant M)$, use $b^{\pi_1}_{r,m}$ reserved coaches.

4. For all $1 \leqslant m \leqslant M$, let $\mathcal{B}^{\pi_1}_m = \left[\underline{b}_m, (\bar{b}_m - b^{\pi_1}_{r,m})\right]$. Let $\mathcal{B}^{\pi_1} = \mathcal{B}^{\pi_1}_1 \times \cdots \times \mathcal{B}^{\pi_1}_M$. Define

$$\vec{b}^{\pi_1}_u = \left(b^{\pi_1}_{u,1}, ..., b^{\pi_1}_{u,M}\right) := \underset{\vec{b} \in \mathcal{B}^{\pi_1}}{\arg\min} \sum_{m=1}^{M} c_o b_m + \sum_{m=1}^{M} \sum_{\ell=1}^{L} c_s \mathbb{E}\left[\overline{Y}^{\pi_1}_{\ell,m}[1,T] - b_m k\right]^+. \qquad (\mathbb{Q}^{\pi_1})$$

Problem $\mathbb{Q}^{\pi_1}$ is a convex program. On train m $(1 \leqslant m \leqslant M)$, use $\left\lfloor b^{\pi_1}_{u,m} \right\rfloor$ unreserved coaches. ■

It is easy to verify that policy π_1 is feasible for problem $\mathbb{P}_S$. We now obtain the performance guarantee offered by policy π_1. To this end, consider a sequence of scaled problems $\mathbb{P}_{S,\kappa}$ indexed by $\kappa \in \mathbb{N}$, where in $\mathbb{P}_{S,\kappa}$ the demand arrival rates and the capacities are all scaled by $\kappa > 0$ (in other words, the number of time periods in the sales horizon is scaled by κ and the maximum number of total coaches allowed on train m is $\kappa \cdot \bar{b}_m$). Throughout this section, for several quantities that we have already defined, we will use the subscript κ to represent their corresponding versions in the κ-scaled system; for example, $\mathcal{J}_{S,\kappa}$ denotes the optimal profit in problem $\mathbb{P}_{S,\kappa}$. Let $\mathcal{J}_S^{\pi_1}$ and $\mathcal{J}_{S,\kappa}^{\pi_1}$ be the expected profit obtained by the static-price flexible-capacity policy in problem $\mathbb{P}_S$ and $\mathbb{P}_{S,\kappa}$, respectively. The following result establishes a bound on the performance of policy π_1.

THEOREM 2. *The difference between the optimal profit of problem $\mathbb{P}_{S,\kappa}$ and the profit obtained by the static-price flexible-capacity policy π_1 in $\mathbb{P}_{S,\kappa}$ is $\mathcal{O}(\sqrt{\kappa})$, i.e., $\mathcal{J}_{S,\kappa} - \mathcal{J}_{S,\kappa}^{\pi_1} = \mathcal{O}(\sqrt{\kappa})$. Consequently, π_1 is asymptotically optimal, i.e., the ratio $\frac{\mathcal{J}_{S,\kappa}^{\pi_1}}{\mathcal{J}_{S,\kappa}} \to 1$ as $\kappa \to \infty$.*

3.2. Static Pricing Policy with Fixed Capacity

In this section, we define our static-pricing policy in which the capacity decisions on all the trains are made at the beginning of the sales horizon. We also establish that this policy offers the same asymptotic convergence rate as that in Theorem 2 for the static-price flexible-capacity policy.

The construction of our static-price fixed-capacity policy is also based on the optimal solution to the deterministic problem $\mathbb{P}_D$. Recall the optimal solution $\left(\vec{x}^*, \vec{y}^*, \vec{b}_u^*\right)$ to problem $\mathbb{P}_D$, where $\vec{x}^* = \left(\vec{x}_1^*, ..., \vec{x}_N^*\right)$, $\vec{x}_n^* = \left(x_{n,1}^*, ..., x_{n,M}^*\right)$ for $1 \leqslant n \leqslant N$, $\vec{y}^* = \left(y_1^*, ..., y_N^*\right)$, and $\vec{b}_u^* = \left(b_{u,1}^*, ..., b_{u,M}^*\right)$. Let $\left(\vec{p}_n^*, q_n^*\right) = \vec{\zeta}_n\left(\vec{x}_n^*, y_n^*\right)$ be the corresponding optimal prices of reserved and unreserved tickets on itinerary n for all $1 \leqslant n \leqslant N$. Also, recall the quantities $\bar{x}_{\ell,m}^* = \sum_{n \in S_\ell} x_{n,m}^*$ and $\bar{y}_{\ell,m}^* = \sum_{n \in S_\ell} \delta_{n,m} \cdot y_n^*$ for all $1 \leqslant \ell \leqslant L$ and $1 \leqslant m \leqslant M$. In the deterministic problem $\mathbb{P}_D$, observe that $\bar{x}_{\ell,m}^* T$ and $\bar{y}_{\ell,m}^* T$ represent the total number of reserved and unreserved customers, respectively, who travel on leg ℓ of train m. Further, $\max\{\bar{x}_{\ell,m}^* T : 1 \leqslant \ell \leqslant L\}$ represents the total number of seats required to satisfy the reserved demand on train m. Therefore, in problem $\mathbb{P}_D$, the number of coaches required to satisfy the reserved demand on train m is equal to $\left\lceil \frac{\max\{\bar{x}_{\ell,m}^* : 1 \leqslant \ell \leqslant L\} T}{k} \right\rceil$.

In our static-price fixed-capacity policy, which we denote henceforth by π_2, we choose the number of coaches on each train based on the capacity consumption in problem $\mathbb{P}_D$ by its optimal solution $\left(\vec{x}^*, \vec{y}^*, \vec{b}_u^*\right)$. Let $b_{u,m}^{\pi_2} = \lfloor b_{u,m}^* \rfloor$ and

$$b_{r,m}^{\pi_2} = \left\lceil \frac{\max\{\bar{x}_{\ell,m}^* : 1 \leqslant \ell \leqslant L\} T}{k} \right\rceil.$$

Algorithm for Static-Price Fixed-Capacity Policy

- **Capacity decisions:**

1. On train m $(1 \leqslant m \leqslant M)$, use $b_{r,m}^{\pi_2}$ reserved and $b_{u,m}^{\pi_2}$ unreserved coaches.

- **Pricing decisions:**

2. In time period $t = 1$, for itinerary n, choose $\big(\vec{p}_n^{\pi_2}(1), q_n^{\pi_2}(1)\big) = (\vec{p}_n^*, q_n^*)$, $1 \leqslant n \leqslant N$.

3. In time periods $t > 1$, for all itineraries ($1 \leqslant n \leqslant N$), do the following: For itinerary n, if $\overline{X}_{\ell,m}^{\pi_2}[1, t-1] = k\, b_{r,m}^{\pi_2}$ for some $\ell \in \bar{S}_n$ and $1 \leqslant m \leqslant M$, then choose the null price vector[4], i.e., $\big(\vec{p}_n^{\pi_2}(t), q_n^{\pi_2}(t)\big) = \big(\vec{p}_{n,\infty}, q_{n,\infty}\big)$; otherwise, choose $\big(\vec{p}_n^{\pi_2}(t), q_n^{\pi_2}(t)\big) = (\vec{p}_n^*, q_n^*)$. ■

Note that the capacity decisions in this policy are made at the beginning of the sales horizon and remain unchanged thereafter. Observe that $b_{r,m}^{\pi_2} + b_{u,m}^{\pi_2} \leqslant \bar{b}_m$, the upper bound on the total number of coaches on train m, and $b_{u,m}^{\pi_2} \geqslant \underline{b}_m$, the minimum number of unreserved coaches on train m, $1 \leqslant m \leqslant M$. Thus, policy π_2 is feasible for problem $\mathbb{P}_S$.

We now establish the performance guarantee offered by policy π_2. To this end, we again consider a κ-scaled system; see the paragraph preceding Theorem 2. Let $\mathcal{J}_{S,\kappa}^{\pi_2}$ be the expected profit obtained by policy π_2 in problem $\mathbb{P}_{S,\kappa}$.

THEOREM 3. *The difference between the optimal profit of problem $\mathbb{P}_{S,\kappa}$ and the profit obtained by the static-price fixed-capacity policy π_2 in $\mathbb{P}_{S,\kappa}$ is $\mathcal{O}(\sqrt{\kappa})$, i.e., $\mathcal{J}_{S,\kappa} - \mathcal{J}_{S,\kappa}^{\pi_2} = \mathcal{O}(\sqrt{\kappa})$. Consequently, π_2 is asymptotically optimal, i.e., the ratio $\frac{\mathcal{J}_{S,\kappa}^{\pi_2}}{\mathcal{J}_{S,\kappa}} \to 1$ as $\kappa \to \infty$.*

REMARK 1. **(Analysis of a Discrete Variant)** In practice, the price of a train itinerary is often restricted to a few convenient values for consumers' ease of understanding and familiarity. Motivated by this observation, in Appendix B, we study a variant of problem $\mathbb{P}_S$ in which the possible prices of reserved and unreserved tickets on each itinerary are restricted to discrete and finite sets. We develop an efficient policy for this problem and establish an attractive guarantee on the performance of the policy (Theorem **??**). ■

We now turn our attention to dynamic policies for problem $\mathbb{P}_S$.

4. Dynamic Pricing Policies

In this section, we develop two dynamic-pricing policies for problem $\mathbb{P}_S$ and establish their asymptotic optimality. As with our static policies, the fundamental difference between our two dynamic policies is the timing of the capacity decisions. We first develop our dynamic-price flexible-capacity policy and show that when the demand arrival rates and the capacities are both scaled by a factor $\kappa \in \mathbb{N}$, the optimality gap of this policy scales proportional to $\log \kappa$, which also establishes that this policy is asymptotically optimal (Theorem 4). Subsequently, we develop our dynamic-price fixed-capacity policy and establish the same results for that policy as well (Theorem 5).

4.1. Dynamic Pricing Policy with Flexible Capacity

Our dynamic-price flexible-capacity policy also uses the optimal solution to problem $\mathbb{P}_D$. Recall the optimal solution $\big(\vec{x}^*, \vec{y}^*, \vec{b}_u^*\big)$ to problem $\mathbb{P}_D$, where $\vec{x}^* = \big(\vec{x}_1^*, \ldots, \vec{x}_N^*\big)$, $\vec{x}_n^* = \big(x_{n,1}^*, \ldots, x_{n,M}^*\big)$ for

[4] Note that the switching criteria to the null price vector in π_2 is different from that in π_1.

$1 \leqslant n \leqslant N$, $\vec{y}^* = \left(y_1^*, ..., y_N^*\right)$, and $\vec{b}_u^* = \left(b_{u,1}^*, ..., b_{u,M}^*\right)$. We now define another quantity that is useful in the development of our policy. Under any policy π, we define $\Delta_{n,m}^{\pi}(t) = X_{n,m}^{\pi}(t) - x_{n,m}^{\pi}(t)$ and $\widehat{\Delta}_n^{\pi}(t) = Y_n^{\pi}(t) - y_n^{\pi}(t)$ for all $1 \leqslant t \leqslant T$, $1 \leqslant n \leqslant N$, and $1 \leqslant m \leqslant M$. In words, $\Delta_{n,m}^{\pi}(t)$ and $\widehat{\Delta}_n^{\pi}(t)$ represent the difference between the realized and expected demand, in period t, of reserved tickets for itinerary n on train m, and unreserved tickets for itinerary n, respectively. As in Jasin (2014), we make the following assumption that is useful in defining our policy:

Assumption 1. *The deterministic problem $\mathbb{P}_D$ has an interior optimum. That is, for $1 \leqslant n \leqslant N$, there exist $\vec{\underline{\psi}}_n = \left(\underline{\psi}_{n,1}, ..., \underline{\psi}_{n,M}\right) > \vec{0}$, $\vec{\overline{\psi}}_n = \left(\overline{\psi}_{n,1}, ..., \overline{\psi}_{n,M}\right) > \vec{0}$, $\hat{\underline{\psi}}_n > 0$, and $\hat{\overline{\psi}}_n > 0$ such that*

$$\left[\left(\vec{x}_n^*, y_n^*\right) - \left(\vec{\underline{\psi}}_n, \hat{\underline{\psi}}_n\right), \left(\vec{x}_n^*, y_n^*\right) + \left(\vec{\overline{\psi}}_n, \hat{\overline{\psi}}_n\right)\right] \in \Lambda_n^C,$$

where Λ_n^C is the convex set of feasible demand-rate vectors for any itinerary n.

The above assumption is useful in the construction of both our dynamic-price flexible- and fixed-capacity policies. Specifically, this assumption ensures that the periodic adjustments to the demand rates are feasible; i.e., lie in the interior of Λ^C. We now define our policy, which we denote by π_3. We first describe our pricing decisions followed by our capacity decisions.

Algorithm for Dynamic-Price Flexible-Capacity Policy

- **Pricing decisions:**

1. In time period $t = 1$, for itinerary n, choose $\left(\vec{p}_n^{\pi_3}(1), q_n^{\pi_3}(1)\right) = \left(\vec{p}_n^*, q_n^*\right) = \vec{\zeta}_n\left(\vec{x}_n^*, y_n^*\right)$, $1 \leqslant n \leqslant N$.

2. In time periods $t > 1$, for all itineraries ($1 \leqslant n \leqslant N$) do the following: For itinerary n, if $\overline{X}_{\ell,m}^{\pi_3}[1, t-1] = k\left(\overline{b}_m - b_{u,m}^*\right)$ for some $\ell \in \bar{S}_n$ and $1 \leqslant m \leqslant M$, then choose the null price vector, i.e., $\left(\vec{p}_n^{\pi_3}(t), q_n^{\pi_3}(t)\right) = \left(\vec{p}_{n,\infty}, q_{n,\infty}\right)$; otherwise, choose $\left(\vec{p}_n^{\pi_3}(t), q_n^{\pi_3}(t)\right) = \vec{\zeta}_n\left(\vec{\tilde{x}}_n(t), \tilde{y}_n(t)\right)$. Here, $\left(\vec{\tilde{x}}_n(t), \tilde{y}_n(t)\right)$ is defined as follows:

$$\tilde{x}_{n,m}(t) = \begin{cases} x_{n,m}^* - \sum_{s=1}^{t-1} \dfrac{\Delta_{n,m}^{\pi_3}(s)}{T-s}, & \text{if } -\overline{\psi}_{n,m} < \sum_{s=1}^{t-1} \dfrac{\Delta_{n,m}^{\pi_3}(s)}{T-s} < \underline{\psi}_{n,m} \\ 0, & \text{otherwise,} \end{cases}$$

$\vec{\tilde{x}}_n(t) = \left(\tilde{x}_{n,1}(t), ..., \tilde{x}_{n,M}(t)\right)$, and

$$\tilde{y}_n(t) = \begin{cases} y_n^* - \sum_{s=1}^{t-1} \dfrac{\hat{\Delta}_n^{\pi_3}(s)}{T-s}, & \text{if } -\hat{\overline{\psi}}_n < \sum_{s=1}^{t-1} \dfrac{\hat{\Delta}_n^{\pi_3}(s)}{T-s} < \hat{\underline{\psi}}_n \\ 0, & \text{otherwise.} \end{cases}$$

- **Capacity decisions:**

3. The capacity to be used on each train is determined from the total demand generated at the end of the sales horizon. Define

$$b_{r,m}^{\pi_3} = \left\lceil \frac{\max\left\{\overline{X}_{\ell,m}^{\pi_3}[1,T] : 1 \leqslant \ell \leqslant L\right\}}{k} \right\rceil \; \forall 1 \leqslant m \leqslant M.$$

On train m $(1 \leqslant m \leqslant M)$, use $b^{\pi_3}_{r,m}$ reserved coaches.

4. For all $1 \leqslant m \leqslant M$, let $\mathcal{B}^{\pi_3}_m = \left[\underline{b}_m, (\bar{b}_m - b^{\pi_3}_{r,m})\right]$. Let $\mathcal{B}^{\pi_3} = \mathcal{B}^{\pi_3}_1 \times \cdots \times \mathcal{B}^{\pi_3}_M$. Define

$$\vec{b}^{\pi_3}_u = \left(b^{\pi_3}_{u,1}, ..., b^{\pi_3}_{u,M}\right) := \underset{\vec{b} \in \mathcal{B}^{\pi_3}}{\arg\min} \sum_{m=1}^{M} c_o b_m + \sum_{m=1}^{M} \sum_{\ell=1}^{L} c_s \mathbb{E}\Big[\overline{Y}^{\pi_3}_{\ell,m}[1,T] - b_m k\Big]^+. \qquad (\mathbb{Q}^{\pi_3})$$

Problem $\mathbb{Q}^{\pi_3}$ is a convex program. On train m $(1 \leqslant m \leqslant M)$, use $\lfloor b^{\pi_3}_{u,m} \rfloor$ unreserved coaches. ■

In the pricing definitions above, the adjustments in the demand rates are similar to those in the linear rate-control policies developed in Besbes and Maglaras (2012) and Jasin (2014). For any itinerary on a train, the demand rate used in a time period is equal to the demand rate in the previous period minus the difference between the realized and expected demand used in the previous period, divided by the remaining duration in the sales horizon. The intuition behind this policy is as follows: if the realized demand in the previous period is higher (lower) than expected, then we decrease (increase) the demand rate used in the current period; the increase or decrease is scaled based on the remaining duration of the sales horizon. These adjustments in the demand rates try to maintain the sales of each itinerary, on every sample path, to be close to that in the optimal solution to the upper-bound problem $\mathbb{P}_D$.

It can be verified that policy π_3 is feasible for $\mathbb{P}_S$. We now proceed to establish the performance guarantee offered by this policy. As before, we consider a κ-scaled system in which the demand (the number of time periods) and capacity are proportionally scaled by κ; we refer to the corresponding problem as $\mathbb{P}_{S,\kappa}$. Without loss of generality, we assume that $T = 1$. Therefore, the number of time periods in the κ-scaled system is equal to κ. Recall that for our analysis of the static policies, we first obtained performance guarantees for our policies in the unscaled system and then analyzed how these guarantees scale in the κ-scaled system. However, in this section, for convenience of analysis, we directly work with the κ-scaled system. Let $\mathcal{J}^{\pi_3}_{S,\kappa}$ be the expected profit obtained by policy π_3 in problem $\mathbb{P}_{S,\kappa}$. The result below bounds the optimality gap of policy π_3:

THEOREM 4. *The difference between the optimal profit of problem $\mathbb{P}_{S,\kappa}$ and the profit obtained by the dynamic-price flexible-capacity policy π_3 in $\mathbb{P}_{S,\kappa}$ is $\mathcal{O}(\log \kappa)$, i.e., $\mathcal{J}_{S,\kappa} - \mathcal{J}^{\pi_3}_{S,\kappa} = \mathcal{O}(\log \kappa)$. Consequently, π_3 is asymptotically optimal, i.e., the ratio $\frac{\mathcal{J}^{\pi_3}_{S,\kappa}}{\mathcal{J}_{S,\kappa}} \to 1$ as $\kappa \to \infty$.*

REMARK 2. **(Restricted Flexible-Capacity Policies)** In practice, even if capacity decisions can be delayed until the end of the sales horizon, there could be restrictions on the number of last-minute changes in the number of coaches, due to limited availability of resources. In Appendix C, we examine restricted flexible-capacity policies in which the number of reserved coaches used on each train is decided at the start of the sales horizon and this number is then allowed to increase or decrease by at most one coach at the end of the sales horizon.

4.2. Dynamic Pricing Policy with Fixed Capacity

In this section, we define our dynamic-pricing policy in which the capacity decisions are made at the beginning of the sales horizon. We establish that this policy also offers the same convergence rate as that of the dynamic-price flexible-capacity policy we analyzed in the previous subsection.

The construction of our dynamic-price fixed-capacity policy, which we denote by π_4, is also based on the optimal solution $\left(\vec{x}^*, \vec{y}^*, \vec{b}_u^*\right)$ to problem $\mathbb{P}_D$. Let $b_{u,m}^{\pi_4} = \lfloor b_{u,m}^* \rfloor$ and

$$b_{r,m}^{\pi_4} = \left\lceil \frac{\max\{\bar{x}_{\ell,m}^* : 1 \leqslant \ell \leqslant L\} T}{k} \right\rceil.$$

Algorithm for Dynamic-Price Fixed-Capacity Policy

• **Capacity decisions:**

1. On train m $(1 \leqslant m \leqslant M)$, use $b_{r,m}^{\pi_4}$ reserved and $b_{u,m}^{\pi_4}$ unreserved coaches.

• **Pricing decisions:**

2. In time period $t = 1$, for all itineraries $(1 \leqslant n \leqslant N)$, choose $\left(\vec{p}_n^{\pi_4}(1), q_n^{\pi_4}(1)\right) = (\vec{p}_n^*, q_n^*) = \vec{\zeta}_n\left(\vec{x}_n^*, y_n^*\right)$.

3. In time period $t > 1$, for all itineraries $(1 \leqslant n \leqslant N)$ do the following: For itinerary n, if we have $\overline{X}_{\ell,m}^{\pi_4}[1, t-1] = k\, b_{r,m}^{\pi_4}$ for some $\ell \in \bar{S}_n$ and $1 \leqslant m \leqslant M$, then choose the null price vector, i.e., $\left(\vec{p}_n^{\pi_4}(t), q_n^{\pi_4}(t)\right) = \left(\vec{p}_{n,\infty}, q_{n,\infty}\right)$; otherwise, choose $\left(\vec{p}_n^{\pi_4}(t), q_n^{\pi_4}(t)\right) = \vec{\zeta}_n\left(\vec{\tilde{x}}_n(t), \tilde{y}_n(t)\right)$. Here, $\left(\vec{\tilde{x}}_n(t), \tilde{y}_n(t)\right)$ is defined as follows:

$$\tilde{x}_{n,m}(t) = \begin{cases} x_{n,m}^* - \sum_{s=1}^{t-1} \frac{\Delta_{n,m}^{\pi_4}(s)}{T-s}, & \text{if } -\overline{\psi}_{n,m} < \sum_{s=1}^{t-1} \frac{\Delta_{n,m}^{\pi_4}(s)}{T-s} < \underline{\psi}_{n,m} \\ 0, & \text{otherwise,} \end{cases}$$

$$\tilde{y}_n(t) = \begin{cases} y_n^* - \sum_{s=1}^{t-1} \frac{\hat{\Delta}_n^{\pi_4}(s)}{T-s}, & \text{if } -\hat{\overline{\psi}}_n < \sum_{s=1}^{t-1} \frac{\hat{\Delta}_n^{\pi_4}(s)}{T-s} < \hat{\underline{\psi}}_n \\ 0, & \text{otherwise,} \end{cases}$$

and $\vec{\tilde{x}}_n(t) = \left(\tilde{x}_{n,1}(t), ..., \tilde{x}_{n,M}(t)\right)$. ■

Note that the capacity decisions in policy π_4 are made at the beginning of the sales horizon and remain unchanged thereafter. We now proceed to establish the performance guarantee offered by policy π_4. To this end, we again consider the κ-scaled system. Let $\mathcal{J}_{S,\kappa}^{\pi_4}$ be the expected profit obtained by policy π_4 in problem $\mathbb{P}_{S,\kappa}$. The proof of the following result uses arguments similar to those in the proof of Theorem 4; a summary of the main steps is provided in the appendix.

Theorem 5. *The difference between the optimal profit of problem $\mathbb{P}_{S,\kappa}$ and the profit obtained by the dynamic-price fixed-capacity policy π_4 in $\mathbb{P}_{S,\kappa}$ is $\mathcal{O}(\log \kappa)$, i.e., $\mathcal{J}_{S,\kappa} - \mathcal{J}_{S,\kappa}^{\pi_4} = \mathcal{O}(\log \kappa)$. Consequently, π_4 is asymptotically optimal, i.e., the ratio $\frac{\mathcal{J}_{S,\kappa}^{\pi_4}}{\mathcal{J}_{S,\kappa}} \to 1$ as $\kappa \to \infty$.*

Based on our work with a railway company in Japan, we now proceed to numerically demonstrate the attractive performance of our policies on a test suite of instances generated from the real-world operations of a popular high-speed railway line and discuss the related insights.

5. Numerical Experience and Insights

The network of Shinkansen high-speed trains (colloquially, "bullet trains") is a major means of transportation in Japan. For our numerical study, we focus on the Tokyo-Shinosaka railway line, one of the busiest Shinkansen lines. Five different types of trains operate on this line; we focus on the "Nozomi" type, which is the fastest and attracts the most demand. Before we discuss our test bed, we recall the main parameters of our model: (i) The number of stations on the railway line $(L+1)$, the number of trains (M) operating on the line, and the number of seats in a coach (k). (ii) The parameters related to the demand function $\lambda_{n,m}$ for reserved tickets on itinerary n of train m, and the demand function λ_n^u for unreserved tickets on itinerary n; $n = 1, 2, \ldots, N; m = 1, 2, \ldots, M$. (iii) The cost for operating a coach (c_o), the per-person penalty cost for travelling on a leg of the journey by standing (c_s), the maximum number of coaches on each train $(\bar{b}_m, m = 1, 2, \ldots, M)$, and the minimum number of unreserved coaches on each train $(\underline{b}_m, m = 1, 2, \ldots, M)$.

5.1. Test Bed

We first examine the performance of our policies on three "full-scale" instances that correspond to the actual parameters of the Tokyo-Shinosaka line (Section 5.1.1). Next, we examine the performance on three full-scale instances based on Yan et al. (2022) (Section 5.1.2). Then, to confirm robustness of the performance and to assess the impact of the penalty cost (for standing passengers) on congestion in unreserved coaches, we also use a larger test bed of relatively smaller instances (Section 5.1.3).

5.1.1. Full-Scale Instances The Tokyo-Shinosaka line connects the following five stations, in that order: Tokyo, Shinyokohoma, Nagoya, Kyoto, and Shinosaka. Thus, the number of legs on this line is $L = 4$, and the number of different itineraries is $N = 10$. Our interactions with practitioners suggest that significant demand substitution occurs only between the trains that operate within the same standard hour; e.g., between trains that operate from 9am–10am or between trains that operate from 11am–12noon. On average, about 5 Nozomi trains operate on this line in a standard hour; we therefore set $M = 5$. The number of seats in a coach is $k = 87$. We use an MNL demand function to model the demand-price relationship for each itinerary. Accordingly, the demand function of a reserved ticket for itinerary n on train m; $1 \leqslant n \leqslant 10$, $1 \leqslant m \leqslant 5$, is given by

$$\lambda_{n,m}(\vec{p}_n, q_n) = \frac{\exp(\alpha_{n,m} - \beta_{n,m} \cdot p_{n,m})}{1 + \sum_{m=1}^{M} \exp(\alpha_{n,m} - \beta_{n,m} \cdot p_{n,m}) + \exp(\alpha_n^u - \beta_n^u \cdot q_n)} \Gamma_n,$$

and the demand function of an unreserved ticket for itinerary n, $1 \leqslant n \leqslant 10$, is given by

$$\lambda_n^u(\vec{p}_n, q_n) = \frac{\exp(\alpha_n^u - \beta_n^u \cdot q_n)}{1 + \sum_{m=1}^{M} \exp(\alpha_{n,m} - \beta_{n,m} \cdot p_{n,m}) + \exp(\alpha_n^u - \beta_n^u \cdot q_n)} \Gamma_n.$$

Here, the β parameters specify price sensitivity and the Γ parameters represent market size. Based on our discussions with practitioners, we set $(\beta_{n,1}, ..., \beta_{n,5})$ = (0.00095, 0.000975, 0.001, $0.001025, 0.105$) and $\beta_n^u = 0.0011$ for $1 \leqslant n \leqslant 10$. The prices of a reserved ticket for the ten itineraries (denoted by $\hat{p}_n, n = 1, 2, \dots, 10$) in Japanese Yen are as follows[5]: 3010 (Tokyo-Shinyokohoma), 11300 (Tokyo-Nagoya), 14170 (Tokyo-Kyoto), 14720 (Tokyo-Shinosaka), 10640 (Shinyokohoma-Nagoya), 13500 (Shinyokohoma-Kyoto), 14390 (Shinyokohoma-Shinosaka), 5910 (Nagoya-Kyoto), 6680 (Nagoya-Shinosaka), 3070 (Kyoto-Shinosaka). Using these prices, we set $\alpha_{n,m} = 0.001 \cdot \hat{p}_n$ for $1 \leqslant n \leqslant 10$ and $1 \leqslant m \leqslant 5$, and $\alpha_n^u = 1.1 \cdot \alpha_{n,m}$ for $1 \leqslant n \leqslant 10$. The average total number of passengers traveling over the five trains on each of the ten itineraries are as follows[6]: 87, 677, 390, 846, 125, 110, 175, 77, 232, 61; we denote these by $\hat{d}_n$, $n = 1, 2, \dots, 10$. The maximum number of coaches (resp., minimum number of unreserved coaches) is $\overline{b}_m = 16$ (resp., $\underline{b}_m = 4$), for $1 \leqslant m \leqslant 5$. The probabilities $\delta_{n,m}$ with which the demand for unreserved passengers realizes on specific trains (see Section 2.2) are set as follows: $\delta_{n,1} = 0.3$, $\delta_{n,2} = 0.25$, $\delta_{n,3} = 0.2$, $\delta_{n,4} = 0.15$, $\delta_{n,5} = 0.1$, for $1 \leqslant n \leqslant 10$. The cost for operating a coach (c_o) on the line is approximately $100,000$ Yen.

We generate three full-scale instances, one each corresponding to low, medium, and high demand, by choosing three values of a demand-scaling factor η (namely, 0.5, 2.0, 5.0) and setting $\Gamma_n = \eta \cdot \hat{d}_n$ for $1 \leqslant n \leqslant 10$. For these instances, we set the per-person per-leg penalty cost for standing (c_s) to 1000. For our larger test bed described below, we vary this penalty cost systematically to examine its impact on congestion in the unreserved coaches, and thereby the firm's net profit.

5.1.2. Full-Scale Instances Based on Yan et al. (2022) In Yan et al. (2022), the authors use a logit model to estimate the passenger-choice behavior of buying high-speed railway tickets from data based on the Beijing-Hohhot high-speed railway in China. The authors consider two products, namely full-price ticket and discounted ticket. They estimate the price-sensitivity parameter for the utility of both the products to be equal to -0.0098. Based on this estimate, we generate another set of three full-scale instances for our setting as follows.

We again consider the Tokyo-Shinosaka railway line with the number of legs $L = 4$, the number of itineraries $N = 10$, and the number of trains $M = 5$. Let $\hat{\beta} = -0.0098$, as obtained from Yan et al. (2022). The price-sensitivity parameters are drawn as follows: For $1 \leqslant n \leqslant 10$, we draw $\beta_{n,m}$ randomly from U$[0.95 \cdot \hat{\beta}, 1.05 \cdot \hat{\beta}]$ for $1 \leqslant m \leqslant 5$. Since unreserved demand is typically more price-sensitive than reserved demand, we set $\beta_n^u = 1.1 \cdot \hat{\beta}$. The base-attraction parameters are set as follows: $\alpha_{n,m} = \hat{\beta} \cdot \hat{p}_n$ for $1 \leqslant n \leqslant 10$ and $1 \leqslant m \leqslant 5$, and $\alpha_n^u = 1.1 \cdot \alpha_{n,m}$ for $1 \leqslant n \leqslant 10$. For

[5] Data obtained from the Central Japan Railway Company; `https://global.jr-central.co.jp/en/info/fare/`.

[6] Data obtained from West Japan Marketing Communications Inc. records (`https://www.jcomm.co.jp/transit/price/databook/pdf/19D_41.pdf`) and Ministry of Land, Infrastructure and Transport (`http://www.mlit.go.jp/common/001193645.xlsx`), and the procedure described in Konno and Raghunathan (2020).

all the other parameters in our model, we choose the same values as those in Section 5.1.1. As discussed above, we generate three full-scale instances corresponding to three different demand levels – low, medium, and high.

5.1.3. A Test Bed of Smaller Instances To demonstrate robustness of the performance of our four policies, we also generate a test bed of smaller instances. For this test bed, we set $L = 2$, $N = 3$, $M = 3$, $c_o = 65,000$, $c_s = 1000$, $\bar{b}_m = 5$, $\underline{b}_m = 1$ for $1 \leqslant m \leqslant 3$, and $\delta_{n,1} = 0.25$, $\delta_{n,2} = 0.35$, $\delta_{n,3} = 0.4$ for $1 \leqslant n \leqslant 3$. The prices of a reserved ticket for the three itineraries (denoted by $\hat{p}_n, n = 1, 2, 3$) are 3010, 11300, and 10640, respectively. The average number of passengers traveling over the three trains on each of the three itineraries (denoted by $\hat{d}_n$, $n = 1, 2, 3$) are 52, 406, and 75, respectively. As before, we generate problem instances with multiple levels of demand; here, we choose six values of a demand-scaling factor η (namely, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0) and set $\Gamma_n = \eta \cdot \hat{d}_n$ for $1 \leqslant n \leqslant 3$. For each value of η, we generate 30 instances by randomly choosing the remaining parameters (that appear in the demand functions for the reserved and unreserved tickets) as follows: $\beta_{n,m}$ is drawn randomly from U[0.00095,0.00105] for $1 \leqslant n \leqslant 3$ and $1 \leqslant m \leqslant 3$, β^u_n is drawn randomly from U[0.00105,0.00115] for $1 \leqslant n \leqslant 6$, $\alpha_{n,m}$ is drawn randomly from U[$0.0009 \cdot \hat{p}_n$,$0.0011 \cdot \hat{p}_n$] for $1 \leqslant n \leqslant 3$ and $1 \leqslant m \leqslant 3$, and $\alpha^u_{n,m}$ is drawn randomly from U[$0.001 \cdot \hat{p}_n$,$0.0012 \cdot \hat{p}_n$] for $1 \leqslant n \leqslant 3$.

5.2. Results and Insights

For each problem instance, we simulate 100 sample paths (where each path corresponds to a specific realization of the arrival of customers over the sales horizon) to compute the expected profit ($\mathcal{J}^\pi_S$) obtained by a policy (say π) in problem $\mathbb{P}_S$. Using the optimal profit ($\mathcal{J}_D$) of the upper-bound problem $\mathbb{P}_D$, we compute the optimality gap of policy π relative to the upper bound as follows:

$$\%\text{Gap}^\pi = \frac{\mathcal{J}_D - \mathcal{J}^\pi_S}{\mathcal{J}_D} \times 100.$$

Table 2 shows the performance of our policies on the three full-scale instances described in Section 5.1.1, and Table 3 shows the performance on the three full-scale instances based on Yan et al. (2022) and described in Section 5.1.2: All the optimality gaps are less than 5%, demonstrating excellent performance.

Table 4 shows the performance of our policies on the smaller instances described in Section 5.1.3. For each of the six values of the demand-scaling factor η and for each of the four

Demand Level	**Static Price Fixed Capacity** (π_1)	**Static Price Flexible Capacity** (π_2)	**Dynamic Price Fixed Capacity** (π_3)	**Dynamic Price Flexible Capacity** (π_4)
Low	3.400 %	1.711 %	5.055 %	4.026 %
Medium	2.352 %	0.621 %	2.206 %	1.472 %
High	2.381 %	2.322 %	1.619 %	1.617 %

Table 2 The performance of our policies on the three full-scale instances

policies, the table shows the average of the percentage optimality gaps over the corresponding 30 instances. For each of the four policies, the average optimality gap over the $6 \times 30 = 180$ instances is less than 5%, again establishing the effectiveness of our policies. In Appendix E, we verify the robustness of the performance of our policies by testing with multiple combinations of values for the cost parameters c_o and c_s. In Appendix F, we demonstrate that the superior theoretical (i.e., asymptotic) performance of our dynamic-pricing policies relative to our static-pricing policies translates into a superior practical (i.e., non-asymptotic) performance.

Remark 3. **(Re-Solving Based Dynamic Pricing Policies)** In our dynamic-pricing policies π_3 and π_4 we examined in Section 4, the price of a ticket on any itinerary can change (via a linear rate-correction algorithm) in every period of the sales horizon. In Appendix D, we also develop and evaluate dynamic-pricing policies for problem $\mathbb{P}_S$ that are based on re-solving its deterministic variant $\mathbb{P}_D$ once at the midpoint of the sales horizon. ■

Value of Dynamic Pricing and Value of Flexible Capacity

We now examine the "value of dynamic pricing" under fixed capacity and under flexible capacity. Then, we also examine the "value of flexible capacity" under static pricing and under dynamic pricing. Finally, among our policies, we discuss the firm's preference between moving from static pricing to dynamic pricing and/or fixed capacity to flexible capacity.

The value of dynamic pricing under fixed capacity is defined as the percentage improvement in profit when the firm moves from static pricing to dynamic pricing, under fixed capacity. The value of dynamic pricing under flexible capacity is analogously defined. For our test bed of instances described in Section 5.1.3, Figure 2 shows the value of dynamic pricing as demand varies, both under fixed capacity and under flexible capacity. Specifically, the left-hand-side plot shows the percentage improvement in profit (average over the 30 instances in the test bed) when the firm moves from static pricing to dynamic pricing, as a function of the demand-scaling factor. There are several interesting observations:

1. When demand is low, our static-pricing policies perform better than our dynamic-pricing policies. We now discuss the intuition. Recall that the prices used in the static-pricing policies are obtained from the optimal solution to the deterministic upper-bound problem $\mathbb{P}_D$. For the static-pricing policies, under uncertain demand, the sales on some sample paths are lower (resp., higher) than those in the optimal solution to $\mathbb{P}_D$, eventually resulting in an expected revenue close to the optimal revenue of $\mathbb{P}_D$. Here, higher sales than those in the solution to $\mathbb{P}_D$ occur on some

Demand Level	**Static Price Fixed Capacity** (π_1)	**Static Price Flexible Capacity** (π_2)	**Dynamic Price Fixed Capacity** (π_3)	**Dynamic Price Flexible Capacity** (π_4)
Low	2.303 %	2.303 %	4.995 %	4.995 %
Medium	2.142 %	0.606 %	1.223 %	0.588 %
High	1.948 %	1.656 %	1.524 %	1.358 %

Table 3 The performance of our policies on the three full-scale instances based on Yan et al. (2022)

Demand-Scaling Factor (η)	**Static Price Fixed Capacity** (π_1)	**Static Price Flexible Capacity** (π_2)	**Dynamic Price Fixed Capacity** (π_3)	**Dynamic Price Flexible Capacity** (π_4)
0.5	5.840 %	4.950 %	7.813 %	7.349 %
1	3.498 %	2.194 %	4.323 %	3.925 %
2	3.493 %	3.017 %	2.962 %	2.838 %
3	4.475 %	4.475 %	2.971 %	2.972 %
4	4.083 %	4.081 %	2.580 %	2.580 %
5	3.564 %	3.562 %	2.403 %	2.403 %
6	3.219 %	3.188 %	2.100 %	2.100 %

Table 4 The average optimality gaps of our policies on the smaller instances

sample paths because there is a sufficient amount of slack capacity in the solution of $\mathbb{P}_D$, due to low demand. By slack capacity, we refer to the extra capacity resulting from the *discreteness of capacity* in our model, since railway coaches can only be added in integer numbers (i.e., capacity can only be added in discrete chunks). We elaborate using an example.

Suppose that the number of seats in a coach is 100 and a maximum of 5 coaches are allowed on a train. Under low demand, say the optimal solution to the deterministic problem uses 320 seats, i.e., 3.2 coaches. Then, due to the discreteness restriction on capacity, we choose $\lceil 3.2 \rceil = 4$ coaches as the capacity in our policies. Thus, although the capacity decision here is endogenous, discreteness leads to a slack of 80 seats in the system. This slack enables the static-pricing policies to achieve higher sales on some sample paths than those in the deterministic solution and hence, an overall revenue that is close to the optimal revenue in $\mathbb{P}_D$. In contrast, our dynamic-pricing policies are heuristics based on linear rate-correction, which start the sales horizon using prices obtained from the solution to $\mathbb{P}_D$ and subsequently adjust these prices through time such that, on each sample path, sales are close to those in the solution to $\mathbb{P}_D$. Therefore, our dynamic-pricing policies do not exploit the slack in the system and hence, may be outperformed by our static-pricing policies under low demand. The slack capacity is, however, absent under high demand. This is because, under high demand, the optimal solution to the deterministic problem uses all available capacity (e.g., 500 seats in the example above). Therefore, our dynamic-pricing policies can outperform our static-pricing policies.

2. The value of dynamic pricing is more pronounced under fixed capacity when compared to that under flexible capacity. The intuition is as follows. If one views static pricing and fixed capacity as a "base" setting, then the firm has two levers to improve profit, namely using dynamic pricing and/or using flexible capacity. If the firm chooses not to exploit one of the levers, then the other lever becomes more potent. Thus, as shown in Figure 2, if the firm chooses to continue with fixed capacity, then the value of dynamic pricing is higher relative to the case where the firm uses flexible capacity.

We now discuss the value of flexible capacity. We define the value of flexible capacity under static pricing as the percentage improvement in profit when the firm moves from fixed capacity to flexible capacity, under static pricing. The value of flexible capacity under dynamic pricing

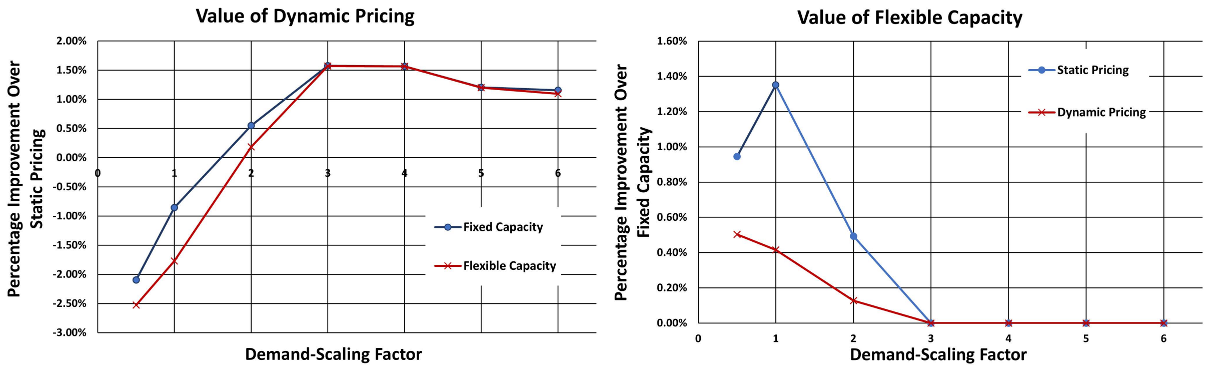

Figure 2 **The value of dynamic pricing and flexible capacity as a function of the demand-scaling factor.**

is analogously defined. For the test bed of instances described in Section 5.1.3, Figure 2 shows the value of flexible capacity as demand varies, both under static pricing and under dynamic pricing. The right-hand-side plot shows the percentage improvement in profit (average over the 30 instances in the test bed) when the firm moves from fixed capacity to flexible capacity, as a function of the demand-scaling factor. We note the following observations:

1. The value of flexible capacity is high when demand is low to moderate. The fixed-capacity policies decide capacity at the beginning of the sales horizon, while the flexible-capacity policies have the luxury of adjusting capacity, either upward or downward, based on the realized demand. The benefit of this flexibility is significant when demand is low. When demand is high, both fixed- and flexible-capacity policies typically use the maximum feasible capacity leaving little or no potency in capacity adjustments and hence, the value of flexible capacity is minimal.

2. The value of flexible capacity is higher under static pricing as compared to that under dynamic pricing. As discussed above, when the firm does not exploit the dynamic-pricing lever, the value of using the flexible-capacity lever is higher.

Among our policies, we now examine the firm's preference between moving from static to dynamic pricing or moving from fixed to flexible capacity (or both), and the effect of demand on

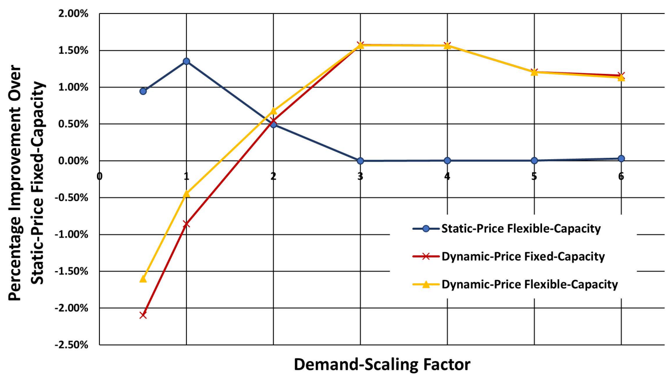

Figure 3 **The relative benefit of dynamic pricing and flexible capacity, as a function of the demand-scaling factor.**

these decisions. If one views the static-price fixed-capacity policy as our "base" policy, then the firm has three options:

(a) Change only the pricing decisions. That is, move from the static-price fixed-capacity policy to the dynamic-price fixed-capacity policy.

(b) Change only the capacity decisions. That is, move from the static-price fixed-capacity policy to the static-price flexible-capacity policy.

(c) Change both the pricing and capacity decisions. That is, move from the static-price fixed-capacity policy to the dynamic-price flexible-capacity policy.

Figure 3 shows the relative performance of these three options for the test bed of instances described in Section 5.1.3. The figure plots the percentage increase in profit (average over the 30 instances in the test bed) as the firm shifts from our static-price fixed-capacity policy to each of our other three policies, as a function of the demand-scaling factor. Under low demand, as discussed in detail above, static pricing is relatively attractive, and the benefit under flexible capacity is high. Thus, our static-price flexible-capacity policy is the most attractive. When demand is sufficiently high, dynamic pricing outperforms static pricing, whereas the marginal benefit offered by flexible capacity reduces since there is little potency in capacity adjustments. Therefore, the performance of the dynamic-price fixed-capacity policy is as good as that of the dynamic-price flexible-capacity policy. Consequently, it is sufficient to switch only the pricing decisions from static to dynamic.

We discuss additional insights in the supplementary appendix. In Appendix I, we examine limited flexible-capacity policies, where capacity decisions are delayed to an intermediate point in the sales horizon. In Appendix J, we study the impact of prohibiting unreserved travel.

5.3. Congestion Management

We examine the impact of the per-person per-leg penalty cost for standing (c_s) on the congestion in the unreserved coaches and also on the profit of the firm. For an instance in our test bed described in Section 5.1.3, Figure 2 illustrates the behavior of the expected profit, revenue, and congestion as a function of c_s, under our static-price fixed-capacity policy (π_1). The left-hand-side plot in Figure 4 shows the behavior of the expected net profit (objective of $\mathbb{P}_S$, i.e., expected revenue less operational costs and total penalty cost from standing), the expected revenue in $\mathbb{P}_S$, and the net upper-bound profit (objective of $\mathbb{P}_D$), as a function of c_s.

An interesting observation is that the expected revenue can increase with an increase in c_s. This can happen due to several reasons, two of which we mention here. Consider two values of the penalty cost, namely c_{s_1} and c_{s_2} with $c_{s_1} < c_{s_2}$: (a) An increase from c_{s_1} to c_{s_2} can lead to an increase in the number of coaches used (to reduce congestion) and also an increase in the sales of unreserved tickets. This increases both the expected revenue and operational cost. (b) An increase from c_{s_1} to c_{s_2} can lead to an increase in the prices used by our policy (obtained from the

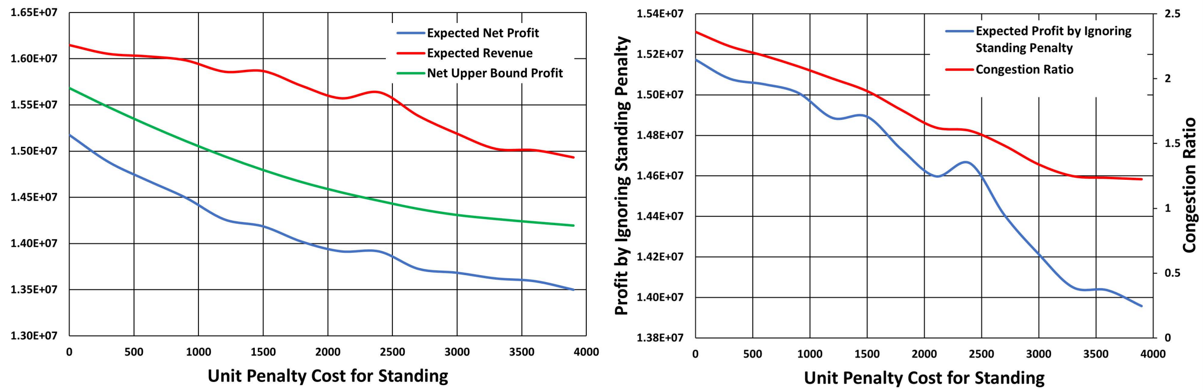

Figure 4 **An examination of congestion and profit as a function of penalty for standing (c_s).**

solution to $\mathbb{P}_D$) keeping the number of coaches unchanged, leading to a decrease in the unreserved demand but possibly an increase in the expected revenue. The non-monotone behavior in the expected revenue also translates to the expected profit computed by ignoring the penalty cost for standing passengers, as shown in the right-hand-side plot of Figure 4.

Let the "congestion ratio" be defined as the expected number of people traveling in unreserved coaches divided by the total number of unreserved seats. Figure 2 also illustrates the behavior of the congestion ratio and the expected profit as a function of c_s. Such plots can serve as useful tools for railway administrators as they attempt to arrive at an appropriate value of c_s to strike a balance between profit and congestion.

6. Concluding Remarks

Revenue management in railways distinguishes itself from that in traditional applications such as airlines, hotels, and fashion retail, in several ways, including the viability of jointly optimizing pricing and capacity, capacity acquisition in discrete chunks, the existence of unreserved capacity, standing travel, and the need for congestion management. Based on our work with a railway company in Japan, we analyze a joint pricing and capacity problem, and develop four asymptotically optimal policies to cater to different practical needs (static or dynamic price, and fixed or flexible capacity). Table 5 offers a succinct summary of our results. Our numerical study demonstrates the excellent performance of our policies on instances based on real-world rail operations.

To demonstrate the broader applicability of our analysis, we briefly discuss two potential applications beyond railways.

Joint Pricing and Capacity Decisions in Delivery Services: Consider a platform that offers product delivery services, say the delivery of groceries, and consider a single day of operations. The platform receives customer requests in advance for deliveries to be made on the focal day. The platform accepts these requests until midnight of the previous day – we refer to the duration until this deadline as the request horizon. Customers have the option to choose one from several available time windows for their delivery; for example, 9am–12noon, 12noon–3pm, or 3pm–6pm. Throughout the request horizon, the platform can dynamically vary the delivery prices for the

	Static Price	Dynamic Price
Flexible Capacity	Policy π_1 **(Section 3.1)** Asymptotic Optimality: Yes Convergence Rate: $\sqrt{\kappa}$ Capacity Decision: End of sales horizon Pricing Decision: State Independent Computation: Uses solution to $\mathbb{P}_D$ and $\mathbb{Q}^{\pi_1}$	Policy π_3 **(Section 4.1)** Asymptotic Optimality: Yes Convergence Rate: $\log\kappa$ Capacity Decision: End of sales horizon Pricing Decision: State Dependent Computation: Uses solution to $\mathbb{P}_D$ and $\mathbb{Q}^{\pi_3}$
Fixed Capacity	Policy π_2 **(Section 3.2)** Asymptotic Optimality: Yes Convergence Rate: $\sqrt{\kappa}$ Capacity Decision: Start of sales horizon Pricing Decision: State Independent Computation: Uses solution to $\mathbb{P}_D$	Policy π_4 **(Section 4.2)** Asymptotic Optimality: Yes Convergence Rate: $\log\kappa$ Capacity Decision: Start of sales horizon Pricing Decision: State Dependent Computation: Uses solution to $\mathbb{P}_D$

Table 5 Overview of our four policies for problem $\mathbb{P}_S$ ($\kappa \in \mathbb{N}$ is the demand and capacity scaling factor)

different time windows, depending on demand. The requests for deliveries arrive randomly at rates that depend on the delivery prices. Simultaneously, on the supply side, the firm reserves delivery drivers to fulfill delivery requests on the focal day. This supply is also uncertain – drivers are available at a rate that depends on the daily wages set by the platform. The platform's objective is to maximize its profit by dynamically optimizing delivery prices for the different delivery windows and the daily driver wages. Several elements in this setting are similar to those in the problem we analyzed in this paper – for example, on the demand-side, delivery windows are similar in spirit to multiple choices of trains for an itinerary, while on the supply-side, delivery windows are conceptually similar to legs in train-travel. As another example, the number of deliveries that can be made by a driver over a time window is akin to the number of seats in a coach. Our analysis can be used as a stepping stone to develop effective policies for this setting.

Joint Pricing and Capacity Decisions in the Movie-Theater Industry: Here, each movie is similar to an itinerary for which tickets are sold. The different time-slots for shows correspond to different trains. For a given time-slot and a given movie, the number of (parallel) screens on which the movie is shown is the capacity decision for that movie and time-slot combination. Finally, pricing decisions are to be made for each movie and time-slot combination; these prices can vary throughout the sales horizon.

Acknowledgements

The authors are grateful to the department editor, the associate editor, and three anonymous reviewers for their many sharp comments and constructive suggestions that significantly improved the paper, in both content and exposition.

References

Amtrak 2022. Coach class seating on unreserved services. `https://tinyurl.com/f5yb56n3`.

Armstrong, A., and J. Meissner. 2010. Railway revenue management: Overview and models.

Atar, R., and M. I. Reiman. 2012. Asymptotically optimal dynamic pricing for network revenue management. *Stochastic Systems* 2 (2): 232–276.

Besbes, O., and C. Maglaras. 2012. Dynamic pricing with financial milestones: Feedback-form policies. *Management Science* 58 (9): 1715–1731.

Bumpensanti, P., and H. Wang. 2020. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science* 66 (7): 2993–3009.

Business Traveller 2020. Dynamic fares arrive on Beijing–Shanghai high–speed rail. `https://tinyurl.com/2zc5jd5b`.

Chen, Q., S. Jasin, and I. Duenyas. 2016. Real-time dynamic pricing with minimal and flexible price adjustment. *Management Science* 62 (8): 2437–2455.

Ciancimino, A., G. Inzerillo, S. Lucidi, and L. Palagi. 1999. A mathematical programming approach for the solution of the railway yield management problem. *Transportation Science* 33 (2): 168–181.

Ferreira, K. J., B. H. A. Lee, and D. Simchi-Levi. 2016. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18 (1): 69–88.

Ferreira, K. J., D. Simchi-Levi, and H. Wang. 2018. Online network revenue management using thompson sampling. *Operations Research* 66 (6): 1586–1602.

Gallego, G., and G. Van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* 45 (1): 24–41.

Hetrakul, P., and C. Cirillo. 2014. A latent class choice based model system for railway optimal pricing and seat allocation. *Transportation Research Part E: Logistics and Transportation Review* 61:68–83.

Hohberger, S. 2020. Dynamic pricing under customer choice behavior for revenue management in passenger railway networks. *Ph.D. thesis. Universität Mannheim*. `https://madoc.bib.uni-mannheim.de/54035`.

Jasin, S. 2014. Reoptimization and self-adjusting price control for network revenue management. *Operations Research* 62 (5): 1168–1178.

Kamandanipour, K., M. M. Nasiri, D. Konur, and S. H. Yakhchali. 2020. Stochastic data-driven optimization for multi-class dynamic pricing and capacity allocation in the passenger railroad transportation. *Expert Systems with Applications* 158:113568.

Ke, J., D. Zhang, and H. Zheng. 2019. An approximate dynamic programming approach to dynamic pricing for network revenue management. *Production and Operations Management* 28 (11): 2719–2737.

Konno, N., and A. Raghunathan. 2020. Data-driven joint optimization of pricing and seat allocation in trains. *Computers in Railways XVII: Railway Engineering Design and Operation* 199:379.

Kunnumkal, S., and K. Talluri. 2016. A note on relaxations of the choice network revenue management dynamic program. *Operations Research* 64 (1): 158–166.

Lei, Y., S. Jasin, and A. Sinha. 2018. Joint dynamic pricing and order fulfillment for e-commerce retailers. *Manufacturing & Service Operations Management* 20 (2): 269–284.

Li, H., and W. T. Huh. 2011. Pricing multiple products with the multinomial logit and nested logit models: Concavity and implications. *Manufacturing & Service Operations Management* 13 (4): 549–563.

Maglaras, C., and J. Meissner. 2006. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management* 8 (2): 136–148.

Price of Travel 2022. Should you buy European train tickets in advance: Yes, and here's how to save. `https://tinyurl.com/5apb66ap`.

Qin, J., Y. Zeng, X. Yang, Y. He, X. Wu, and W. Qu. 2019. Time-dependent pricing for high-speed railway in China based on revenue management. *Sustainability* 11 (16): 4272.

Rail Mitra 2019. Dynamic fare pricing in premium trains of the Indian Railways. `https://tinyurl.com/45j9t6cu`.

Talluri, K. T., and G. J. Van Ryzin. 2006. *The Theory and Practice of Revenue Management*. Springer.

The Urbanist 2019. Going the distance on Amtrak. `https://tinyurl.com/2p8t334y`.

Trains.com 2019. Amtrak cites safety in turning away CTrail ticket holders. `https://tinyurl.com/yckdj5xc`.

Wang, X., H. Wang, and X. Zhang. 2016. Stochastic seat allocation models for passenger rail transportation under customer choice. *Transportation Research Part E: Logistics and Transportation Review* 96:95–112.

Wang, Y., and H. Wang. 2022. Constant regret resolving heuristics for price-based revenue management. *Operations Research* 70 (6): 3538–3557.

Yan, Z., M. Jian, X. Li, and J. Cao. 2022. Modelling the passenger choice behaviour of buying high-speed railway tickets. *Promet-Traffic&Transportation* 34 (3): 455–465.

You, P.-S. 2008. An efficient computational approach for railway booking problems. *European Journal of Operational Research* 185 (2): 811–824.

Zhai, X., J. Zhao, and Q. Chen. 2018. Optimization of the assignment of tickets for railway networks with large passenger flows. *Journal of Rail and Rapid Transit* 232 (2): 632–642.

Zhang, D., and L. Weatherford. 2017. Dynamic pricing for network revenue management: A new approach and application in the hotel industry. *INFORMS Journal on Computing* 29 (1): 18–35.

Zhu, F., S. Liu, R. Wang, and Z. Wang. 2021. Assign-to-Seat: Dynamic capacity control for selling high-speed train tickets. *Available at SSRN 3477339*.