Predictive and Prescriptive Analytics for Airport Slot Allocation

by

Phillip D. Schmedeman

B.S., Engineering Management, United States Military Academy (2011)

Submitted to the System Design and Management Program
in partial fulfillment of the requirements for the degree of

Master of Science in Engineering and Management

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2021

Author .....
System Design and Management Program
May 14, 2021

Certified by .....
Alexandre Jacquillat
Assistant Professor, Operations Research and Statistics
Thesis Supervisor

Accepted by .....
Joan S. Rubin
Executive Director, System Design and Management Program

Predictive and Prescriptive Analytics for Airport Slot Allocation

by

Phillip D. Schmedeman

Submitted to the System Design and Management Program
on May 14, 2021, in partial fulfillment of the
requirements for the degree of
Master of Science in Engineering and Management

Abstract

Slot allocation is the primary form of strategic demand management practiced at airports globally to address congestion and reduce delay. To perform slot allocation, airport schedulers must account for detailed requests from hundreds of airlines for thousands of flights over a six-month season while adhering to variable airport capacities and the Worldwide Airport Slot Guidelines (WASG). This represents a highly complex combinatorial scheduling problem that has vast implications for airlines and passengers. While previous research has proposed a range of optimization models to support slot allocation, they commonly assume a flight-centric approach, which may extend or eliminate passenger connections without accounting for the costs.

This thesis develops an original approach to airport slot allocation that incorporates passenger considerations. The proposed multi-objective optimization model allocates slots according to the WASG and airport capacity constraints while minimizing one flight-centric metric—schedule displacement—and two passenger-centric metrics—infeasible connections and connection time. Since this approach requires passenger forecasts to account for costs, we use historical itinerary data and machine learning methods to predict passenger flows across a network of flights. We apply this predict-then-optimize framework using real-world data from Singapore Changi Airport to create slot assignments that achieve Pareto optimality in acceptable computation times. The results indicate that schedule-coordinated airports can reduce passenger costs from slot allocation, with relatively small adjustments to schedule displacement. Ultimately, the proposed multi-objective formulation provides a new paradigm that can create more attractive flight schedules at major airports worldwide, based on airport-level considerations, airline-level considerations, and, for the first time, passenger-level considerations.

Thesis Supervisor: Alexandre Jacquillat
Title: Assistant Professor, Operations Research and Statistics

Acknowledgments

This thesis is the result of a two-year journey at MIT, which would not have been possible without the people supporting me in Cambridge and around the world.

First, I would like to thank my thesis supervisor, Professor Alexandre Jacquillat, without whom this thesis would not have been possible. Alex has been an invaluable source of support and inspiration. Alex is vigorous in his desire to tackle the most challenging problems and brilliant in his ability to make them understandable. His passion and enthusiasm for analytics gave me a boost of energy during every encounter. Thank you for taking a chance on me and for investing your time in our work together.

Special thanks to Nuno Ribeiro and Sebastian Birolini for being so welcoming and introducing me to the world of academic research. Nuno and Sebastian are rising stars in the field of operations research, and collaborating with them has been an honor. They were generous with their time and patient as I learned. I look forward to our continued work on this project and hope to see you in person soon.

I am grateful for the opportunity to roam the corridors of MIT and for the faculty members who shared their time and knowledge. Specifically, Professors Olivier de Weck, Amedeo Odoni, and James Orlin provided wisdom and opened my mind to the ideas that shaped this work.

Many thanks to my friends in the SDM program. Learning alongside you was a joy, and our off-campus adventures provided the perfect respite. I feel incredibly fortunate to have met you and been surrounded by this diverse group of talented leaders. Thank you for your help and friendship. Good luck with your next endeavors, and I hope to see you all soon.

I would also like to thank my family. Thank you to my parents for prioritizing my education. Thank you to Kate and Henry for providing tranquility in a chaotic world and the perfect pandemic sanctuary. Thank you to my son, Asher, who was born three months before the completion of this thesis. I promise to teach you everything I know and continue learning from you.

Last but most important thanks to my beautiful, intelligent, and loving wife, Rashina. We climbed this mountain together, as we will the next one. Thank you for your patience, strength, and support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Airport Congestion

> "Airport and air traffic congestion is a growing problem on an international scale and is widely viewed as one of the principal constraints to the future growth of the global air transportation industry" [32].

In addition to providing a means of worldwide rapid transit, air travel plays a central role in the global economy by facilitating employment, trade, and tourism. Nevertheless, the world's busiest airports are experiencing severe challenges due to escalating congestion driven by an increasingly constrained relationship between capacity and demand. While the coronavirus pandemic has halted air traffic growth, it has also plunged the aviation system into a new era of uncertainty as the industry slowly scales operations back to pre-pandemic levels. In these times of disruptions and uncertainties, the efficient use of scarce airport resources becomes even more critical.

Airport capacity is primarily determined by the capacity of the airfield and particularly the runway system [32]. Runways create a bottleneck as three-dimensional air traffic flows compress into a single file for arrival and departure. The most influential factor in runway system capacity is the number and geometric layout of runways. It is prohibitively difficult to expand runways at the most congested airports. Runway

expansion requires additional land and generates significant economic and environmental effects that necessitate complex and lengthy approval processes.

The most direct consequence of congestion is travel delay. Excepting the pandemic, travel delay is particularly severe at airports in North America, Europe, and the Pacific Rim. In Europe, from 2014 to 2019, approximately 39% of flights were delayed on arrival, with an average delay of approximately 29 minutes [13, 14, 15, 16, 17, 18]. These metrics are often worse in the United States, and airports in Asia that were previously delay-free observed significant delays in the years preceding 2020. Since delays propagate throughout the air transportation network, delays at one congested airport spread to others, even those with excess capacity.

Travel delays create significant costs for airlines and travelers. In the most thorough accounting of the costs associated with delay, researchers estimated that the total cost of domestic air travel delays in the United States was $31.2 billion in 2007 [3]. That total includes $8.3 billion in additional operating costs for airlines, $16.7 billion associated with passenger delays, and $6.2 billion in other indirect costs to the economy. Although we are not aware of a similarly comprehensive study of the costs of air travel delays in other regions, recent estimates of annual costs in Europe are on the order of $10 billion [35]. An attempt to measure the global cost to airlines and passengers by Amadeus estimated the cost at $60 billion in 2016, corresponding to 8% of worldwide airline revenue.

1.2 Background on Slot Allocation

The significant effects of congestion, and the inability of many airports to make capacity increases in the near term, have prompted research and application of strategic demand management. Demand management aims to maintain the efficiency of air travel in the face of congestion through administrative or economic interventions. While there are many approaches to strategic demand management, they commonly aim to mitigate congestion by limiting airline access to the airport.

The motivation for demand management emerges from the observation that the

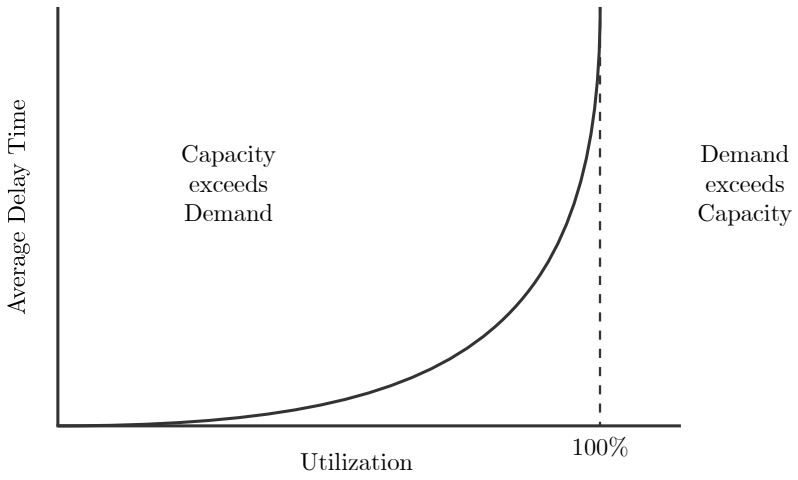

Figure 1-1: Nonlinear relationship between delay and utilization [32]

relationship between the utilization ratio ρ and delay is nonlinear [32]. Under steady-state conditions, the average expected delay is proportional to $\frac{1}{1-\rho}$, where ρ is the ratio of demand over capacity. The utilization ratio is known in queuing theory as a fundamental measure of level of service. To illustrate this relationship we plot delay as a function of ρ in Figure 1-1. Notice that delay is highly sensitive to changes in demand or capacity when an airport is operating near saturation. As demand increases (e.g., a wave of departures) or capacity decreases (e.g., adverse weather conditions), delay increases disproportionately. Supply-side efforts to increase capacity should be explored as a first option, but in situations where capacity improvements are not possible in the near term, demand management provides a reasonable alternative. Demand management offers strategies to keep air travel accessible and efficient for a wide range of customers, and it is relatively quick and inexpensive to implement.

Airports in the United States are unique in that they apply a laissez-faire approach to demand management. Aircraft operators are free to schedule arrivals and departures at any time provided that there is capacity to access the terminal buildings. By leaving airport scheduling largely unconstrained, this approach relies upon the assumption that delay costs will be internalized by the airlines, thereby preventing delays from exceeding levels viewed as tolerable by the airlines. As a result, airports in the United States achieve higher levels of throughput than comparable airports in

Europe and Asia, but worse on-time performance [31].

A counter-argument to the laissez-faire approach to scheduling is that the airlines do not fully account for costs to other stakeholders—namely passengers—when left to self-regulation. At highly congested airports outside of the United States, there is widespread agreement that the benefits of well-executed demand management outweigh the costs [32]. In particular, demand management creates a more evenly distributed demand profile, thereby mitigating travel delays.

Active approaches to strategic demand management fall into two main categories: economic and administrative. Research indicates that economic approaches, such as congestion pricing and slot auctions, can achieve economically efficient outcomes in theory by allocating scarce airport resources to the flights that generate the most value. However, stakeholder resistance to the monetary transfers and effects on competition has prevented the widespread adoption of economic approaches.

The prevailing administrative approach to strategic airport demand management is schedule coordination. The International Air Transportation Authority (IATA), a trade association, governs schedule coordination through slot guidelines. According to IATA, 204 airports representing 43% of global traffic conducted schedule coordination under the slot guidelines in 2019 [21]. A "slot" grants permission to arrive at or depart from an airport at a specified date and time. If the requests for operations exceed an airport's declared capacity at a given time, some itineraries are displaced to a time with surplus capacity. The 2020 edition of IATA's Worldwide Airport Slot Guidelines (WASG) states that the "prime objective of airport slot coordination is to ensure the most efficient declaration, allocation and use of available airport capacity in order to optimize benefits to consumers, taking into account the interests of airports and airlines" [1].

Each schedule-coordinated airport conducts slot coordination independently according to the general timeline represented in Figure 1-2. The process begins around one year before a season, with each slot-coordinated airport reporting their declared capacities, which specify the number of slots to allocate by the time of day. Independently, airlines conduct fleet and network planning to determine their desired flight

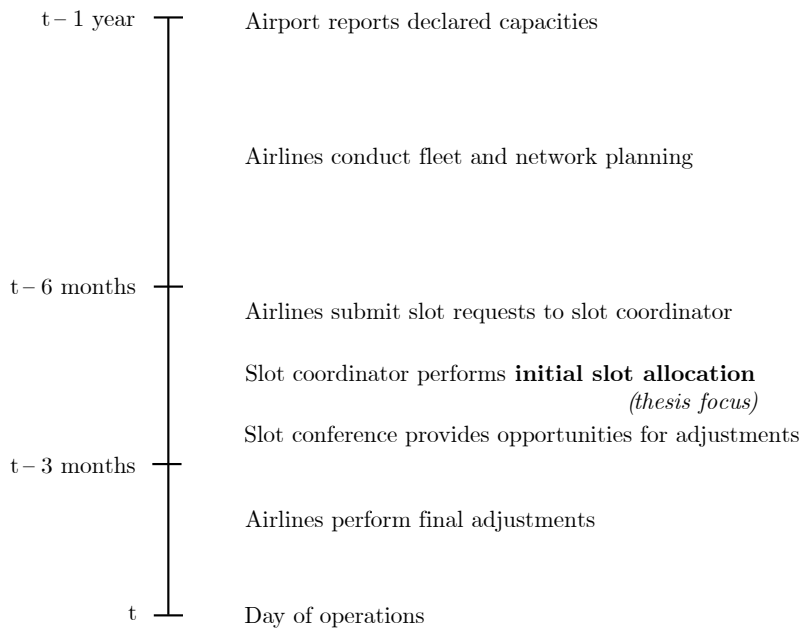

Figure 1-2: Slot allocation timeline

schedules. Approximately five months before a season, the airlines submit requests to slot coordinators at each airport for arrival and departure times. Approximately four months before a season, the slot coordinators perform initial slot allocation, which is the focus of this thesis. Slot coordinators aim to meet the airlines' requests while respecting the airport's capacity constraints. Requests that exceed the airport's declared capacity are displaced (i.e., shifted to times with excess capacity) according to priorities designated in the WASG, but the method of displacement is unregulated. Following initial slot allocation, the airlines have opportunities to make adjustments, including at bi-annual slot conferences, but initial allocation remains the most important stage of the process.

To perform slot allocation, airport slot coordinators must account for variable airport capacities and detailed itinerary requests from hundreds of airlines for thousands of flights over a six-month season (summer or winter) while adhering to the WASG. This represents a highly complex combinatorial problem that has vast implications for passengers, airlines, and airports. In practice, coordinators apply a variety of approaches, often assisted by special-purpose software that processes requests se-

quentially according to their priority class [37]. Since this approach does not consider the full range of possibilities, it results in sub-optimal slot assignments that sacrifice the efficiency and utilization of scarce airport resources.

In response, researchers have developed a range of optimization models to support slot allocation. While these previous models have explored various approaches to account for the complexity of the problem, they commonly assume a flight-centric approach, which minimizes schedule displacement metrics. However, minimizing displacement does not necessarily translate into the most attractive itineraries for passengers. By focusing strictly on displacement, previous optimization models may eliminate important passenger connections or make connections longer. This increases costs for passengers and may also decrease revenue for the airlines.

1.3 Thesis Outline

The objective of this thesis is to improve decision-making in airport demand management through a novel approach to airport slot allocation. We develop an original optimization model that assigns slots according to airport capacities and the WASG while balancing the multi-objectives of minimizing schedule displacement and the costs incurred by passengers. However, a major challenge is that this model requires estimated passenger flows, which are unknown at the time of initial slot allocation (i.e., months before the day of operations). Therefore, we use historical data, and machine learning models to *predict* passenger flows, which are then used as inputs to *optimize* slot allocation. This predict-then-optimize approach enables the incorporation of passenger-level metrics, which guide decision-making with a more accurate representation of the costs associated with slot allocation.

Chapter 2 provides a review of the literature on airport slot allocation optimization, passenger cost modeling, demand forecasting, and the integration of predictive and prescriptive analytics. We briefly review optimization models from previous research that form the basis for our proposed model and discuss research that has utilized passenger-level metrics in related topics to improve decision-making. We

reference studies that demonstrate the potential for machine learning in passenger forecasting and the integration of predictive and prescriptive analytics.

Chapter 3 describes the modeling framework which integrates predictive and prescriptive analytics. We address the challenges involved with incorporating passenger considerations into slot allocation. We introduce our methodology and the experimental conditions used to validate the approach.

Chapter 4 presents predictive analytics for passenger flow forecasting, an input for optimization. The chapter begins with a description of the historical passenger itinerary and socioeconomic data used for prediction. We engineer numeric, categorical, and binary features that are strong predictors of passenger flows. We estimate parametric and non-parametric prediction models for Singapore Changi Airport and evaluate their performance over a year of operations. We report how the models perform overall, and specifically on new itineraries that did not exist in the training observations. Random forests achieves the best performance and is therefore selected to provide input for the optimization model in the following chapter.

Chapter 5 proposes an optimization formulation and solution approach for the slot allocation problem. The chapter begins with our methodology and a simple example to demonstrate the concept of including passenger considerations in slot optimization. We then explain the process for integrating the passenger prediction and slot allocation optimization models. We provide a mathematical formulation and solution algorithm for the proposed multi-objective mixed-integer program.

Chapter 6 presents the results from applying the proposed framework using real-world data to generate slot allocation solutions for Singapore Changi Airport. After accounting for passenger costs, we find that a model focused exclusively on schedule displacement would have eliminated 2,065 passenger connections and extended passenger connection time by 30,444 hours over the season. Using the proposed model, we identify new solutions along the Pareto frontier, which facilitates a discussion of trade-offs between objectives and indicates the potential for large reductions in passenger costs for relatively small increases in displacement. One such solution restores 52% of the infeasible passenger connections and reduces 16% of the connection

time incurred from slot allocation, with only a 1% increase in schedule displacement. Given the airport characteristics, we recommended solutions and provide a synthesis of the benefits. This chapter closes with a sensitivity analysis of important modeling parameters, demonstrating that the benefits are robust to uncertainty.

Chapter 7 synthesizes the results, summarizes major contributions, and proposes areas for future work. Computational challenges limited the scope of our optimization model, so we suggest methods for strengthening and expanding the prescriptive analytics. Additionally, the COVID-19 pandemic is catalyzing unprecedented change in the aviation system, which will test the flexibility and robustness of the prediction and optimization models.

Chapter 2

Literature Review

This thesis relates to the literature on airport slot allocation optimization, passenger cost modeling, demand forecasting, and the integration of predictive and prescriptive analytics. In this chapter, we provide a brief review of these fields.

Optimization models developed for the slot allocation problem have historically been formulated as integer programs. One of the earliest models was developed by Zografos et al., in 2012 [45]. This model allocated slots for a regional airport according to an objective function that minimizes schedule displacement. When applied, it achieved optimality in under five minutes and demonstrated improvements of 14-95% over the existing slot schedule. To reduce complexity, this model dealt with each priority class of slots separately, which prevents the consideration of all possible combinations. Additionally, implementation of this and similar models using exact methods remains intractable for the largest schedule-coordinated airports.

In 2018, Ribeiro et al. [37] proposed the Priority-based Slot Allocation Model (PSAM), which considers all slot classes simultaneously, enabling the full range of scheduling combinations. The PSAM uses a multi-objective function to capture trade-offs between various measures of schedule displacement (e.g., total displacement, maximum displacement, number of slots displaced or rejected). The PSAM includes constraints to account for the WASG (i.e., IATA guidelines) and strengthened the formulation with valid inequalities that enable implementation at medium-sized airports. In 2019, Ribeiro et al. [36] developed a large-scale neighborhood search

algorithm that can solve PSAM for large airports. This model achieved optimality in 10 hours for a season at a large international airport. The model proposed by this thesis is an extension of the PSAM formulation.

These previous models demonstrated that state-of-the-art optimization methods can achieve solutions for the largest schedule-coordinated airports over an entire season and that the improvement over existing methods generally grows with the airport's size. This trend matches intuition since larger airports are more likely to exceed the limitations of the existing methods utilized by slot coordinators. More recent models have introduced alternative objectives such as inter-airline fairness and equity [23, 24, 26, 43, 44], airline acceptability of schedule displacement [42, 44], airline flexibility preferences [25], and blocking (i.e., scheduling flexibility) [19]. Despite capturing a wide range of objectives, previous models take a flight-centric approach. By neglecting passenger costs associated with increasing or eliminating connections, both passengers and airlines lose value.

Literature from other fields of aviation research demonstrates how incorporating passenger considerations alters the decision-making regarding airport resources. In 2014, Barnhart et al. [4] estimated that the average delay per passenger is roughly twice the average delay per flight. Flight delay is typically calculated as the difference between a single flight's scheduled and observed arrival time. However, the delay time for passengers often relies on connections between multiple flight-legs, creating multiple opportunities for delay. Flight cancellations and missed connections cause significant disruptions (and costs) for passengers but are not considered by flight delay metrics. Additionally, as load factors increase at large transfer airports, which are often congested, it becomes more challenging to re-book passengers who miss connections, further compounding passenger delay times.

In the field of airline scheduling and disruption recovery, it is well established that there is no one-to-one mapping between flight-level and passenger-level considerations. Early models for airline disruption recovery focused exclusively on minimizing airline operating costs. These models neglected costs associated with disrupted passengers, who must be reaccommodated on alternate itineraries because of missed connections

or flight leg cancellations. In 2006, Bratu and Barnhart [8] extended the conventional airline disruption recovery model to minimize jointly airline operating costs and passenger delay and disruption costs. The results from applying this model to a case study of 83,869 passengers over a single day of operations indicated that airlines could reduce passenger delay by 19.4% without significant increases in operating costs. In 2017, Marla et al. [29] formulated a model to incorporate passenger delay costs into flight planning. This research showed that adjustments to flight speed through increased fuel burn could reduce passenger disruptions by 66-83% and achieve savings of 5.7%–5.9% for the airline.

Researchers have also incorporated passenger considerations into decision-making for tarmac delay policies [41], airline timetable development and fleet assignment [40], and air traffic flow management (ATFM) [22]. Collectively, these studies indicate that passengers should be incorporated into decision models to account for costs. By including passenger costs in decision-making, airlines and airports can improve profitability through increased customer retention. To our knowledge, this study is the first to incorporate passenger considerations directly into the optimization of slot allocation.

In the ATFM study [22], Jacquillat demonstrates the predict-then-optimize framework that is also applied in this thesis. This approach, which leverages analytics to predict an unknown variable that is used as an input for optimization, has recently experienced widespread application in academic literature under the umbrella of data-driven decision-making [5]. Since future passenger flows are required to evaluate the passenger costs associated with slot allocation decisions, we employ predictive analytics for forecasting.

Transportation demand has traditionally been estimated using logit models, following one of the earliest works by McFadden in 1973 [30]. More recent studies employ discrete choice models in the context of air transportation to capture supply-demand interactions [2, 6, 7, 12, 40]. The application of choice models in passenger prediction depends upon (i) information regarding competing itineraries and (ii) features that explain market share, such as price and time-of-day information. However, since ini-

tial slot allocation occurs months before the beginning of a season, these attributes are unavailable.

As an alternative, we employ machine learning methods for passenger prediction, leveraging several socioeconomic variables and itinerary features. Recent studies demonstrate the advantages of machine learning for passenger forecasting, as in Lhéritier et al. 2019 [27]. Logit models lack the flexibility to handle collinear attributes and correlations between alternatives, but machine learning methods are less restrictive and able to leverage many variables with complex relationships. This suggests that the integration of predictive and prescriptive analytics provides opportunities to improve slot allocation. By predicting passenger flows and incorporating passenger-level metrics into a multi-objective optimization model we can more accurately account for the costs of slot allocation decisions.

Chapter 3

Modeling Framework

As with previous slot optimization models, the proposed model takes as inputs the airline slot requests and the airport's declared capacities. It then produces a slot assignment accordingly, with the primary objective of minimizing the displacement relative to the requests while complying with the airport's capacity and the Worldwide Airline Slot Guidelines (WASG). However, our approach in this thesis differs from the previous literature by accounting for the impact of slot allocation on passenger itineraries—therefore balancing flight-level and passenger-level considerations. One complication is that passenger itineraries are unknown at the time of slot allocation, which occurs months before the flights' operations, prior to passenger bookings. Therefore, accounting for costs requires predictions of passenger flows in flight networks. Ultimately, the proposed slot allocation approach integrates two analytics models: a predictive model of passenger flows in flight networks and a prescriptive optimization model of slot allocation. Figure 3-1 provides an overview of the modeling framework.

3.1 Predictive Modeling

Forecasting passenger flows in commercial air travel is a challenging time series problem. Each itinerary has a unique level, trend, seasonality, and noise [38]. Level describes the central tendency, which can be used to differentiate high-volume itineraries

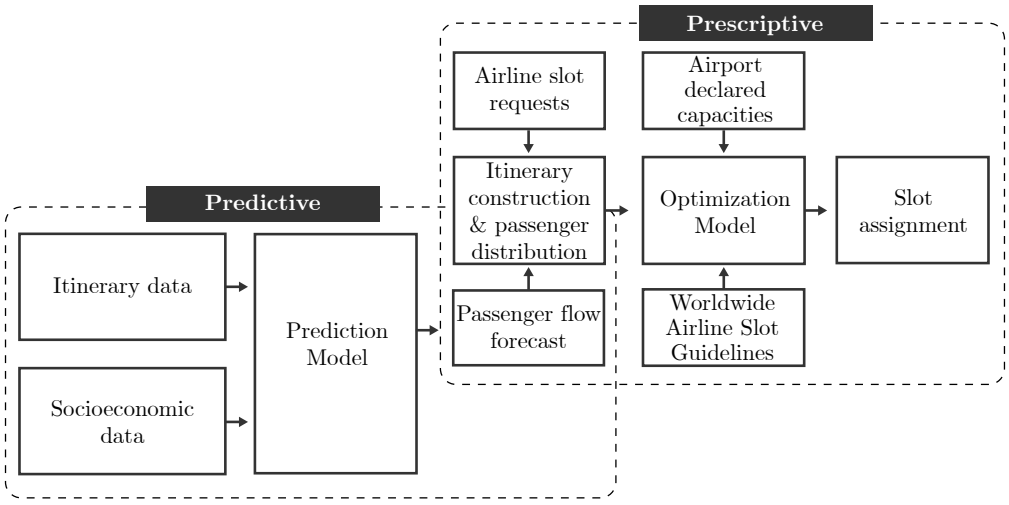

Figure 3-1: Modeling framework

from those that typically carry medium or low levels of passengers. Trend is the change from one year to the next, as itineraries experience different degrees of growth or stagnation. Seasonality is short-term cyclic behavior, which can be subtle for itineraries serving a diverse customer base or more pronounced for business or leisure-dominated markets. Noise is the random variation or measurement error that obscures these characteristics, motivating the need for strong features that adequately describe passenger flows over time.

There are a vast number of fluctuating itinerary combinations, which generate categorical variables with many categories (i.e., hundreds of airports and airlines). Some series of itineraries are consistent, while others are intermittent, which causes problems for some prediction models. Previous studies demonstrate favorable results forecasting demand for existing itineraries, but a greater challenge is to predict demand for new itineraries (i.e., new combinations of airports and airlines). In this case, we are not simply extrapolating but predicting the passenger flows for an emergent series. For this subset of itineraries, the most important predictor—historical passenger counts—is unavailable. In the dataset used for this thesis, approximately 10% of the medium to high-volume itineraries each year are new. Low-volume itineraries (less than 30 passengers per month) are even more sporadic. This prompts the de-

velopment of continuous features that are not dependent upon historical traffic and remain available for new itineraries.

The most reliable sources of historical itinerary data are aggregated by month, which means we are unable to leverage patterns over shorter time scales. The unavailability of information required to estimate market share (e.g., price and time-of-day features) and incomplete set competing itineraries precludes the use of discrete choice modeling. Nevertheless, large datasets and supervised machine learning methods provide opportunities to produce accurate and valuable forecasts of passenger flows.

For prediction, we leverage a dataset from the Official Airline Guide (OAG) that reports historical itinerary-level information in commercial aviation. A supplemental socioeconomic dataset provides population and GDP information for catchment areas surrounding the airports. Using these datasets, we engineer features that are valuable for predicting passenger flows. After evaluating several machine learning methods, we find that random forests yield the best results. To ensure the model is flexible to changes in the transportation network, we conduct stratified testing on new itineraries. We then deploy the best model to make passenger forecasts for itineraries across the period of slot assignment.

3.2 Prescriptive Modeling

Slot allocation is a highly complex combinatorial scheduling problem. The model must consider airline requests for every operation (i.e., arrival and departure) at an airport over a six-month season. The requests include the type of operation, date, time, and frequency for recurring operations (e.g., daily, weekly, weekdays). Airports may use various timescales (e.g., 5 minutes, 15 minutes, 60 minutes) to declare their capacities, and the limitations may differ depending on the type of operation. The WSAG defines the priority of slots and provides rules to maintain equal treatment for all requests in the same class. Rules regarding schedule regularity are intended to facilitate crew and aircraft planning. These guidelines create an interdependence across days in a season, thus considerably increasing the combinatorial complexities

of the slot allocation problem.

To balance the costs of slot allocation decisions, the optimizer must account for the perspectives of multiple stakeholders—namely airlines and passengers. When incorporating passenger-level metrics such as connection time, we find it is coupled with allocation decisions, thereby creating a nonlinearity due to the interdependence of two decision variables. The problem's computational complexity demands state-of-the-art algorithmic approaches to approximate the complete problem and achieve optimal solutions in acceptable times.

The proposed optimization model is an extension of the priority-based slot allocation model (PSAM) presented in [37]. This integer programming formulation includes constraints for airport declared capacities and the WSAG. We expand the PSAM formulation to minimize schedule displacement and two measures of passenger costs associated with slot allocation. We account for passenger costs using forecasts from the prediction model to measure increased or eliminated passenger connections due to displacement decisions. We develop a solution approximation approach based on coordinate descent and approximate the nonlinear objective using a weighted sum approach. We apply the ϵ-constrained method to balance the three objectives and generate solutions along the Pareto frontier.

The integration of the prediction and optimization models faces several challenges as there is no one-to-one mapping between the models. First, passenger itineraries are unknown at the time of initial slot allocation, which occurs months before the day of operations. This requires an itinerary construction procedure to establish all potential itineraries that can materialize from the airline slot requests. Second, the dataset used for the prediction model is aggregated by month, whereas the optimization model treats itineraries according to day and time. To disaggregate the forecasts, we test four variations of distributing the forecasted flows across the potential itineraries using a weight-based approach.

In the following chapters, we demonstrate the proposed approach using real-world data to produce a slot assignment for Singapore Changi Airport. Changi Airport is a major connection hub in Southeast Asia where many users (passengers and airlines)

depend upon connecting itineraries. For the predictive analytics, we use historical itinerary data from 2016-2019, with the final year reserved exclusively for model evaluation (i.e., test dataset). We find that random forest achieves excellent predictive performance, with a mean absolute error of 121 passengers per month. For prescriptive analytics, we use airline slot requests for the summer 2019 season (April-October). The proposed solution approach generates Pareto optimal solutions in a few hours of computation time. Analysis of the results indicates that schedule-coordinated airports can reduce passenger costs from slot allocation, with relatively small adjustments to schedule displacement. Additionally, sensitivity analysis of important modeling parameters indicates the benefits are robust to uncertainty.

Chapter 4

Predictive Modeling: Passenger Flow Forecasting

This chapter presents predictive analytics for passenger flow forecasting—a critical input for passenger-centric slot allocation optimization. The chapter begins with a description of the historical passenger itinerary and socioeconomic data used for prediction. We engineer numerical, categorical, and binary features that are strong predictors of passenger flows. We then train parametric and non-parametric prediction models and evaluate their performance over four years of data from Singapore Changi Airport. We report how the models perform overall, with a particular focus on new itineraries (i.e., new combinations of airports and airlines that did not exist in the training observations). Random forests achieve the best performance and are therefore selected to provide inputs for the optimization model in the subsequent chapter.

To define the prediction problem, let z be an itinerary serving an airport pair (o, d) (origin, destination). Itineraries can be either nonstop or connecting. Connecting itineraries involve transferring at an intermediate airport, denoted by h, between flight-legs operated by the same or different airlines. As such, an itinerary $z = (o, h, d, a_1, a_2)$ describes a combination of airports and airlines. For example, a nonstop itinerary from SIN to LHR operated by Singapore Airlines is denoted as $z_1 = (\text{SIN}, -, \text{LHR}, \text{SQ}, -)$; and a connecting one-stop itinerary from SIN

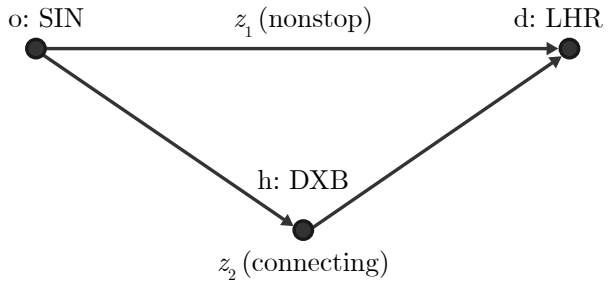

Figure 4-1: Visualization of the airports included in two example itineraries

to LHR through DXB with both legs operated by Emirates is denoted as $z_2 = (\text{SIN}, \text{DXB}, \text{LHR}, \text{EK}, \text{EK})$. Figure 4-1 provides an illustration of these itineraries. The dependent variable for the prediction problem is $y_{z,m}$, the number of passengers traveling on itinerary z in month m.

4.1 Predictive Data

The itinerary dataset contains information for commercial flights arriving or departing Singapore Changi Airport from 2016-2019. The initial dataset included 1.49M observations, but we use 343,829 observations for predictive modeling after filtering low-volume itineraries. Independent variables include the month, year, airport, airline, alliance membership, and airline classification as a low-cost carrier (LCC). The dependent variable is itinerary passengers, a count variable aggregated by month. The predicted passengers for itineraries in 2019 will be used as an input for optimizing slot allocation.

To supplement the itinerary data, we create socioeconomic variables using a geographic information system (GIS)-based procedure. For each airport, we define a circular catchment area with a radius of 100km using coordinates from the OpenFlights Airport Database. This enables the computation of relative population and GDP for each airport in the itinerary dataset based on high-resolution (30 arcsec) global spatial datasets. Table 4.1 describes the columns in the itinerary and socioeconomic datasets prior to feature engineering.

Table 4.1: Initial itinerary and socioeconomic datasets

Variables	Description
Itinerary Passengers	Dependent variable; number of passengers traveling on itinerary i defined by a combination of airports (o, h, d) and airlines, aggregated by month; constrained demand (i.e., minimum of supply and demand); Mean: 575; Std dev: 2,234; Range [30, 48,691]
Market Passengers	The total number of itinerary passengers that travel between a city-pair, including all airlines, route types (nonstop, connecting), and nearby airports in a month; Mean: 6,550; Std dev: 15,637; Range [30, 220,192]
Airports	Origin, hub (if connecting), destination (o, h, d); Count: 893
Airlines	Single airline if nonstop, two for connecting itineraries; Count: 234
Alliance Membership	Alliance membership of the primary airline; Star Alliance 47.7%, SkyTeam 8.2%, Oneworld 13.3%, Other 30.8%
LCC	Dummy variable classifying the primary airline as a low-cost carrier (LCC) 15.7%
Date	Month, Year [January 2016 - December 2019]
Distance	Flying distance (km) for each leg of the itinerary; Mean: 3,988; Std dev: 3,478; Range [101, 17,577]
Population	Values for origin and destination airports at 100km radii; Mean: 8.4M; Std dev: 7.3M; Range [1,219, 33.2M]
GDP	Values for origin and destination airports at 100km radii; gross domestic product (GDP) in USD, purchasing power parity (PPP); Mean: 301B; Std dev: 252B; Range [71M, 1,288B]

The presence of many low-volume itineraries in the OAG dataset creates a challenge for prediction. There are many seasonal or multi-stop itineraries that are operated infrequently and have low passenger counts (less than 30 passengers per month). Since the purpose of our prediction model is to provide accurate forecasts for medium and high-volume itineraries, we prioritize accuracy on those itineraries. Including low-volume itineraries while fitting our model would introduce noise into our predictions. We filter the dataset, removing itinerary observations with less than 30 passengers or involving more than one connection. The remaining dataset represents 23% of the initial observations but 97% of the passengers. This procedure reduces noise and improves our predictive performance on the significant itineraries.

4.2 Feature Engineering

Aggregation and decomposition are useful methods for reducing complexity in a system by adjusting the scale at which we conduct analysis [10]. While engineering features, the level of aggregation (or decomposition) of the predictors is informed by that of the dependent variable we aim to predict. Observations in the OAG dataset are aggregated by itinerary over a given month, so we focus on features near this level of aggregation.

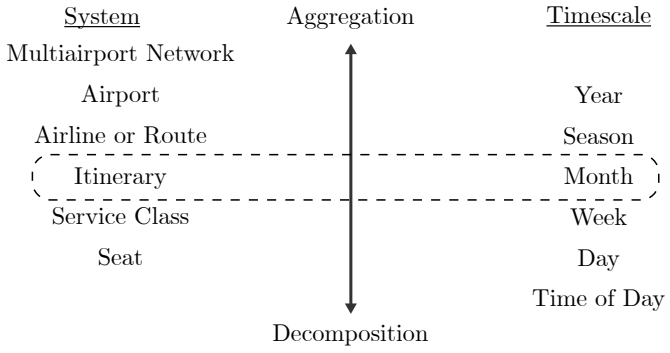

Figure 4-2: System aggregation and decomposition

Using the itinerary and socioeconomic datasets, we create features that are valuable for predicting passenger flows. We describe the most important features below, and a complete list of features tested is provided in Appendix A.

4.2.1 Historical Traffic

Historical metrics of passenger flows are useful predictors, especially in a developed transportation market such as Singapore. In general, the passengers observed for an itinerary will increase gradually from year to year, resulting from steady growth in demand and capacity. We create several lagging features, which leverage historical information, the most important of which is the lagging dependent variable (i.e., lagging itinerary passengers). This variable captures the historical passenger count for each itinerary (set of airports and airlines over a given month) in the previous year. Similarly, we create a lagging variable for the monthly market passengers corresponding to each itinerary in the previous year.

We also consider a lagging variable that captures the annual traffic at the origin and destination airports (o and d). Since this variable observes a higher level of aggregation than the variables for itinerary and market passengers, it is less precise, but the observations are more consistent. As commercial air travel exhibits a high level of symmetry (i.e., flows $o \to d$ are approximately equal to those from $d \to o$), we calculate the geometric mean of the traffic values at the origin and destination airports for each itinerary. Equation (4.1) calculates the annual airport traffic for itinerary i as a function of the traffic levels over the previous year ($\text{yr} - 1$).

$$\text{traffic}_{i,t} = \sqrt{\text{traffic}_{o,\text{yr}-1} \times \text{traffic}_{d,\text{yr}-1}} \tag{4.1}$$

Next, we consider a lagging variable for growth. Basic temporal variables (e.g., year and month) are weak proxies for growth, as they neglect differences between itineraries. To create a predictor of growth for each itinerary, we calculate the 12-month percent growth of the dependent variable between the two most recent years (i.e., the final two years of the training dataset). Equation (4.2) calculates the lagging percent growth for itinerary i as a function of the itinerary passenger counts y in the two previous years.

$$\text{lagging growth}_{i,m} = \frac{y_{i,m-12} - y_{i,m-24}}{y_{i,m-12}} \tag{4.2}$$

Since temporal patterns differ between airport pairs, month is a weak predictor as it is ignorant of the route. Instead, we use a seasonality variable that is specific for each airport pair (o, d). We calculate seasonality as the percent deviation from the annual mean passengers for the airport pair, denoted by μ_{od}. This standardization is carried out for each airport pair, such that an itinerary has a seasonality value of zero in month m if the mean monthly passengers traveling $o \rightarrow d$ in month $m-12$, denoted by $\nu_{od,m-12}$, is equal to $\mu_{od,\text{yr}-1}$. Negative and positive values denote deviations from the annual mean in the respective direction. Equation (4.3) calculates the seasonality for itinerary i as a function of the historical monthly and annual means, for a given airport pair (o, d).

$$\text{seasonality}_{i,m} = \frac{\nu_{od,m-12}}{\mu_{od,\text{yr}-1}} \tag{4.3}$$

A vulnerability of historical traffic variables is that they experience gaps whenever there is an inconsistent or new series, meaning the combination of airports and airlines did not exist during the same month in the previous year. For instance, at Changi Airport, approximately 10% of the itineraries each year are new (excluding low-volume itineraries), so the historical passenger count (i.e., lagging dependent variable) for those observations is zero. This provides an opportunity for a dummy variable (binary feature) that identifies if the itinerary is new. However, in the case of the remaining historical traffic variables (market passengers, airport traffic, passenger growth, and seasonality), there is not one sensible constant to apply when data is unavailable. For these variables, we found it best to impute the value using the other independent variables as predictors. After evaluating several imputation approaches, we applied k-nearest neighbors for imputation, as it yielded the best results during validation on the downstream prediction model.

4.2.2 Itinerary Characteristics

A dummy variable indicating whether the itinerary is nonstop or connecting is a simple yet influential predictor of passenger flows. In this dataset, we find that

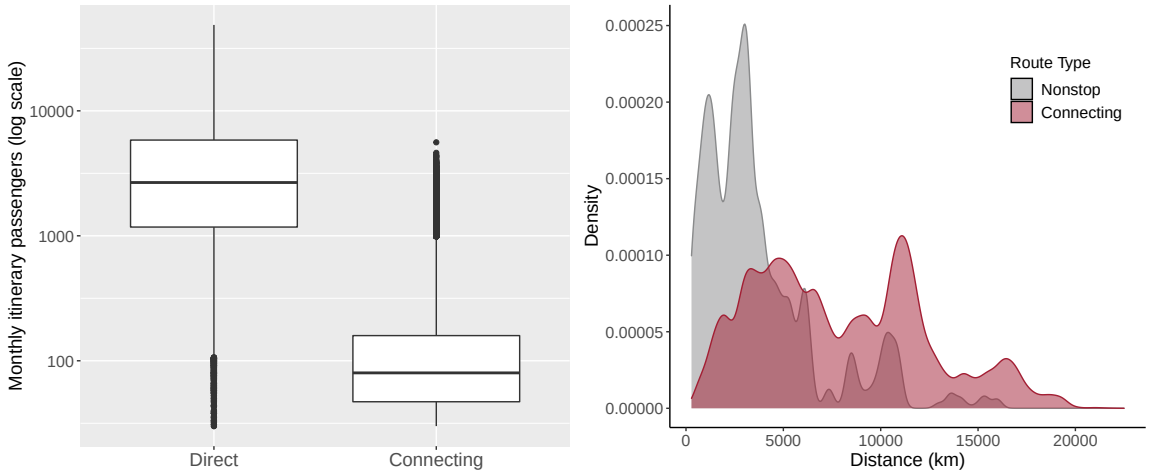

(a) Passengers by itinerary type

(b) Flying distance density plot

Figure 4-3: Descriptive feature visualizations

nonstop itineraries have higher passenger counts than their corresponding connecting itineraries (Figure 4-3a). There are many connecting options for every nonstop route, which increases competition.

Flying distance is a continuous feature that possesses a nonlinear relationship with passenger flows, as we observe in Figure 4-3b. At very short distances (under 500 km), the number of itinerary passengers is low due to competition with ground modes of travel. The density plot has a bi-modal distribution, with the first peak indicating flights to mainland Asia and the second peak representing flights to Europe and the Americas.

To account for the efficiency of connecting itineraries, we calculate a routing factor as described by [33]. The routing factor for a given itinerary is defined as the ratio of the flying distance to the corresponding nonstop distance (i.e., great circle distance). This calculation yields a routing factor of one for nonstop itineraries and a continuous value greater than one for connecting itineraries. Equation (4.4) calculates the routing factor for itinerary i as a function of the distances between each airport.

$$\text{routing factor}_i = \frac{\text{dist}_{oh} + \text{dist}_{hd}}{\text{dist}_{od}} \tag{4.4}$$

This variable serves as a proxy for the efficiency of a connecting itinerary. Routes

through hubs that are centrally located between the origin and destination have a routing factor marginally higher than one, whereas more inefficient routes yield higher values. Figure 4-4 displays two connecting itineraries, with similar $o \rightarrow d$ distances, but with different routing factors. BKK-SIN-CGK has a routing factor of 1.001 and a mean of 740 passengers per month across all airlines. In contrast, BKK-SIN-MNL has a routing factor of 1.726 and a mean of 54 passengers per month. These itineraries match the general trend that itineraries with lower routing factors carry more passengers than similar itineraries offering less efficient routes. This observation matches intuition since the most frequent travelers tend to be more time-sensitive.

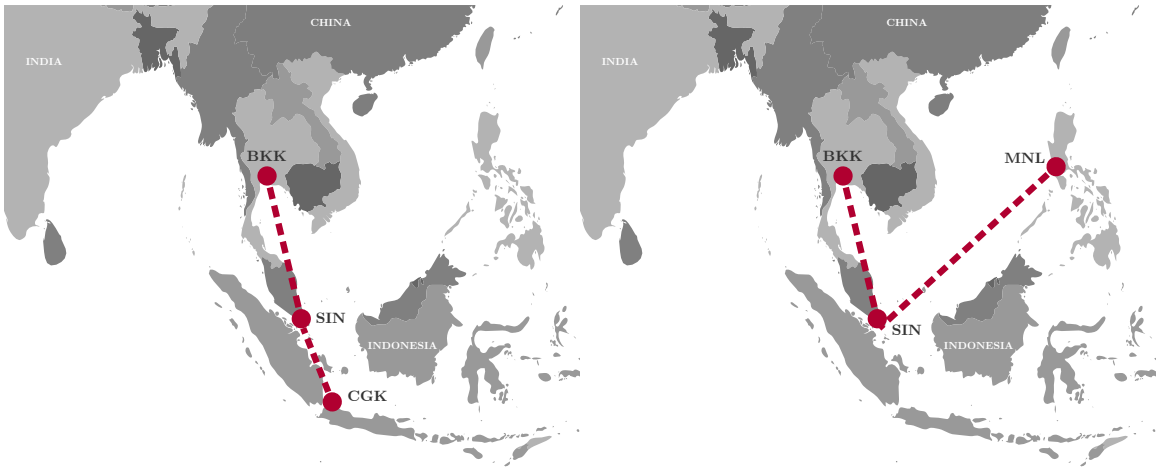

(a) Itinerary with 1.001 routing factor (b) Itinerary with 1.726 routing factor

Figure 4-4: Comparison of two connecting itineraries by routing factor

4.2.3 Socioeconomic

For socioeconomic features, we consider the population, gross domestic product (GDP), and GDP per capita of the areas surrounding the airports (o and d) for each itinerary. Again, since flows are symmetric, we limit the number of variables by considering the geometric mean of the origin and destination airports.

To elucidate interactions between airport pairs, we also consider the percent difference between the origin and destination GDP. With this approach, we intend to capture unique relationships in passenger flows between developed and undeveloped

markets. We find that higher passenger counts are correlated with itineraries that service airports with similar GDP levels.

4.2.4 Airlines and Airports

Airline and airport information is inherently available for all itineraries; however, they are difficult to leverage as features due to the vast number of different airports and airlines. The Changi itinerary dataset contains 893 and 234 unique airports and airlines, respectively. To create practical features, we explore several methods of grouping these variables into reasonable categories. Two methods of grouping airlines are by alliance membership (i.e., Star Alliance, SkyTeam, Oneworld) and class (LCC versus full-service). While effective for the 2016-2019 dataset, the LCC classification will need to be re-engineered over time as airline business strategies are undergoing transformation and many carriers do not fall neatly into one category. We also create dummy variables for the most frequently operating airlines at Singapore Airport: Singapore Airlines (SQ), Scoot (TR), and SilkAir (MI). We experimented with grouping airports geographically by origin and destination continent, but only one category (Asia to Asia) was significant for this dataset.

4.3 Prediction Models

Machine learning models are advantageous for forecasting passenger flows since they can leverage large datasets with numerous features and are flexible to evolution in the transportation network. Airlines create new itineraries as they adapt and grow, and machine learning models can leverage patterns—however inconsistent—within the features available to them. Given the problem size and user context, we focus on simple and scalable methods.

We begin with linear regression, including regularization methods such as subset selection, ridge regression [20], and least absolute shrinkage and selection operator (LASSO) [39]. Regularization is a set of automated techniques for feature selection that manage the trade-off between underfitting and overfitting. The aim is to bal-

ance in-sample performance and robustness in order to maximize out-of-sample performance. This is achieved by adding a penalty term to the ordinarily least-squares function used for regression.

Nevertheless, this prediction problem lends itself to models that can leverage non-linear relationships with variables that are available for new itineraries. For this reason, we also evaluate k-nearest neighbors (k-NN) for regression and random forests [9]. k-NN is a data-driven non-parametric method that predicts by identifying similar records (i.e., "neighbors") in the training data. Neighbors are determined by standardizing the features and calculating the Euclidean distance between itineraries.

Random forests is an ensemble method based on trees that segment the predictor space into regions. To create homogeneous regions, the model selects the feature and a threshold that will minimize impurity. Random forests generate diversified regression trees by randomly selecting subsets of features and training on different subsets of observations. Ensemble learning is grounded in the *wisdom of crowds*, whereby aggregate predictions of a group often outperform individual—even expert—predictions.

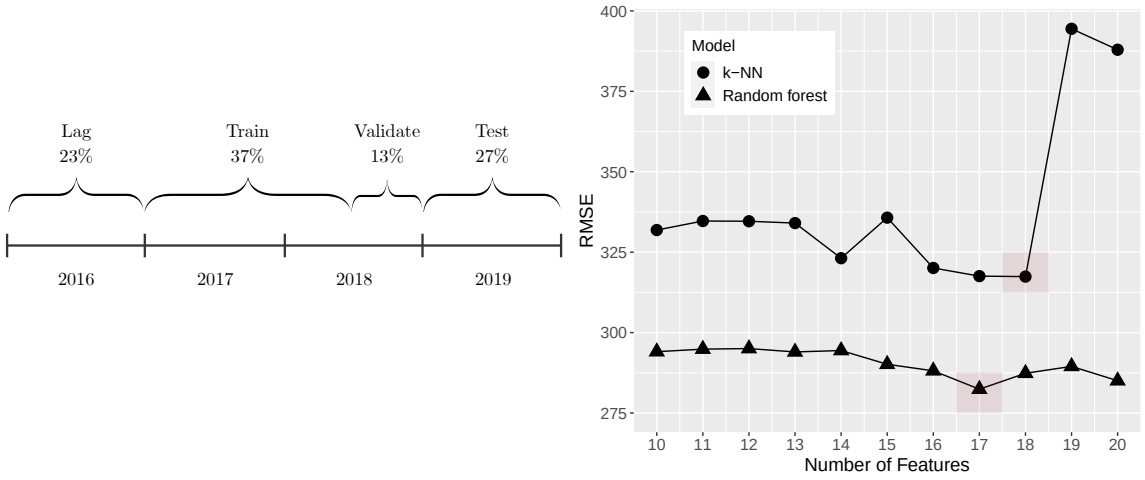

(a) Data Partitioning

(b) Feature selection for k-NN & random forests models on validation data

Figure 4-5: Predictive model validation

To avoid overfitting and assess the predictive performance of the models on unseen data, we first partition the data into training, validation, and test sets. Since we are

dealing with time series data, we partition the observations sequentially. The earlier period is for training and the later periods are for validation and testing, as shown in Figure 4-5a. Since the lagging features require data from the previous year, we are unable to use the first year in the dataset (2016) for modeling. We conduct 10-fold cross-validation using the validation set to select features and tune model parameters. After the model specifications are finalized, we retrain the best model from each class on the complete training and validation sets (2017-2018). Finally, we evaluate out-of-sample performance on the test set (2019), which provides a fair assessment of out-of-sample performance for model selection.

While k-NN and random forests do not have strong issues with collinearity as linear models do, we must still approach feature selection with caution, as they are susceptible to uninformative variables. Figure 4-5b illustrates the use of the validation set to identify the optimal subset of features. We find that k-NN and random forests achieve their best performance on the validation set with 18 and 17 features, respectively. As with the other methods, we seek the most parsimonious model that does not sacrifice predictive performance.

4.4 Prediction Results

After finalizing the model from each method, we evaluate out-of-sample performance to select the best model. We report the predictive performance metrics in Table 4.2. Due to continuous changes in the transportation network, a key requirement of our model is to predict well on new itineraries. Robustness to new combinations of itineraries is especially important for a fast-growing airport such as Singapore Changi. To assess this aspect of predictive performance, we perform a stratified evaluation on new itineraries.

To provide a benchmark for evaluation, we reference a naïve linear model that considers only an intercept value and one feature—the lagging dependent variable (i.e., historical itinerary passengers). Considering only each itinerary's historical passenger values, the naïve model generates an out-of-sample R^2 of 0.960, mean absolute

error (MAE) of 137 passengers per month, and root mean squared error (RMSE) of 436 passengers per month. When regressing on new itineraries, the lagging dependent variable is zero, and the model predicts the coefficient (55 passengers for the Changi dataset). Without additional features to inform the prediction, performance decreases significantly on new itineraries.

The regularized linear models (linear regression via subset selection, LASSO, and ridge regression) limit themselves to focus on three features: the lagging dependent variable, a dummy variable for connecting (versus nonstop), and a dummy variable indicating if the itinerary is new (versus continuing from the previous year). As with the naïve model, the remaining linear models perform well on the population as a whole (RMSE of 424), but performance deteriorates on new itineraries.

Table 4.2: Final model performance on test data

	Entire Population				Existing Itineraries			New Itineraries		
	c.v.	o.o.s.			o.o.s.			o.o.s.		
Method	R^2	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE
Naïve (LDV only)	0.956	0.960	137	436	0.974	129	392	0.119	206	709
Linear regression	0.961	0.963	144	424	0.974	129	400	0.392	271	589
Ridge regression	0.962	0.963	153	424	0.970	140	402	0.409	259	581
LASSO	0.962	0.963	153	424	0.970	140	402	0.409	259	581
k-NN	0.991	0.966	122	402	0.971	117	392	0.594	169	482
Random forests	0.998	0.969	121	388	0.972	120	386	0.708	129	408

We obtain the best performance on the test set with random forests. The random forests model yields excellent predictive performance with an out-of-sample R^2 of 0.969, mean absolute error (MAE) of 121 passengers per month, and root mean squared error (RMSE) of 388 passengers per month. Most noticeably, performance for this model drops only slightly, to an MAE and RMSE of 129 and 408 passengers per month respectfully, on new itineraries. In contrast to simpler methods, random forests provide the flexibility to leverage linear and nonlinear relationships with features that are available for new itineraries. Since many of the continuous features have complex nonlinear relationships with passenger flows, the random forests model is best able to leverage them to make precise predictions on itineraries involving new combinations of airports and airlines.

As a data-driven, non-parametric method, k-NN also performs reasonably well; however, it falls short of random forests on new itineraries, with an RMSE of 482 passengers per month.

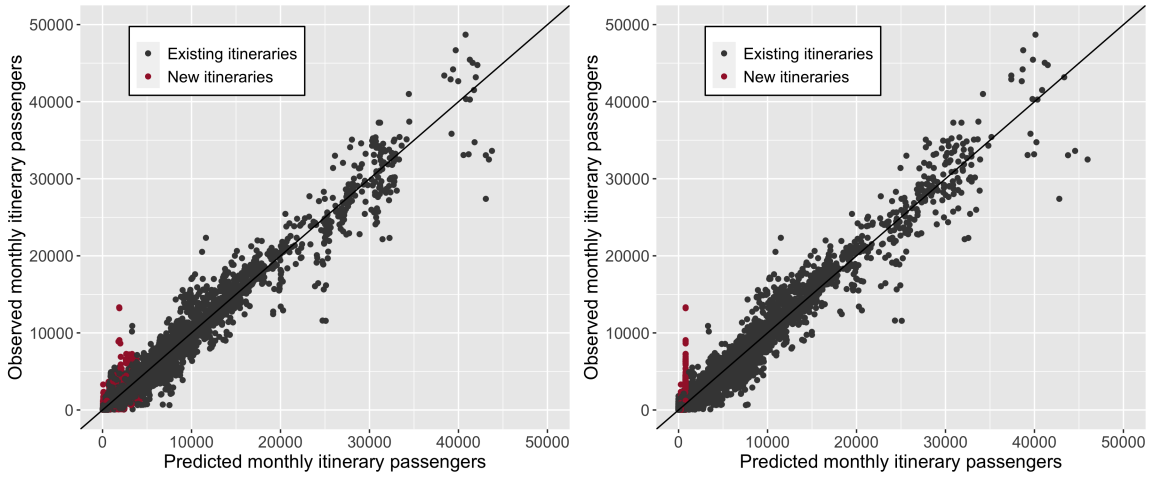

(a) Random Forests (b) Best Linear Model

Figure 4-6: Predicted versus observed values for new itineraries

The superiority of random forests becomes visible by plotting the predicted versus observed values and differentiating by new and existing itineraries (Figure 4-6). For random forests (Figure 4-6a), the values for new itineraries are scattered appropriately along the 45-degree line. In contrast, the best linear model (Figure 4-6a) exhibits a clear pattern of underestimating the number of passengers for new itineraries. The simplicity of the linear models—their dependence on the lagging dependent variable, supplemented only by two dummy variables—restricts their ability to make differentiated predictions for new itineraries. Random forests provide accurate predictions and the flexibility to account for new combinations of airports and airlines. Collectively, these results prompt us to select random forests for integration with the optimization model.

A common criticism of random forests, ensemble, and broadly machine learning models is that they are a "black box." This critique stems from the tension that exists between model accuracy and interpretability. In general, to improve accuracy, models grow increasingly complex, thereby decreasing interpretability. The field of game

theory offers a method to illuminate machine learning models by calculating Shapley Additive Explanation (SHAP) values [28]. SHAP values improve interpretability by attributing to each feature the change in the predicted value when conditioning on that feature.

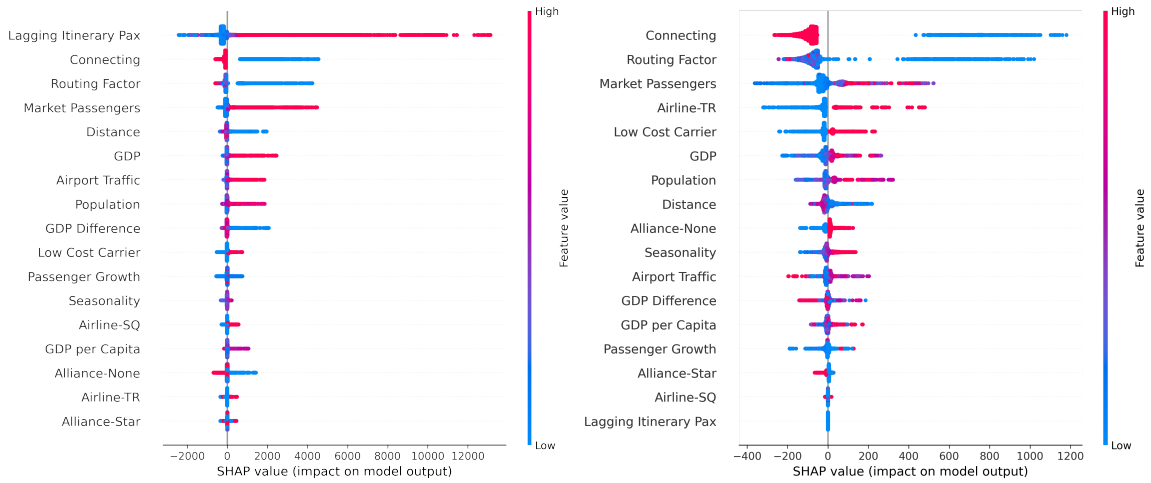

(a) Entire Population (b) New Itineraries

Figure 4-7: Feature importance and impact for random forests model

We plot the SHAP values for random forests in Figure 4-7. The features are placed in descending order by their importance to predictive performance. When evaluating the model on the entire test set population (Figure 4-7a), lagging itinerary passengers is the most important feature. However, Figure 4-7b illustrates that lagging itinerary passengers becomes the least important feature when predicting flows for new itineraries. This observation reinforces that models which rely upon historical passenger flows will perform poorly on new itineraries.

Applying a color scale to the SHAP values provides a reference for feature input values, and the horizontal axis depicts whether the effect is to increase or decrease the target value. For each feature, observations with relatively high input values are depicted in red, and observations with relatively low input values are depicted in blue. As an example, if the route is nonstop (Connecting = 0), the low input value will increase the predicted passenger count. Conversely, a connecting itinerary (Connecting = 1) will decrease the prediction. As another example, a high market

passenger value (shown in red) will increase the predicted value. These plots do not prove causality, but they do provide model interpretability by illustrating the correlation between model input and output values.

Chapter 5

Prescriptive Modeling: Multi-objective Slot Allocation Optimization

This chapter proposes an optimization formulation and solution approach for the slot allocation problem. The chapter begins with our methodology and a simple example to demonstrate the concept of including passenger considerations in slot optimization. We then explain the process for integrating the prediction and prescription models. We describe and provide a mathematical formulation for the proposed multi-objective integer programming model. The interdependence of two decision variables creates computational challenges, so we develop a solution approximation algorithm based on coordinate descent.

5.1 Methodology

As a resource scheduling problem, slot allocation is well suited for optimization and can be modeled with integer programming. The slot allocation problem is stated as: Given a set of airline requests for arrivals and departures over a season and a set of constraints representing the airport's resource limitations, identify an optimal combination of slot assignments, with the primary objective of minimizing the difference

in time between the requested and proposed schedules. As an unconstrained problem (i.e., without capacity constraints), each request would be granted as requested by the airlines. However, constraints enforce the realities of the airport's declared capacity, WSAG, and other operational considerations to moderate the demand profile at a level that keeps delays low.

As reviewed in Chapter 2, previous studies formulate the problem as an integer programming model that minimizes flight-level objectives (e.g., total displacement, maximum displacement, number of slots displaced or rejected). However, this myopic approach may lead to solutions that extend or eliminate passenger connections between flights, thereby increasing costs for passengers. In this section, we propose a new modeling approach that incorporates passenger considerations into slot allocation decision-making.

The proposed model minimizes three objectives in order of priority: total schedule displacement, infeasible passenger connections, and increase in passenger connection time. The primary objective, total displacement, is the most commonly accepted flight-level metric used by previous studies. Prioritizing total displacement promotes the acceptability of a solution by producing a schedule that is representative of the airlines' requests. The second objective advocates that important connections are not eliminated while conducting slot allocation. The third objective aims to maintain the efficiency of connections.

Incorporating passenger-level objectives (e.g., infeasible connections, connection time) to slot allocation creates interdependencies between flight-level decisions (e.g., flight displacement) and passenger-level considerations (e.g., aircraft capacities, passenger connection times). We find that capturing these interdependencies creates nonlinearities in the objective function, thereby considerably increasing the complexity of the optimization model. To overcome this challenge, we propose a novel algorithm based on coordinate descent that optimizes flight-level and passenger-level decisions sequentially and iterates until convergence. In the following chapter, we demonstrate the effectiveness of our formulation and algorithmic approach using real-world data from Singapore Changi Airport.

5.2 Proof of Concept

To introduce the effects of considering passengers during slot allocation optimization, we provide an illustrative toy example. The scenario involves five arrival and three departure requests, represented as arrows (5-1). Let us imagine these requests are submitted to a hub airport such as Changi.

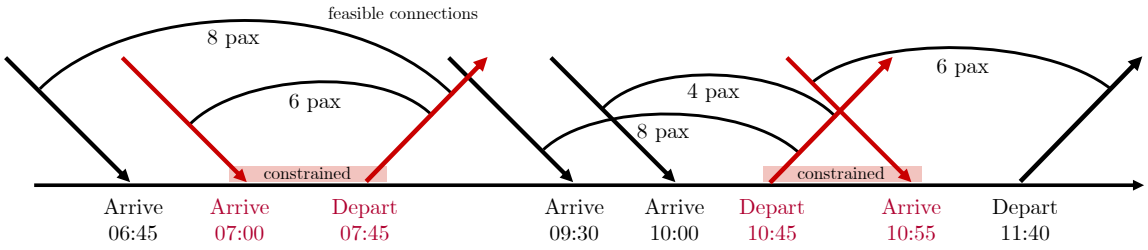

Figure 5-1: Example scenario for slot allocation optimization

There are desired connections between five flight pairs, represented as arcs. The forecasted number of passengers for each connection are displayed below the connection arc (e.g., "8 pax" corresponds to 8 passengers forecasted to make the connection). If the scheduler granted the slots as requested (i.e., ignoring capacity constraints), there would be 32 passenger connections. Note that we consider 45 minutes to be the minimum connection time (i.e., time between arrival and departure) to enable a passenger connection. However, let us also consider that two blocks of time are capacity constrained from itineraries in a higher priority class (e.g., historic slots). Therefore, the four movements that are requested during those times must be displaced outside of the constrained periods, indicated with red boxes. The following cases will illustrate solutions (i.e., slot schedules) under three different allocation strategies.

First, we evaluate the strategy commonly used by slot coordinators and existing slot allocation models, focusing exclusively on the minimization of schedule displacement measures such as total displacement, denoted as Case (A). To minimize displacement while respecting the capacity constraints, we displace each conflicting itinerary to the closest slot with available capacity. The resulting schedule is represented in Figure 5-2, with the displaced itineraries shown in blue. The resulting

schedule has a total displacement of 20 minutes, which is relatively low. However, since passenger-level metrics are not considered, 10 passenger connections become infeasible (less than 45 minutes), and the connections that remain feasible are increased by a total of 100 passenger minutes.

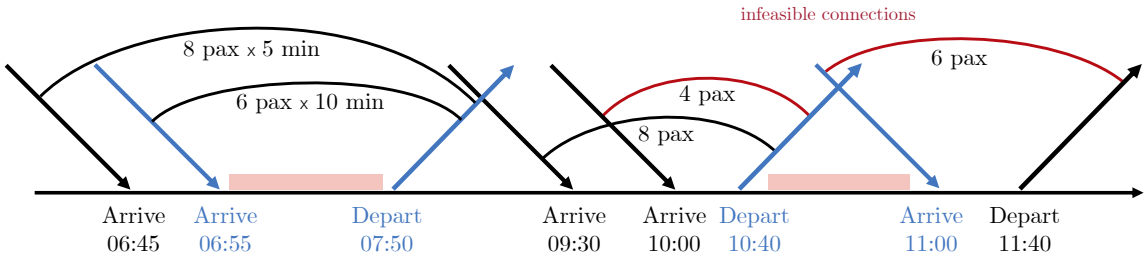

Figure 5-2: Minimizing displacement, Case (A)

Next, we evaluate an allocation strategy that minimizes displacement *and* infeasible connections, denoted as Case (B). To minimize displacement while maintaining connections, we displace conflicting itineraries in a direction that maintains their connections. This results in a schedule that has arrival and departure waves, with at least 45 minutes between connections to maintain their feasibility (Figure 5-3). The resulting schedule has a total displacement of 40 minutes, which is twice that obtained for Case (A). Since passenger connections are maintained, zero connections become infeasible. However, connection time is increased by 370 passenger minutes, as this solution sacrifices efficiency to preserve connections.

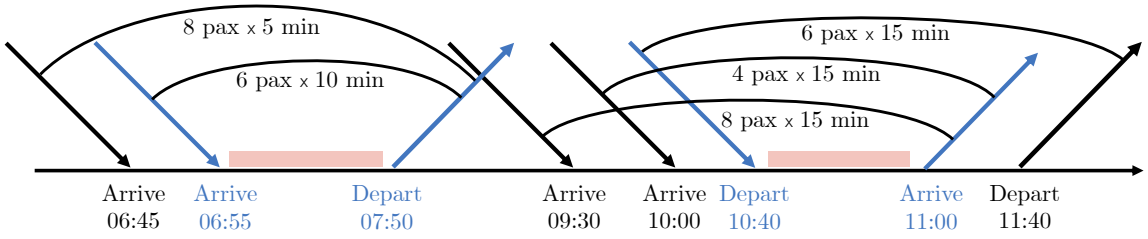

Figure 5-3: Minimizing displacement *and* infeasible connections, Case (B)

Last, we evaluate a strategy that minimizes displacement, infeasible connections,

and increasing connection time, denoted as Case (C). This strategy considers the number of passengers for each connection to determine how the conflicting itineraries should be displaced. Since the 10:45 departure receives only four connecting passengers from the 10:00 arrival, the optimizer chooses to sacrifice this connection rather than lengthen the connection time for the eight passengers who arrive at 09:30. The schedule that results from this objective function strikes a balance between the two previous examples (Figure 5-4). The resulting schedule has a total displacement of 30 minutes, which splits the difference between Case (A) and Case (B). Four passenger connections become infeasible, and connection time is increased by 190 passenger minutes. Table 5.1 summarizes the objective values under each strategy.

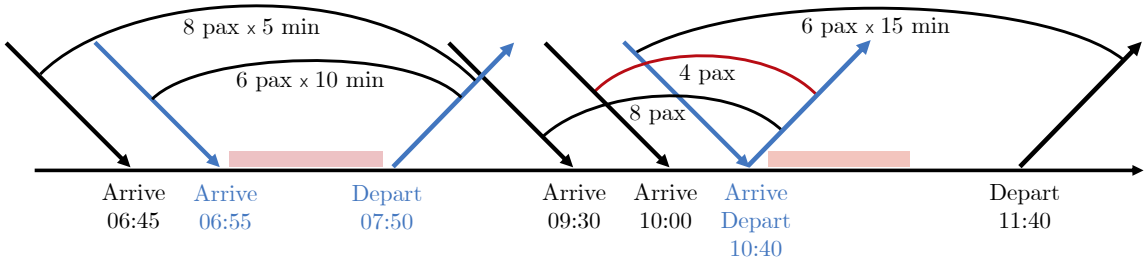

Figure 5-4: Minimizing displacement, infeasible connections, *and* increase in connection time, Case (C)

Table 5.1: Objective values for example scenario

Case	Objective Terms	Total Disp. (min)	Inf. Conn. (pax)	Inc. Conn. Time (pax-min)
A	Disp.	20	10	100
B	Disp. + Inf. Conn.	40	0	370
C	Disp. + Inf. Conn. + Inc. Conn. Time	30	4	190

This is a simplified example to demonstrate that minimizing flight displacement metrics alone does not necessarily translate into the most beneficial itineraries for passengers. To fully account for the costs associated with slot allocation, we must incorporate passenger considerations, but this is more complex than represented in the example. In reality, passengers may have alternatives if a connection between a

flight-pair is extended or eliminated. For example, in Case (C) (Figure 5-4), imagine that the departures at 10:40 and 11:40 are operated by the same airline alliance and have the same destination. Under these circumstances, the passengers arriving at 10:00 can be reaccommodated to the 11:40 departure. This new solution reduces infeasible connections to zero but increases the connection time by 240 passenger minutes. Aircraft capacity is another important consideration. If sufficient capacity does not exist to reaccommodate passengers on alternative flights, then those passenger connections will still be considered infeasible. Adding these realistic considerations increases the feasible region, creates interdependence amongst modeling decisions and involves trade-offs between objectives. The model must address these complexities to fully account for costs and create valuable solutions.

5.3 Integration

The proposed optimization model requires three primary inputs: (i) slot requests from airlines, (ii) airport declared capacities, and (iii) forecasted passenger flows. The third input, which enables the calculation of passenger costs associated with displacement decisions, requires integrating the prediction and optimization models. The first challenge with integration is that passenger itineraries are unknown at the time of initial slot allocation. An itinerary construction procedure is necessary to establish all potential itineraries that can exist given the set of slot requests. A second challenge is that the passenger flow forecasts are aggregated by month due to the aggregation of the historical itinerary data source (more detailed itinerary data for slot-coordinated airports is unavailable). Since the optimizer must assign passengers to specific operations (i.e., with date and time) to calculate cost, we must distribute the passenger forecasts over the potential itineraries.

5.3.1 Itinerary Construction

For optimization, itineraries are uniquely identified by the following features: origin airport, intermediate airport (if connecting), destination airport, airline(s), and

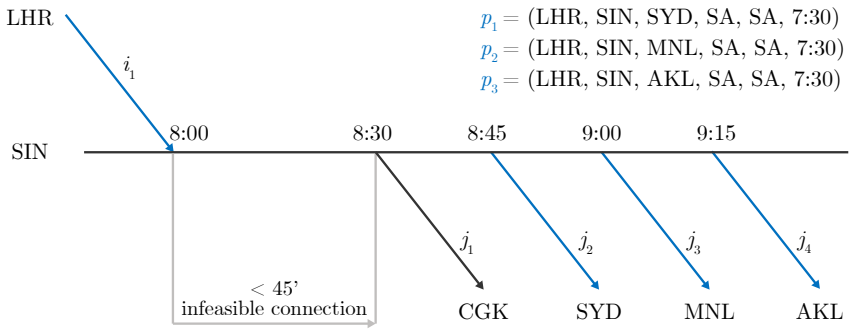

Figure 5-5: Itinerary construction example

reference time. For nonstop itineraries, reference time corresponds to the time of operation at the slot-coordinated airport (i.e., arrival or departure time). For connecting itineraries, reference time corresponds to the arrival time of the first flight leg.

To define the set of itineraries $\mathcal{P}$ (indexed p), we begin with the set of airline slot requests $\mathcal{S}$. The slot requests are indexed by i (or j), and each request represents a specific flight operation, characterized by arrival time (or departure time), airline, origin airport (or destination airport), and priority according to the WASG. Nonstop itineraries are self-evident, as any single slot request constitutes a feasible nonstop itinerary. We define the subset of nonstop itineraries as $\mathcal{P}^N \subset \mathcal{P}$.

To construct the set of connecting itineraries, defined as $\mathcal{P}^C \subset \mathcal{P}$, we pair arrival and departure slot requests (i, j) based on the following feasibility criteria: (i) arrival and departure flights operated by the same airline alliance, (ii) connection time between 45 minutes and 6 hours, and (iii) maximum routing factor of 1.5.

Figure 5-5 demonstrates how a single arrival slot request can generate many feasible connecting itineraries. In this example, the $\mathcal{S}^C(i, j) = \{(i_1, j_2), (i_1, j_3), (i_1, j_4)\}$ is the set of (3) feasible connections for the arrival slot request i_1. The 8:30 departure slot request j_1 does not create a feasible connecting itinerary with i_1 since it allows less than 45 minutes for transfer. While only a single arrival is represented in this example, each departure can similarly create multiple itineraries with different combinations of arrivals.

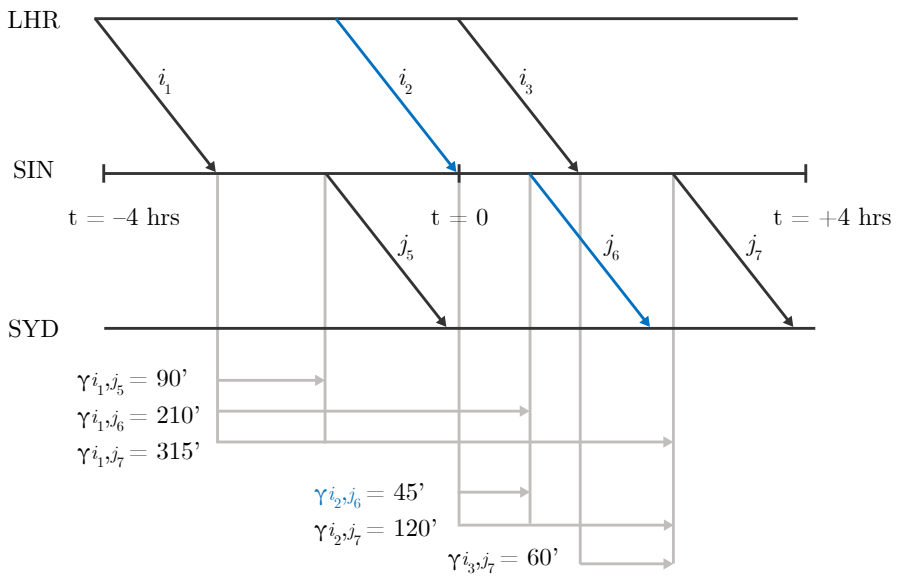

Figure 5-6: Identification of alternative itineraries example

To account for reaccomodating passengers, we define the set of all feasible alternatives for each itinerary p. This involves alternate single slot requests for nonstop itineraries $\mathcal{S}_p^N(i)$ or alternate pairs of slot requests for connecting itineraries $\mathcal{S}_p^C(i,j)$ that can accommodate passengers traveling on the itinerary p. We define these sets by assuming a time window of ± 4 hours; that is, given itinerary p, characterized by time t, we consider slot requests characterized by reference time $t' \in [t-4, t+4]$ as feasible alternatives.

Figure 5-6 provides an illustration of alternatives for itinerary p_1, constructed from the slot request i_2 (to SYD). In this example, (i_2, j_6) represents the minimum connection time alternative. If this alternative becomes infeasible due to displacement decisions or capacitated from reaccomodations, passengers can be recaptured by the following alternatives: $\mathcal{S}_p^C(i,j) = \{(i_1,j_5),(i_1,j_6),(i_1,j_7),(i_2,j_7),(i_3,j_7)\}$, assuming that they meet the feasibility criteria defined above.

To account for the costs of displacement and reaccomodation, we define the connection time of each alternative as $\gamma_{i,j}$. For each day d and itinerary p, we then define δ_{pd} as the minimum connection time across the set of alternatives $(i,j) \in \mathcal{S}_p^C(i,j)$ (e.g., (i_2, j_6) in Figure 5-6). Equation (5.1) expresses this relationship, where $\gamma_{i,j}$

gives the connection time between slot requests i and j, and B_{id} (B_{jd}) is equal to 1 if slot request i (j) is requested on day d, 0 otherwise. As will be described in Section 5.4, the optimization model will consider δ_{pd} and the set of alternatives $\mathcal{S}_p(i,j)$ to control and measure the cost of reaccommodating passengers on itinerary p.

$$\delta_{pd} = \min_{(i,j)\in\mathcal{S}_p^C(i,j)} \gamma_{i,j} B_{id} B_{jd} \tag{5.1}$$

5.3.2 Passenger Distribution

A major challenge for integrating the predictive and prescriptive models is a dichotomy in the unit of analysis surrounding passenger itineraries. Historical data on passenger itineraries are available at the monthly level, and so is our prediction. However, the optimization model considers passenger flows for each specific itinerary, corresponding to a specific day and time of day. To circumvent this challenge, we must distribute forecasted monthly passengers across the set of all possible itineraries $\mathcal{S}$ constructed from the slot requests. Using Equation (5.2), we perform a weighted distribution across all possible itineraries based on a weighting factor ζ_{pd}, which is forced to zero if itinerary p is not provided on day d.

$$N_{pd} = \frac{\zeta_{pd}}{\sum_{d\in D_m}\sum_{p\in P_z}\zeta_{pd}} y_{zm} \tag{5.2}$$

Where the notation is defined as:

Prediction Model

z = origin o, hub h, destination d, and airline(s)

y_{zm} = number of passengers forecasted to travel o, h, d in month m

Optimization Model

p = origin o, hub h, destination d, airline(s), and time t

N_{pd} = number of passengers forecasted to travel on itinerary p in day d

$\mathcal{M}$ = set of months indexed by m, $\{1, \ldots, M\}$

$\mathcal{D}_m$ = set of days in month m

ζ_{pd} = weight of itinerary p in day d

Equation (5.2) states that the number of passengers forecasted for each time-specific itinerary p on day d is given by a weighted distribution of the aggregate monthly forecast y_{zm} across all days the itinerary is requested. As such, the prediction model directs the aggregate forecast for each itinerary p across all days in a month, while the disaggregate time-specific forecasts N_{pd} are allowed to deviate from y_{zm} (with a zero-sum) based on the weight. Since we cannot estimate the weights empirically due to the data aggregation, we experiment with four different approaches.

1. **Uniform** - Monthly passengers for each itinerary p are distributed equally across all days d and times t that the itinerary is offered (i.e., uniform distribution). For example, if a given itinerary has a forecast of $y_{zm} = 1,000$ and the airline slot requests generate ten occurrences of the itinerary during a month, then each itinerary is allocated a forecast of $N_{pd} = 100$ passengers regardless of the day and time.

$$\zeta_{pd} = 1 \tag{5.3}$$

2. **Seat Capacity** - Monthly passengers are distributed proportionally according to the seat capacity requested for itinerary p on day d at time t. This method assumes the airlines have information regarding the demand profile for each requested flight (across days and time) and plan seat capacity accordingly.

$$\zeta_{pd} = \text{SeatCap}_{pd} \tag{5.4}$$

3. **Time-of-Day A** - Monthly passengers are allocated according to a time-of-day distribution empirically estimated in flight scheduling literature [34], that represents morning and evening departure peaks. We use α to supplement the weight ζ_{pd} according to the distribution. The more favorable periods of the day for departures (i.e., peak periods) are assigned higher α values.

$$\zeta_{pd} = \begin{cases} 1 & \text{if itinerary } p \text{ not operated during peak periods} \\ 1+\alpha & \text{if } p \text{ is operated during peak periods} \end{cases} \tag{5.5}$$

4. **Time-of-Day B** - This method uses the same structure as method 3, Equation (5.5), however we apply a distribution from choice modeling literature [2], in which only a morning departure peak is represented. In effect, the two time-of-day methods allocate greater portions of passengers to itineraries that are assumed to have higher utilities.

While the later methods have clear advantages over the uniform approach, they all exhibit weaknesses. The seat capacity method neglects the fact that demand is not the only factor airlines consider when pairing aircraft with slot requests. Additional considerations arise from fleet and network planning. The time-of-day distributions may have regional bias and do not capture day-of-the-week factors. In Section 6.4 we conduct sensitivity analysis of the four distribution methods. We found that the improvements provided by the proposed optimization model are not sensitive to the differences between these methods. We derived the experimental results reported in Section 6 using the seat capacity method.

5.4 Model Description

The proposed optimization model is an extension of the Priority Slot Allocation Model (PSAM) [37]. PSAM is an integer program with flight-centric objectives. In this section, we extend the PSAM formulation to consider two passenger-level objectives. Henceforth, the proposed model is referred to as the priority-based slot allocation model with passenger considerations (PSAM-Pax).

PSAM-Pax minimizes three objectives in order of priority: (1) total displacement, (2) infeasible passenger connections, (3) and increase in passenger connection time. Objectives (2) and (3) are *not* equivalent to maximizing the absolute number of connections and minimizing the absolute connection time, which would encourage solutions radically different than the airlines' requested schedules. Instead, the model uses the reference itineraries constructed in Section 5.3.1 to measure and minimize the passenger costs that result from slot allocation. In effect, the model explores minor deviations from solutions that minimize displacement to capture improvements to the passenger-level objectives. The model ensures solution feasibility through constraints that account for airport declared capacities, WASG (i.e., IATA guidelines), and passenger connections. The WASG constraints define rules for schedule regularity, aircraft turnaround times, and priorities across groups of slot requests: historic, change-to-historic, new entrant, and other. The change-to-historic group is further subdivided into "CR" and "CL" requests, necessitating additional constraints.

Minimize	Total Displacement + Infeasible Connections
	+ Increase in Connection time
Subject to	Airport capacities
	WASG constraints
	Passenger connections

In addition to assigning airline requests to slots, the model treats the assignment of passengers to itineraries as a decision variable. We assume that each passenger has an a priori reference itinerary assignment (and reference connection time if connect-

ing) established during the passenger distribution procedure, as described in Section 5.3.2. The reference passenger assignment corresponds to the unconstrained solution. Once constrained, the model might decide to reaccommodate passengers to alternate itineraries in conjunction with displacement decisions after considering all costs. For instance, if a reference connecting itinerary becomes infeasible due to displacement, then all passengers of that connection will be reaccomodated to an alternate itinerary or rejected if sufficient capacity on a feasible alternate itinerary is unavailable. In some cases, aircraft capacities may force the model to reaccommodate passengers to different itineraries.

5.5 Model Formulation

We now present the proposed model formulation, including a description of the sets, parameters, variables, objectives, and constraints.

5.5.1 Sets

$\mathcal{T}$ = set of time periods indexed by t, $\{1, \ldots, T\}$

$\mathcal{D}$ = set of days indexed by d, $\{1, \ldots, D\}$

$\mathcal{S}$ = set of slot requests, indexed by i and j, $\{1, \ldots, S\}$

$\mathcal{S}_{arr}/\mathcal{S}_{dep} \subset \mathcal{S}$ = subset of arrivals / departures

$\mathcal{S}_H \subset \mathcal{S}$ = subset of "historic" slots

$\mathcal{S}_{CH} \subset \mathcal{S}$ = subset of "change-to-historic" slots

$\mathcal{S}_{NE} \subset \mathcal{S}$ = subset of "new-entrant" slots

$\mathcal{S}_{OS} \subset \mathcal{S}$ = subset of "other" slots

$\mathcal{S}_{CR}/\mathcal{S}_{CL} \subset \mathcal{S}_{CH}$ = subset of "CR" / "CL" requests among change-to-historic slots

$\mathcal{N} \subset (\mathcal{S} \times \mathcal{S})$ = subset of pairs $(i \times j) \in \mathcal{S} \times \mathcal{S}$ such that there is an aircraft turnaround between i and j

$\mathcal{C}$ = set of capacity timescales, indexed by c (e.g. capacities for 15-minute periods, 60-minute periods, etc.), $\{1, \ldots, C\}$

$\mathcal{P}$ = set of passenger itineraries indexed by p, $\{1, \ldots, P\}$

$\mathcal{P}^N/\mathcal{P}^C \subset \mathcal{P}$ = subset of nonstop / connecting itineraries

$\mathcal{S}_p \subset \mathcal{S}$ = subset of slot requests $i \in \mathcal{S}$ that can accommodate passengers on itinerary $p \in \mathcal{P}$

$\mathcal{S}_p^C(i,j) \subset \mathcal{S}_p$ = subset of slot requests $i \in \mathcal{S}_p$ that can accommodate passengers on itinerary $p \in \mathcal{P}^C$

$\mathcal{S}_p^N(i,j) \subset \mathcal{S}_p$ = subset of slot requests $i \in \mathcal{S}_p$ that can accommodate passengers on itinerary $p \in \mathcal{P}^N$

PSAM-Pax adds to PSAM a new set of passenger itineraries $\mathcal{P}$, which is given by the itineraries that are built through the itinerary construction procedure. From $\mathcal{P}$ two subsets are considered: (i) the subset of nonstop itineraries $\mathcal{P}^N$ (i.e., the slot

coordinated airport is the origin or destination) and (ii) the subset of connecting itineraries $\mathcal{P}^C$ (i.e., the coordinated airport is used to connect). Within the subset of slot requests $\mathcal{S}$, we add two subsets to designate the slot requests that can accommodate passengers from itinerary $p \in \mathcal{P}$, specifically the subsets $\mathcal{S}_p^C$ and $\mathcal{S}_p^N$ for nonstop and connecting itineraries respectively.

5.5.2 Parameters

$$A_{it} = \begin{cases} 1 & \text{if slot } i \in \mathcal{S} \text{ is requested to operate no earlier than period } t \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases}$$

$$H_{it} = \begin{cases} 1 & \text{if slot } i \in \mathcal{S} \text{ was operated in the previous year no earlier than period } t \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases}$$

$$B_{id} = \begin{cases} 1 & \text{if slot } i \in \mathcal{S} \text{ is requested to operate on day } d \in \mathcal{D} \\ 0 & \text{otherwise} \end{cases}$$

$CR_{tdc}^{\text{dep}}/CR_{tdc}^{\text{arr}}/CR_{tdc}^{\text{tot}}$ = departure / arrival / total declared runway capacity at the airport in period $t \in \mathcal{T}$, day $d \in \mathcal{D}$, and timescale $c \in \mathcal{C}$

L_C = length of timescale $c \in \mathcal{C}$ (e.g., 15 minutes, 60 minutes)

$T^{\max}/T^{\min}$ = maximum / minimum change in the aircraft turnaround time between i and j, as compared to the requested aircraft turnaround time

ΔH_{ij} = historic aircraft turnaround time between i and j

ΔA_{ij} = requested aircraft turnaround time between i and j

Q_i = number of seats in slot request i

N_{pd} = number of passengers predicted to travel in itinerary $p \in \mathcal{P}$ on day $d \in \mathcal{D}$

δ_{pd} = reference connection time for itinerary $p \in \mathcal{P}_C$ on day $d \in \mathcal{D}$, given by the slot pair $i, j \in \mathcal{S}_p^C$ with the smallest connection time

σ_p = minimum feasible connection time for passengers traveling on itinerary $p \in \mathcal{P}$

PSAM-Pax adds four parameters to PSAM, providing information about the seat capacity of each slot request Q_i, the number of passengers traveling on each itinerary N_{pd}, the reference connection time δ_{pd}, and the minimum feasible connection time σ_p. For this study, we consider $\sigma_p = 45$ minutes.

5.5.3 Decision Variables

$$Y_{it} = \begin{cases} 1 & \text{if slot } i \in \mathcal{S} \text{ is rescheduled to operate no earlier than period } t \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases}$$

X_i^+ / X_i^- = displacement of slot $i \in \mathcal{S}$ if rescheduled to a later / earlier time

$$W_i^+ / W_i^- = \begin{cases} 1 & \text{if slot } i \in \mathcal{S} \text{ is displaced to a later / earlier time} \\ 0 & \text{otherwise} \end{cases}$$

Z_{pijd}^C = number of passengers on itinerary $p \in \mathcal{P}^C$ assigned to the arrival slot i and departure slot $j \in \mathcal{S}_p^C$ requested to operate on day $d \in \mathcal{D}$

Z_{pid}^N = number of passengers on itinerary $p \in \mathcal{P}^N$ assigned to the slot $i \in \mathcal{S}_p^N$ requested to operate on day $d \in \mathcal{D}$

Z_{pd} = number of passengers on itinerary $p \in \mathcal{P}$ assigned to the "sink" itinerary on day $d \in \mathcal{D}$; i.e., passengers without a feasible connection

$$\lambda_{pij} = \begin{cases} 1 & \text{if the connection } i, j \in \mathcal{S}_p^C \text{ is infeasible for itinerary } p \\ 0 & \text{otherwise} \end{cases}$$

τ_{pijd} = increase in connection time of passengers for itinerary $p \in \mathcal{P}^C$ on day $d \in \mathcal{D}$, if accommodated to arrival slot i and departure slots $j \in \mathcal{S}_p^C$

From the PSAM formulation, the first three decision variables specify the pairing of slot requests to slot times. PSAM-Pax adds the remaining variables to control three passenger-level decisions: the assignment of passengers from itineraries to slots, infeasible connections, and increase in connection time.

5.5.4 Objective Function and Constraints

The proposed optimization model, referred to as PSAM-Pax, is formulated as follows:

$$\min \quad \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}(X_i^+ + X_i^-)B_{id} + \sum_{p\in\mathcal{P}}\sum_{d\in\mathcal{D}} Z_{pd} + \sum_{p\in\mathcal{P}^C}\sum_{i,j\in\mathcal{S}_p^C}\sum_{d\in\mathcal{D}} \tau_{pijd} Z^C_{pijd} \tag{5.6}$$

$$\text{s.t.} \quad Y_{i1} = 1 \qquad \forall i \in \mathcal{S} \tag{5.7}$$

$$Y_{it} \geq Y_{i,t+1} \qquad \forall i \in \mathcal{S}, t \in \mathcal{T} \tag{5.8}$$

$$\sum_{t\in\mathcal{T}}(1 - A_{it})Y_{it} = X_i^+ \qquad \forall i \in \mathcal{S} \tag{5.9}$$

$$\sum_{t\in\mathcal{T}} A_{it}(1 - Y_{it}) = X_i^- \qquad \forall i \in \mathcal{S} \tag{5.10}$$

$$X_i^+ = X_i^- = 0 \qquad \forall i \in \mathcal{S}_H \tag{5.11}$$

$$X_i^+ \leq H_{it} - A_{it} \qquad \forall i \in \mathcal{S}_{CR}, t \in \mathcal{T} \tag{5.12}$$

$$X_i^- \leq A_{it} - H_{it} \qquad \forall i \in \mathcal{S}_{CR}, t \in \mathcal{T} \tag{5.13}$$

$$W_i^+ \geq Y_{it} - A_{it} \qquad \forall i \in \mathcal{S}_{CL}, t \in \mathcal{T} \tag{5.14}$$

$$W_i^- \geq -Y_{it} + A_{it} \qquad \forall i \in \mathcal{S}_{CL}, t \in \mathcal{T} \tag{5.15}$$

$$X_i^+ = (H_{it} - A_{it})W_i \qquad \forall i \in \mathcal{S}_{CL}, t \in \mathcal{T} \tag{5.16}$$

$$X_i^- = (A_{it} - H_{it})W_i \qquad \forall i \in \mathcal{S}_{CL}, t \in \mathcal{T} \tag{5.17}$$

$$\sum_{i\in\mathcal{S}_{\text{arr}}}\sum_{t=s}^{s+L_c}(Y_{it} - Y_{i,t+1})B_{id} \leq CR^{\text{arr}}_{sdc} \qquad \forall t \in \mathcal{T} \mid t < T - L_c + 1, d \in \mathcal{D}, c \in \mathcal{C} \tag{5.18}$$

$$\sum_{i\in\mathcal{S}_{\text{dep}}}\sum_{t=s}^{s+L_c}(Y_{it} - Y_{i,t+1})B_{id} \leq CR^{\text{dep}}_{sdc} \qquad \forall t \in \mathcal{T} \mid t < T - L_c + 1, d \in \mathcal{D}, c \in \mathcal{C} \tag{5.19}$$

$$\sum_{i\in\mathcal{S}}\sum_{t=s}^{s+L_c}(Y_{it} - Y_{i,t+1})B_{id} \leq CR^{\text{tot}}_{sdc} \qquad \forall t \in \mathcal{T} \mid t < T - L_c + 1, d \in \mathcal{D}, c \in \mathcal{C} \tag{5.20}$$

$$\Delta A_{ij} = \sum_{t\in\mathcal{T}}(A_{jt} - A_{it}) \qquad \forall (i,j) \in \mathcal{N} \tag{5.21}$$

$$\sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it}) - \Delta A_{ij} \geq T^{\min} \qquad \forall (i,j) \in \mathcal{N} \tag{5.22}$$

$$\sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it}) - \Delta A_{ij} \leq T^{\max} \qquad \forall (i,j) \in \mathcal{N} \tag{5.23}$$

$$\Delta H_{ij} = \sum_{t\in\mathcal{T}}(H_{jt} - H_{it}) \qquad \forall (i,j) \in \mathcal{N} \cap (\mathcal{S}_{CH} \times \mathcal{S}_{CH}) \tag{5.24}$$

$$\sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it}) \geq \min(\Delta A_{ij}, \Delta H_{ij}) \qquad \forall (i,j) \in \mathcal{N} \cap (\mathcal{S}_{CH} \times \mathcal{S}_{CH}) \tag{5.25}$$

$$\sum_{t\in\mathcal{T}}(Y_{jt}-Y_{it}) \le \max(\Delta A_{ij}, \Delta H_{ij}) \qquad \forall (i,j)\in \mathcal{N}\cap(\mathcal{S}_{CH}\times\mathcal{S}_{CH}) \qquad (5.26)$$

$$\sum_{t\in\mathcal{T}}(Y_{jt}-Y_{it}) - \sigma_p + M\lambda_{pij} \ge 0 \qquad \forall p\in\mathcal{P}^C, (i,j)\in\mathcal{S}_p^C \qquad (5.27)$$

$$Z^C_{pijd} \le N_{pd}(1-\lambda_{pij}) \qquad \forall p\in\mathcal{P}^C, (i,j)\in\mathcal{S}_p^C, d\in\mathcal{D} \qquad (5.28)$$

$$\sum_{i,j\in\mathcal{S}_p^C} Z^C_{pijd} + \sum_{i\in\mathcal{S}_p^N} Z^N_{pid} + Z_{pd} = N_{pd} \qquad \forall p\in\mathcal{P}, d\in\mathcal{D} \qquad (5.29)$$

$$\sum_{p\in\mathcal{P}^C}\sum_{j\in\mathcal{S}_p^C} Z^C_{pijd} + \sum_{p\in\mathcal{P}^D} Z^N_{pid} \le Q_i \qquad \forall i\in\mathcal{S}_{\text{arr}}, d\in\mathcal{D} \qquad (5.30)$$

$$\sum_{p\in\mathcal{P}^C}\sum_{i\in\mathcal{S}_p^C} Z^C_{pijd} + \sum_{p\in\mathcal{P}^D} Z^N_{pjd} \le Q_j \qquad \forall j\in\mathcal{S}_{dep}, d\in\mathcal{D} \qquad (5.31)$$

$$\tau_{pijd} \ge \sum_{t\in\mathcal{T}}(Y_{jt}-Y_{it}) - \delta_{pd} \qquad \forall p\in\mathcal{P}^C, (i,j)\in\mathcal{S}_p^C, d\in\mathcal{D} \qquad (5.32)$$

$$\tau_{pijd} \ge 0 \qquad \forall p\in P^C, (i,j)\in S_p^C, d\in D \qquad (5.33)$$

$$X_i^+, X_i^- \in \mathbb{N}_0 \qquad (5.34)$$

$$Y_{it}, W_i^+, W_i^- \in \{0,1\} \qquad (5.35)$$

Equation (5.6) formulates the multi-objective function, which minimizes three terms: (1) total displacement, (2) number of infeasible passenger connections, and (3) increase in connection time. Objective (3) is nonlinear due to the multiplication of two decision variables $\tau_{pijd}Z^C_{pijd}$ (the increase in connection time and number of passengers assigned to each itinerary), thereby significantly increasing the model's computational complexity. In section 5.7 we demonstrate a solution approach that approximates the nonlinear formulation with a sequence of mixed-integer programs.

Constraints (5.7) through (5.26) are taken from the PSAM model [37]. Constraint (5.7) ensures that all slots are assigned to a time period. Constraint (5.8) ensures that Y_{it} is non-increasing in t. Constraints (5.9) and (5.10) define the logical relationships between the displacement variables (i.e., each slot request can be scheduled at the requested time, displaced to a later slot, or displaced to an earlier slot). These constraints provide valid equalities that provide tight linear programming relaxations [37]. Constraint (5.11) ensures that historic slots are not displaced. Constraints (5.12) and (5.13) specify that change-to-historic slots with a "CR" code are assigned to a time slot between the historical and the requested time slots. Constraints (5.14) to

(5.17) specify that change-to-historic slots with a "CL" code are assigned to either the historic time slot or the requested time slot. Constraints (5.18) to (5.20) ensure that the number of arrivals, departures, and total number of movements, respectively, scheduled in any time period does not exceed the corresponding runway capacities. Constraints (5.21) to (5.23) ensures that the aircraft turnaround time (i.e., arriving and departing flights using the same aircraft) does not increase or decrease by more than the allowable limits. Constraints (5.24) to (5.26) ensures that the aircraft turnaround time between two change-to-historic slots remains between the requested aircraft turnaround time (A_{ij}) and the historic aircraft turnaround time (H_{ij}), as specified by the WASG.

Constraints (5.27) to (5.33) are the additional constraints necessary to incorporate passenger-level metrics. Constraint (5.27) defines the feasibility of connections between slot pairs. Note that λ_{pij} is a decision variable that is equal to 1 if a connection between slots i and j is infeasible, and 0 otherwise. The parameter M denotes a large number whose minimum value is $T + \sigma_p$. Constraint (5.28) ensures that passengers are only assigned to feasible connections. For example, from constraint (5.27), if connection ij is infeasible (i.e., $\sum_{t \in \mathcal{T}}(Y_{jt} - Y_{it}) - \sigma_p < 0$) then $\lambda_{pij} = 1$. Consequently from constraint (5.28), $Z^C_{pijd} \leq 0$, meaning that zero passengers from itinerary p are assigned to connection ij on any day. Constraint (5.29) ensures that all passengers are assigned to slots (connecting or nonstop) if possible, otherwise to the "sink" option. Constraints (5.30) and (5.31) ensure that aircraft seat capacities are not exceeded for arrival and departure slot requests, respectively. Finally, constraint (5.32) computes the increase in connection time of connecting itineraries. The decision variable τ_{pijd} is given by the difference between the allocated connection time $\sum_{t \in \mathcal{T}}(Y_{jt} - Y_{it})$ and the reference connection time δ_{pd} for itinerary $p \in \mathcal{P}$. We restrict τ_{pijd} to non-negative values to ensure the model does not benefit from minimizing connection time beyond the requested reference connections. Constraints (5.33) to (5.35) define the domain of the decision variables.

5.6 Size of the Formulation

We compare the PSAM-Pax model to the baseline PSAM model, which considers total displacement exclusively and omits passenger-level objectives and constraints. With the addition of passenger variables and constraints, the size of PSAM-Pax is significantly larger than PSAM. Table 5.2 reports the number of variables and constraints for both models. The size of PSAM-Pax scales up with the number of passenger itineraries resulting in an additional PS^2 binary variables, $PS^2D + PSD + PD$ continuous variables, and $PS^2 + 3PS^2D + 2SD + PD$ constraints relative to PSAM. Note that, in theory, P scales up quadratically with the number of slot requests S. However, the set of passenger itineraries $\mathcal{P}$ is sparse because only a small subset of flight pairs in $\mathcal{S} \times \mathcal{S}$ creates feasible connecting itineraries, according to the criteria in Section 5.3.1.

Table 5.2: Size of the PSAM and PSAM-Pax models

Metric	PSAM	PSAM-Pax
# binary variables	$ST + 2S$	$ST + 2S + PS^2$
# integer variables	$2S$	$2S$
# continuous variables	$-$	$PS^2D + PSD + PD$
# constraints (upper bound)	$4S + 7TS + 3TDC + 6S^2$	$4S + 7ST + 3TDC + 6S^2 + PS^2 + 3PS^2D + 2SD + PD$

5.7 Solution Approximation

The nonlinear combination of decision variables $\tau_{pijd} Z^C_{pijd}$ in Objective (3) (i.e., increase in passenger connection time) causes challenges for solving PSAM-Pax. Since the variables are binary and continuous, we explored the possibility of linearizing the term by replacing it with a single variable and adding a set of big-M constraints. However, this transformation increased the size of the problem, and the model was unable to solve realistic problem instances in acceptable computation times. Therefore, we propose a solution approach that applies a coordinate descent procedure to approximate the optimal solution of PSAM-Pax. The procedure relies on the following three steps:

1. **Minimize displacement** - The flight-centric model, referred to as PSAM, is solved as a single objective optimization without passenger considerations. Slots are allocated while minimizing total displacement exclusively, corresponding to PSAM-Pax Objective (1), with the objective value denoted as X^*. While this step produces a complete solution in terms of a slot schedule, the PSAM-Pax objective values are only partially identified since passenger costs are unknown.

2. **Account for passenger costs** – A passenger-centric model, referred to as Pax, is solved in this step. Given the schedule solution obtained in Step (1), passengers are assigned to a slot request or slot request pair (for nonstop and connecting itineraries respectfully) while minimizing PSAM-Pax Objectives (2) and (3). A weight is applied to prioritize (2) infeasible passenger connections over (3) increase in passenger connection time. Displacement remains fixed to the solution from Step (1). Since slots remain fixed according to the Step (1) solution the only constraints considered are: (5.28)–(5.31). These constraints ensure passengers are only allocated to feasible connections and aircraft capacities are enforced. Using the variable Y_{it}, we compute λ_{pij} and τ_{pijd}, which are treated as fixed parameters during this step. The decision variables during this step are Z^C_{pijd}, Z^N_{pid}, and Z_{pd}. This step results in an initial feasible solution for PSAM-Pax that minimizes total displacement. The value for total displacement remains X^* from Step (1), and the corresponding number of infeasible passenger connections is denoted as Z^* and the total increase in passenger connection time Γ^*.

3. **Solve PSAM-Pax approximation** – The schedule solution obtained from the previous steps is used as a warm start for the complete multi-objective PSAM-Pax formulation. The ϵ-constrained method is applied, and Objectives (1) and (2) are implemented as constraints. The objective values X^* and Z^*, obtained from Steps 1 and 2 respectfully, specify the objective values of reference for the first and second terms, that is, (1) the total displacement and (2) number of infeasible passenger connections. We then consider θ as the allowable percentage

increase in total displacement and ϕ as the allowable percentage decrease in the number of infeasible passenger connections. Equations (5.36) and (5.37) are added to PSAM-Pax formulation to constraint the two objectives. By looping over vectors of θ and ϕ, denoted as V^{θ} and V^{ϕ}, we identify solutions along the Pareto frontier.

$$\sum_{i \in S} \sum_{d \in D} (X_i^+ + X_i^-) B_{id} \leq \theta X^* \tag{5.36}$$

$$\sum_{p \in P} \sum_{d \in D} Z_{pd} \leq \phi Z^* \tag{5.37}$$

As such, PSAM-Pax becomes a single-objective minimization of the increase in passenger connection time, given by the product of two decision variables: (i) τ_{pijd}, the increase in connection time for itinerary $p \in \mathcal{P}^C$ on day $d \in \mathcal{D}$, if accommodated to slots i and $j \in \mathcal{S}_p^C$, and (ii) Z_{pijd}^C, the number of passengers on itinerary $p \in \mathcal{P}^C$ accommodated to the arrival slot i and departure slot $j \in \mathcal{S}_p^C$ requested to operate on day $d \in \mathcal{D}$. To reduce the computational expense, we replace the nonlinear objective function with the following weighted sum expression:

$$\min \alpha \sum_{p \in P^C} \sum_{i,j \in S_p^C} \sum_{d \in D} \tau'_{pijd} Z_{pijd}^C + (1 - \alpha) \sum_{p \in P^C} \sum_{i,j \in S_p^C} \sum_{d \in D} \tau_{pijd} Z_{pijd}^{\prime C} \tag{5.38}$$

Where τ'_{pijd} and $Z_{pijd}^{\prime C}$ are approximations for τ_{pijd} and Z_{pijd}^C, respectively, obtained by Step (2), and α is a calibration parameter. Note that, when $\alpha = 1$, then the second term of Equation (5.38) becomes zero, and the increase in connection time of itinerary connections τ_{pijd} is approximated by the values computed in Step (2). This may provide an adequate approximation for small increases in displacement since, in this case, slots will not deviate much from the PSAM solution (i.e., displacement is minimized). However, when larger increases are allowed, then more flexibility to slot displacements is given, and therefore the value of τ_{pijd} computed in Step (2) may differ significantly from

the optimal value of the nonlinear PSAM-Pax. Conversely, when $\alpha = 0$ the first term of Equation (5.38) becomes zero, and the number of passengers assigned to each connection Z^C_{pijd} is approximated by the values computed in Step (2). This tends to provide impractical solutions since passengers can freely be assigned to remote slots at no cost. To evaluate the impact of small values of alpha, we consider a minimum value of 0.001. In Section 6.4 we present results that indicate the best solutions are achieved with intermediate α values, as they provide a balance between the two approximated terms.

Step (3) may be repeated over multiple iterations, using the convergence of Γ^* or maximum iterations as stopping criteria. The solution obtained by the previous iteration is used for the warm start, and the parameters τ'_{pijd} and Z'^C_{pijd} are updated. Note that we are iterating over multiple mixed-integer linear programs to solve a non-convex (bilinear) mixed-integer quadratic program. However, we find that the benefits of iteration are marginal, as a near-optimal solution is identified by the first approximation.

Algorithm 1 summarizes our coordinate descent procedure to solve PSAM-Pax. Corresponding to Step (3), the for-loop builds the Pareto frontier by looping through a vector of ϵ-constraints for Objectives (1) and (2), while the approximation of Objective (3) is minimized. Within this process, the while-loop iterates over the approximation, seeking improvements to the solutions.

Algorithm 1 Algorithm to solve PSAM-Pax approximation and identify solutions along the Pareto frontier

- Inputs: airport declared capacities, slot requests, passenger itineraries, passenger flow predictions, calibration parameter α, vectors V^θ and V^ϕ establishing the ϵ-constrained conditions
- Solve PSAM (PSAM-Pax Objective 1):

$$\min_{X^+,X^-} \sum_{i\in\mathcal{S}} \sum_{d\in\mathcal{D}} \left(X_i^+ + X_i^-\right) B_{id} \quad \text{s.t.: Eq. (5.7)–(5.26), (5.34)–(5.35)}$$

- Update the variable $Y_{it}\ \forall i \in S, t \in T$ and total displacement X^*
- Using Y_{it}, estimate passenger costs:

$$\lambda_{pij} = \begin{cases} 1 & \text{if} \quad \sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it}) - \sigma_p > 0 \\ 0 & \text{otherwise} \end{cases} \qquad \text{(Equation (5.27))}$$

$$\tau_{pijd} = \max\left(\sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it}) - \delta_{pd}, 0\right) \qquad \text{(Equations (5.32) and (5.33))}$$

- Solve Pax (PSAM-Pax Objectives 2 & 3), where M is a large number and Y_{it}, λ_{pij}, and τ_{pijd} are treated as parameters:

$$\min_{Z_{pd},Z^C_{pijd}} M \sum_{p\in\mathcal{P}} \sum_{d\in\mathcal{D}} Z_{pd} + \sum_{p\in\mathcal{P}^C} \sum_{i,j\in\mathcal{S}_p^C} \sum_{d\in\mathcal{D}} \tau_{pijd} Z^C_{pijd} \quad \text{s.t.: Eq. (5.28)–(5.33)}$$

- Update the number of passengers assigned to each itinerary $Z^C_{pijd}, Z^N_{pid}, Z_{pd}$, the total number of infeasible passenger connections Z^*, and the total increase in passenger connection time Γ_1^*
- Initialize $n = 1$, $Z'^C_{pijd} = Z^C_{pijd}$, and $\tau'_{pijd} = \tau_{pijd}$

  for $\theta = V^\theta$ and $\phi = V^\phi$ **do**

  while $\Gamma_n^* \neq \Gamma_{n-1}^*$ or $n < n_{\max}$ **do**

  - Solve PSAM-Pax approximation (Equation 5.38):

$$\min_{Z^C_{pijd},\tau_{pijd}} \alpha \sum_{p\in P^C} \sum_{i,j\in S_p^C} \sum_{d\in D} \tau'_{pijd} Z^C_{pijd} + (1-\alpha) \sum_{p\in P^C} \sum_{i,j\in S_p^C} \sum_{d\in D} \tau_{pijd} Z'^C_{pijd} \quad \text{s.t.: Eq. (5.7)–(5.37)}$$

  - Update the total increase in passenger connection time Γ_n^* and PSAM-Pax decision variables $D_n^{\theta\phi}\left(Y_{it}, Z^C_{pijd}, Z^N_{pid}, Z_{pd}, \lambda_{pij}, \tau_{pijd}\right)$
  - Update $Z'^C_{pijd} = Z^C_{pijd}$ and $\tau'_{pijd} = \tau_{pijd}$
  - $n = n + 1$

  end while

  - Return $D_n^{\theta\phi}$ with the minimum Γ_n^*

  end for

Chapter 6

Experimental Results

In this chapter, we implement the PSAM-Pax approximation, as described in Section 5.7, at Singapore Changi Airport using slot requests for the 2019 summer season (April-October). The results indicate the potential for PSAM-Pax to provide better solutions than a model focused exclusively on schedule displacement. By controlling the trade-off between objectives, we use PSAM-Pax to analyze the effects of increases in displacement on the passenger-level objectives. Recommended solutions are emphasized with bold font in the tables and a shaded region on the Pareto frontier. A synthesis of the benefits provided by these solutions is provided in Section 6.3.

6.1 Computation Performance

Solving the complete slot allocation problem for a season at one of the busiest schedule-coordinated airports requires considerable computational resources. To reduce the expense while still providing meaningful results, we consider only a single declared capacity: total movements for one hour. In reality, Changi has additional capacities for different types of movements (i.e., arrivals versus departures) and for different timescales (e.g., 5 minutes, 15 minutes). Considering these additional capacities increases computation time beyond 24-hours. Since the single capacity version of the problem is less constrained, there will be less displacement, and therefore should provide a lower bound on the benefits of considering passenger-level metrics.

Under single-capacity conditions, the computation time for Step (1) of the solution approximation (i.e., PSAM) is approximately 2 hours. Computation time for the remaining steps (2) and (3) ranges from 36 to 523 minutes, depending on the modeling constraints. The PSAM solution only needs to be computed once to provide an initial feasible solution, so the tables in this chapter report only the additional computation time necessary for PSAM-Pax.

As shown by the following tables, the model requires less time to return solutions when the ϵ (i.e., objective) constraints are more strict. As the objective constraints are relaxed, the feasible region becomes larger, thereby increasing the space the algorithm must search for solutions. We found that the best solutions correspond to configurations that have a computation time of 64 to 201 minutes. The results provided in this section are obtained without iterating over Step 3, as the benefits of iteration were marginal.

6.2 Main Results

To understand the relationship between the three PSAM-Pax objectives—(1) total displacement, (2) infeasible passenger connections, (3) and increase in passenger connection time—we solve PSAM-Pax while using the ϵ-constrained method to allow for incrementally higher levels of schedule displacement. Table 6.1 reports the objective values calculated relative to the unconstrained problem and percent changes in comparison to the baseline PSAM solution. For these results, we weight Objective (2) much higher than Objective (3). The first row provides a reference for comparison as it describes the PSAM solution (obtained by Step (1) of the solution approximation), which minimizes displacement exclusively without considering passenger objectives. While the PSAM solution minimizes displacement, 2065 passenger connections become infeasible, and connection time is increased by 1,826,660 passenger minutes.

In the subsequent rows of Table 6.1, displacement is relaxed from the baseline to capture improvements in Objectives (2) and (3). For the second row, we constrain displacement to the same value generated by Step (1), and optimize the two passenger

Table 6.1: Objective values for PSAM-Pax solutions compared to the baseline PSAM solution

Model	Tot. Disp. (min)		Inf. Conn. (pax)		Inc. Conn. Time (pax-min)		CPU Time (min)
PSAM	157,390	0%	2,065	0%	1,826,660	0%	-
PSAM -Pax	157,390	+0.00%	2,057	-0.39%	1,826,660	0.00%	46
	157,783	+0.25%	1,714	-17.00%	1,819,155	-0.41%	48
	158,177	+0.50%	1,352	-34.53%	1,518,490	-16.87%	71
	158,964	**+1.00%**	**992**	**-51.96%**	**1,534,890**	**-15.97%**	**61**
	160,538	+2.00%	871	-57.82%	1,623,690	-11.11%	48
	163,686	+4.00%	406	-80.34%	1,751,925	-4.09%	69
	166,833	+6.00%	148	-92.83%	1,666,485	-8.77%	65
	169,981	+8.00%	106	-94.87%	1,991,910	+9.05%	103
	173,129	+10.00%	46	-97.77%	2,051,855	+12.30%	114
	180,999	+15.00%	0	-100.00%	1,508,095	-17.44%	84
	204,607	+30.00%	0	-100.00%	914,850	-49.92%	115
	236,085	+50.00%	0	-100.00%	680,780	-62.73%	124

objectives. The improvement to infeasible connections (-0.39%) indicates that for the same level of displacement (157,390 minutes) it is possible to find a slightly better solution using the PSAM-Pax approach. In terms of the Pareto frontier, the solution in row two dominates the solution produced by optimizing displacement alone. By considering passengers, it is possible to reduce passenger-level costs without increasing displacement. However, the improvements are minor, with only eight passengers regaining connections that were rendered infeasible when optimizing displacement alone. The model demonstrates more substantial benefits to the passenger objectives with further relaxations of displacement in the subsequent rows.

In rows 3–13 of Table 6.1, we relax the constraint on displacement and move along the Pareto frontier away from the anchor point identified by Step (1) of the solution approximation. Allowing a 1% increase in total displacement, we find a solution that decreases infeasible connections by approximately 52% and reduces the increase in connection time by approximately 16%. This is a recommended solution, which we will discuss in Section 6.3.

Since we are weighting Objective (2) over Objective (3), we observe an inflection

point between rows 4 and 5, where continued improvements to infeasible connections are only possible by increasing connection time. The solutions beyond row 5 sacrifice efficiency to maintain connections, as seen in Case (B) (Figure 5-3) of the demonstration of concept. For instance, if we increase total displacement by 10%, we achieve a nearly 98% reduction for Objective (2), but Objective (3) increases by approximately 12%. Airline acceptability of the schedule is an additional concern at higher levels of displacement. These results underscore the importance of identifying an appropriate balance between the multiple objectives.

Table 6.2: Multi-objective tradespace

	Infeasible Connections								
	2,065	1,859	1,652	**1,239**	**826**	413	0	CPU (min)	
Tot. Disp.	0%	-10%	-20%	**-40%**	**-60%**	-80%	-100%	Min	Max
157,390	1,826,660	inf	inf	inf	inf	inf	inf	36	36
+0%	+0%								
157,783	1,711,620	1,715,445	inf	inf	inf	inf	inf	38	39
+0.25%	-6.30%	-6.09%							
158,177	1,379,325	1,392,475	1,399,620	inf	inf	inf	inf	46	59
+0.5%	-24.49%	-23.77%	-23.38%						
158,570	1,305,880	1,314,015	1,335,550	**1,362,345**	inf	inf	inf	43	104
+0.75%	-28.51%	-28.06%	-28.69%	**-25.42%**					
158,963	1,211,475	1,219,560	1,235,340	**1,301,575**	inf	inf	inf	44	110
+1%	-33.68%	-33.24%	-32.37%	**-28.75%**					
160,537	1,082,340	1,088,990	1,111,370	**1,147,755**	**1,245,875**	1,753,740	inf	47	201
+2%	-40.75%	-40.38%	-39.16%	**-37.17%**	**-31.79%**	-3.99%			
163,685	911,785	920,095	940,245	977,515	1,057,750	1,262,945	inf	46	125
+4%	-50.08%	-49.63%	-48.53%	-46.49%	-42.09%	-30.86%			
166,833	774,920	778,150	785,540	862,315	947,960	1,093,670	inf	51	105
+6%	-57.58%	-57.40%	-57.00%	-52.79%	-48.10%	-40.13%			
169,981	714,525	720,095	726,380	759,110	855,575	968,610	inf	55	116
+8%	-60.88%	-60.58%	-60.23%	-58.44%	-53.16%	-46.97%			
173,129	626,830	663,605	663,610	713,585	791,990	890,020	2,051,855	57	338
+10%	-65.68%	-63.67%	-63.67%	-60.93%	-56.64%	-51.28%	+12.33%		
180,998	533,535	536,125	583,325	617,790	649,155	732,440	1,508,095	69	255
+15%	-70.79%	-70.65%	-68.07%	-66.18%	-64.46%	-59.90%	-17.44%		
204,607	262,425	295,455	282,565	339,325	385,120	461,975	914,850	84	453
+30%	-85.63%	-83.83%	-84.53%	-81.42%	-78.92%	-74.71%	-49.92%		
236,085	184,745	184,745	184,720	199,295	255,120	261,350	680,780	71	523
+50%	-89.89%	-89.89%	-89.89%	-89.09%	-86.03%	-85.69%	-62.73%		
Min CPU (min)	36	43	59	64	86	84	321		
Max CPU (min)	88	84	92	119	453	338	523		

Table 6.2 displays the minimized connection time increases when constraining displacement and infeasible connections to specified levels. To develop these results, we applied the ϵ-constrained method as described in Step (3) of Section 5.7. The results indicate that the relationship between displacement and the passenger objectives is

nonlinear. We are able to achieve significant improvements to infeasible connections and connection time when traded for relatively small increases in displacement. For example, with a 1% increase in displacement, we are able to reduce connection time by 33% while holding infeasible connections fixed. By adjusting the ϵ-constrained conditions, we find another solution at 1% displacement that provides a 40% reduction of infeasible connections and approximately 29% reduction in connection time. It is relatively easy to reduce the number of infeasible connections to 80% by increasing total displacement to 2%. However, making all connections feasible comes at the cost of an increase of at least 10% in total displacement and connection time increase of 13%. Recommended solutions are emphasized in bold and will be discussed in Section 6.3.

Figure 6-1 provides a three-dimensional tradespace to visualize the results provided by Table 6.2. While the utopia for displacement and connection time is at the bottom left, infeasible connections improve as we move toward the top right. As we improve any two of the objectives, the third becomes worse. The lines represent the Pareto frontier, with possible solutions represented as points. Recommended solutions are shown in the yellow shaded region and correspond to the bold solutions in Table 6.2.

The same trends observed in aggregate results are visible when examining individual itineraries. To observe the effects on individual itineraries, we analyze a subset of those with the highest passenger costs when displacement is optimized in isolation. Of the itineraries that have an increase in connection time greater than 100 minutes under the baseline PSAM solution, we select 15 that are expected to transport the most passengers over the season. In Table 6.3 we report the increase in connection time for this set of itineraries. The columns describe the scenarios as we increase displacement and infeasible connections relative to the baseline.

The itineraries originating at EWR involve a high-profile flight operated by Singapore Airlines known as the longest commercial flight in distance and time. Given the circumstances of this flight, the gravity of displacement decisions that indiscriminately eliminate or lengthen connections becomes clear. Direct flights from North

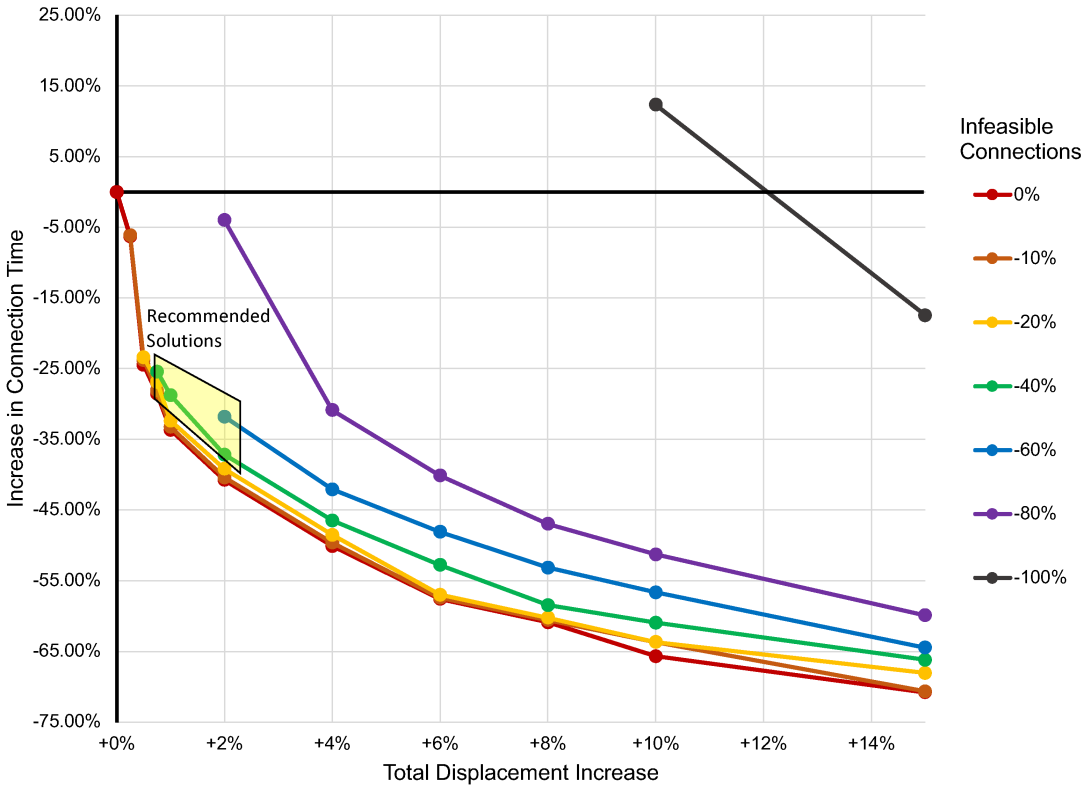

Figure 6-1: Three-dimensional Pareto frontier

America to many locations in Southeast Asia are unavailable, so maintaining feasible connections for key flights such as this becomes even more important.

The first column of results (displacement 0% and infeasible connections 0%) reflects the baseline PSAM solution obtained by Step (1) of the solution approximation, which focuses exclusively on minimizing total displacement. Under this solution, four important connecting itineraries become infeasible (i.e., EWR-DPS, MFM-PDG, IPH-SUB, and MFM-SUB), and the connections that remain feasible have relatively large increases to their connection time. Under the second scenario, displacement increases to 2%, and we find that only two connections are infeasible and the increases in connection time are substantially less. In Figure 6-2 we plot the first two scenarios for all itineraries and label those in the subset. We observe the subset itineraries in the top right of Figure 6-2a, with high values for Objective (3) and many passengers. When we relax displacement in Figure 6-2b, we observe improvements to Objective

Table 6.3: Results for high-impact itinerary subset

			Objective Value Scenarios						
	Displacement:		0%	2%	4%	6%	8%	10%	10%
	Infeasible Conn.:		0%	0%	0%	0%	0%	0%	20%
Origin	Dest	Pax	Increase in Conn. Time (min)						
EWR	CGK	738	175	0	0	0	0	0	0
EWR	DPS	626	inf	0	0	0	0	0	0
EWR	BKK	609	125	0	0	0	0	0	0
PEN	SUB	477	105	45	inf	inf	inf	inf	0
EWR	BKK	438	80	0	0	0	0	0	0
SGN	SUB	385	100	40	40	40	0	0	0
MFM	PDG	147	inf	105	105	45	105	0	105
HKT	SUB	140	105	45	inf	45	0	inf	0
SEA	CGK	138	190	inf	inf	inf	inf	inf	190
IPH	SUB	137	inf	45	45	45	0	0	0
MFM	SUB	137	inf	45	45	45	0	0	0
TAO	SUB	118	105	45	45	45	0	0	0
SEA	BKK	114	130	130	130	130	130	95	130
MNL	PDG	113	105	inf	inf	inf	inf	inf	105
TSN	SUB	102	105	45	45	45	0	0	0

(3) for all itineraries, including some that were infeasible in the baseline solution. In general, this trend continues in the subsequent columns of Table 6.3 as displacement is relaxed. In the final scenarios, we reduce infeasible connections by 20% while maintaining displacement at 10%, which causes the connection time increase for some itineraries to return to their baseline values.

Note that for this set of itineraries, the metrics at 4% displacement are worse than at 2% displacement. While this may seem counter-intuitive, this sample does not fully capture the trade-offs that the optimizer is making. At 4% displacement, the model found it more valuable to improve other itineraries at the expense of some of those included in the sample.

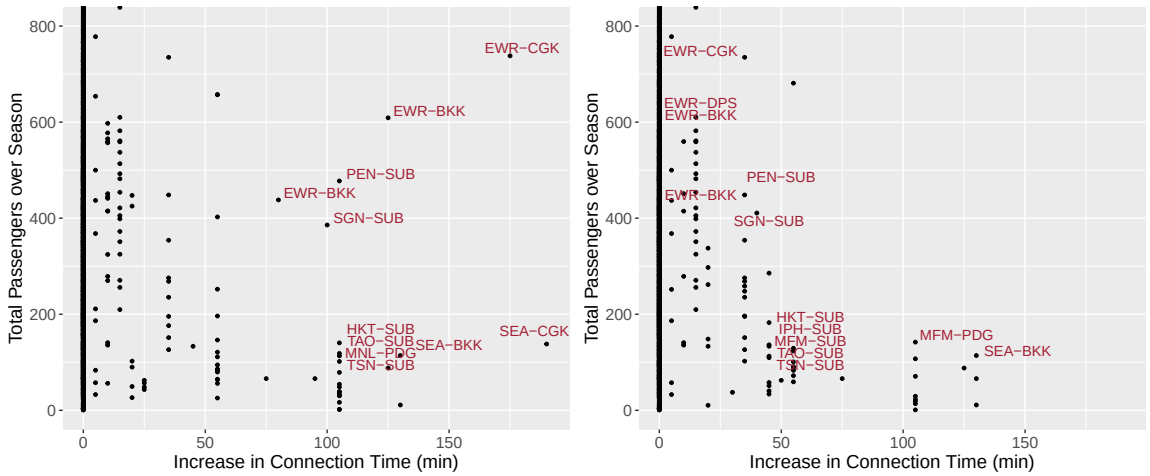

(a) 0% displacement, 0% infeasible connections (b) 2% displacement, 0% infeasible connections

Figure 6-2: PSAM versus PSAM-Pax solutions for high-volume itinerary subset

6.3 Recommended Solutions

While it is impractical to suggest a single solution, we emphasize a range of solutions in Table 6.1, Table 6.2, and Figure 6-1. Note that the modeling parameters that produced these solutions are specific to Changi Airport. As a large hub airport with a high percentage of connections and relatively high schedule displacement, Changi has the potential to trade small percentages of displacement to decrease costs and enable better itineraries for a large number of connecting passengers. The recommended solutions have a total displacement increase of 0.75-2% relative to the baseline (i.e., PSAM) solution. Given the small percentage increase in displacement, these solutions reduce infeasible passenger connections by 40-60% and cut connection time increases by 16-37%.

The recommended solutions were selected for several reasons. First, the solutions are Pareto optimal, meaning that one of the objectives cannot be improved without worsening one or both of the other objectives. Second, we believe a 2% or lower optimality gap for displacement would be acceptable to the airlines. Third, the recommended solutions lie in a region of the tradespace where the relationships between displacement and the passenger-centric metrics are highly nonlinear, as shown in Figure 6-1. This means that we achieve the largest reductions in passenger costs for

small unit increases to displacement.

6.4 Robustness Checks

To evaluate the impact of important model parameters on PSAM-Pax performance, we conduct sensitivity analysis of the calibration parameter α. We use the α parameter in a weighted-sum approach during Step (3) of the solution approximation to balance two terms representing the nonlinear objective (i.e., Objective (3), increase in passenger connection time). Table 6.4 displays the percent decrease in Objective (3) with respect to the baseline PSAM solution under different values of α. The largest improvements to Objective (3) for each level of displacement are shown in bold. We observe that the solutions are not very sensitive to α for values between 0.3 and 0.6, but solution quality decreases as α approaches the extremes (i.e., 0 and 1). The best solutions are achieved with an intermediate α value, as it provides a balance between the two approximated terms.

Table 6.4: Sensitivity of the calibration parameter α

	α value										
Tot. Disp.	0.001	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
+0.25%	**-6%**	**-6%**	**-6%**	**-6%**	**-6%**	**-6%**	**-6%**	-4%	-2%	-3%	-3%
+0.5%	-7%	-10%	-10%	**-24%**	**-24%**	**-24%**	**-24%**	**-24%**	**-24%**	**-24%**	-23%
+0.75%	-11%	-12%	-28%	-29%	**-30%**	**-30%**	**-30%**	**-30%**	-27%	-27%	-25%
+1%	-12%	**-34%**	**-34%**	**-34%**	**-34%**	**-34%**	**-34%**	-30%	-27%	-26%	-24%
+2%	-21%	-22%	-41%	-41%	**-42%**	-41%	**-42%**	-39%	-36%	-22%	-20%
+4%	-31%	-32%	-50%	-50%	-51%	-51%	**-52%**	-48%	-43%	-34%	-19%
+6%	-37%	-38%	**-58%**	**-58%**	**-58%**	**-58%**	**-58%**	-57%	-47%	-43%	-5%
+8%	-41%	-61%	-61%	-61%	**-62%**	-59%	-59%	-59%	-52%	-46%	-1%
+10%	-57%	-58%	-64%	**-66%**	-65%	-62%	-62%	-64%	-57%	-55%	32%
+15%	-52%	-56%	-71%	-71%	**-72%**	-71%	-69%	-68%	-62%	-61%	9%
+30%	-67%	-61%	-60%	**-86%**	-85%	-83%	-78%	-76%	-68%	-72%	7%
+50%	-72%	-82%	-90%	-90%	-90%	-89%	-88%	-85%	-79%	-82%	1%
Min CPU (min)	40	51	32	36	43	30	39	42	39	45	71
Max CPU (min)	116	101	105	88	104	103	95	78	82	110	298

As described in Section 5.3.2, there are several logical distributions we can use to disaggregate the passenger flow predictions. To determine if the results are robust to uncertainty in this process, we evaluate an array of solutions under the alternative dis-

tributions and display our findings in Table 6.5. The table shows the objective values for five levels of increasing displacement and two levels of infeasible passenger connections. First, we report the results for the Seat Capacity distribution, which was used to produce the results reported in this chapter. While keeping the solution fixed, we recalculate the objective values under the alternative distributions: Uniform, Time-of-Day A, and Time-of-Day B. The results are shown to have insignificant variation between the four distribution methods.

Table 6.5: Sensitivity of passenger distribution method

	Seat Capacity		Uniform		Time-of-Day A		Time-of-Day B	
Tot. Disp.	Inf. Con. (pax)	Inc. Con. (pax-min)	Inf. Con. (pax)	Inc. Con. (pax-min)	Inf. Con. (pax)	Inc. Con. (pax-min)	Inf. Con. (pax)	Inc. Con. (pax-min)
0%	2,065	182,660	1,805	1,865,015	1,800	1,824,510	1,806	1,937,575
		1,082,340		1,043,925		1,066,285		1,080,685
2%		-40.75%		-44.03%		-40.66%		-44.24%
		911,785		875,985		895,815		911,665
4%		-50.08%		-53.03%		-49.80%		-52.95%
	2,065	774,920	1,805	725,340	1,800	748,105	1,806	748,405
6%	0%	-57.58%	0%	-61.11%	0%	-57.72%	0%	-61.37%
		714,525		654,495		693,160		671,360
8%		-60.88%		-64.91%		-60.66%		-65.35%
		626,830		628,880		684,975		629,160
10%		-65.68%		-66.28%		-61.10%		-67.53%
		1,111,370		1,104,545		1,129,810		1,140,785
2%		-39.16%		-40.77%		-38.08%		-41.12%
		940,245		936,610		959,340		971,765
4%		-48.53%		-49.78%		-47.42%		-49.85%
	1,652	785,540	1,444	778,975	1,439	805,915	1,445	801,875
6%	-20%	-57.00%	-20%	-58.27%	-20%	-55.85%	-20%	-58.61%
		726,380		709,925		752,950		726,695
8%		-60.23%		-61.95%		-58.73%		-62.49%
		663,610		647,875		683,630		660,430
10%		-63.67%		-65.26%		-62.53%		-65.91%

The consistency in the results demonstrates that the improvements provided by PSAM-Pax are not sensitive to the differences between the distributions. This is a positive indicator of the robustness of the optimization model.

Chapter 7

Conclusion

7.1 Summary of the Results

Airport congestion is an escalating global problem that manifests itself in delays that generate costs for airlines and passengers on the order of billions of dollars annually. Unable to expand capacity due to resource constraints, many of the most congested airports globally turn to slot allocation. Slot allocation reduces delay by leveling the demand profile, recognizing that delay increases nonlinearly as airport utilization approaches saturation [32]. IATA has formalized slot guidelines (i.e., WASG) for schedule-coordinated airports, but the execution of initial slot allocation is often oversimplified and can have significant implications for airlines and passengers.

As a highly complex scheduling problem, slot allocation is well suited for integer programming. While previous models have successfully captured the constraints associated with slot allocation and produced solutions for the largest schedule-coordinated airports, they commonly focus on flight-centric objectives. By minimizing displacement, previous models can produce schedules that align closely with the airline requests but do not entirely account for the costs of slot allocation. Focusing exclusively on displacement will indiscriminately increase passenger connection times or eliminate connections altogether, which increases costs for passengers.

The results of our research indicate that the relationship between schedule displacement and passenger costs are nonlinear. A limited increase in displacement

from the point of single-objective optimality can yield significant improvements to passenger-level metrics, such as infeasible connections and connection time. In this case, a slot optimization model that incorporates passenger considerations can enable better itineraries—reducing costs for passengers and improving airline revenues—especially at large hub airports.

While incorporating passenger-level metrics requires forecasting, our results indicate that passenger flows can be accurately predicted. Historical itinerary and socioeconomic data is available to train machine learning models, which can leverage large datasets with numerous features, and are flexible to changes in the transportation network. The most important predictors for existing itineraries are historical passenger counts at the itinerary and market levels, route efficiency, distance, GDP of the airport catchment areas, and itinerary type (nonstop versus connecting). Itinerary type becomes the most important predictor for new itineraries, as historical passenger counts are unavailable.

Of the prediction methods tested, random forests provide the best performance. Applying cross-validation for feature selection and parameter tuning, we found that random forests performed best with a subset of 17 features, but performance remained high with as few as ten features. When tested on new data, random forests achieved a mean absolute error (MAE) of 121 passengers per month, suggesting excellent predictive performance. The performance of random forests dropped only slightly to an MAE of 129 passengers per month on new itineraries, providing an informative forecast even when historical passenger counts are unavailable. k-NN performed nearly as well as random forests, but the linear models performed poorly on new itineraries. In contrast to linear models, the parametric models provide flexibility to leverage linear and nonlinear relationships with continuous features available for new itineraries.

Integrating the prediction and optimization models requires constructing all possible itineraries, identifying reference itineraries to measure the costs of reaccomodating passengers on alternatives, and disaggregating the forecasts. After experimenting with four methods of distributing aggregate passengers over specific itineraries, we found that the optimization results are robust to this procedure.

The experimental results demonstrated that focusing exclusively on flight-level metrics when conducting slot allocation will generate sub-optimal solutions from the standpoint of passengers. Incorporating passenger-level objectives illuminates opportunities for trade-offs that would otherwise be overlooked. Analysis of the 2019 summer season at Singapore Changi Airport demonstrated that a slot optimizer focused exclusively on total displacement would have eliminated 2,065 passenger connections and increased connection time by 1,826,660 passenger minutes compared to the unconstrained assignment. Minimizing displacement alone resulted in a Pareto-dominated solution that could be improved from the passengers' perspective without increasing displacement. By including objectives to minimize infeasible passenger connections and increasing passenger connection time, we found that schedule displacement exhibits a nonlinear relationship with the passenger-level objectives. This enabled us to trade small increases in displacement for considerable improvements to the passenger-level metrics. Applying the PSAM-Pax approach, we were able to decrease the number of infeasible passenger connections by approximately 52% and reduce passenger connection time increases by approximately 16% with only a 1% increase in total displacement, compared to the baseline solution offered by PSAM.

The interdependence of displacement and passenger decisions created a nonlinear term in the PSAM-Pax objective function that prevented direct implementation of the model. However, we were able to produce solutions by approximating the objective function with a coordinate descent procedure. The solution approximation produces solutions along the Pareto frontier for a large airport in less than five hours of computation time. We used a weighted sum approach to approximate the nonlinear term and found that the best solutions are achieved when the calibration parameter is at intermediate values.

7.2 Contributions of the Thesis

The major contributions of this thesis are:

- Using a predict-then-optimize framework, we propose an original approach to

airport slot allocation that incorporates passenger considerations.

- Machine learning methods are applied to predict passenger flows across a network of flights. We achieve excellent predictive performance using random forests, even on new itineraries that did not exist in the training observations.

- To measure passenger costs associated with slot allocation, months before the day of operations, we demonstrate procedures for distributing passenger forecasts across potential itineraries constructed from slot requests.

- The proposed multi-objective optimization model allocates slots according to airport capacity constraints and the WASG while minimizing one flight-centric metric—schedule displacement—and two passenger-centric metrics—infeasible connections and connection time.

- A solution approach is formulated to efficiently solve the proposed optimization model for the largest schedule-coordinated airports.

- The proposed optimization model, referred to as PSAM-Pax, is implemented using real-world data from Singapore Changi Airport to create slot allocation solutions that achieve Pareto optimality in acceptable computation times. The results indicate that PSAM-Pax can provide significant improvements to slot allocation by trading small increases in schedule displacement for considerable improvements to passenger-level metrics.

7.3 Further Research

As are most real-world problems, airport slot allocation is multi-objective, which leads to many acceptable solutions arrayed along the Pareto frontier. While the PSAM-Pax model quantifies the trade-offs between passenger-level and flight-level objectives, the proper balance remains an open research question. Future studies could apply PSAM-Pax to investigate and recommended proper weighting between objectives based on specific airport features. Relevant features may include annual

traffic, delay, percentage of connecting itineraries, schedule displacement, and passenger characteristics. For a large transfer airport with high displacement, such as Singapore Changi, passenger-level metrics probably deserve greater weight. In contrast, a non-transfer airport or one with low displacement may allocate greater weight towards minimizing schedule displacement. By identifying solutions along the Pareto frontier, PSAM-Pax provides a means to facilitate conversations with stakeholders, but additional research is required to identity the "best" solution given specific airport characteristics.

For the application of PSAM-Pax in this thesis, we reduce computational expense by considering a single declared capacity for the airport. The solution approach should be strengthened to solve PSAM-Pax with multiple capacity constraints, such as those for various timescales and types of operation. Heuristic search techniques, which use a combination of randomness and rules to guide the search for global optima, have demonstrated success in related studies. The large-scale neighborhood search method developed in [36] provides a basis for this solution approach.

With improved algorithms and advanced computing power, it becomes possible to expand the modeling framework to account for multi-airport systems. Since the PSAM-Pax is limited to a single airport, it neglects network effects. Future work could analyze multiple airports to understand interactions and work to optimize groups of schedule-coordinated airports simultaneously.

Another direction for future research is to improve the treatment of passenger demand. The modeling framework used in this thesis does not account for supply-demand interactions. Rather than treating passenger flows as an elastic function of demand, PSAM-Pax accepts the forecast as a fixed parameter. In reality, passenger demand influences airline slot requests, and displacement decisions influence demand. This thesis demonstrated how incorporating passenger considerations into slot allocation can enable improved itineraries, so it is reasonable that the quality of displacement decisions will affect the overall demand level and how the demand is distributed across itineraries.

One option to incorporate supply-demand interactions is through a feedback loop

in the modeling framework. This would involve expanding the prediction model to include supply-related variables, which are output by the optimization model. The prediction model uses these variables to evaluate the quality of itineraries and updates the forecasts accordingly. The feedback loop enables the models to iterate between supply optimization and demand estimation until convergence. Decisions regarding the number of iterations and convergence criteria are challenges of this approach. A second option is to incorporate a parametric function into the optimization formulation mathematically. Since the functional form of a realistic demand model is nonlinear, this approach will increase the computational complexity.

PSAM-Pax is subject to the uncertainty of a forecast. The pandemic resulting from COVID-19 reinforced the lesson that "the forecast is always wrong" [11]. To avoid overconfidence in forecasts, we must convey confidence intervals alongside predictions. We then use the confidence intervals to conduct robust optimization, which will produce a slot allocation that is not sensitive to small changes in the forecasts. This will not prevent errors when it comes to black swan events like COVID-19, but it is a step in the right direction. Additionally, future travel patterns may look very different than those represented by the data used in this thesis. The usefulness of predictive features varies as the transportation network evolves, which is currently occurring at an unprecedented rate. As one example, the binary indicator for LCC may no longer be valuable as airlines merge elements of low-cost and full-service business models, as was the trend even before COVID-19. Future studies need to incorporate current data and explore new features as airlines and airports redefine their identities in the post-pandemic world.

Finally, additional research should include a comparison of slot allocation under PSAM-Pax with other demand management approaches, such as congestion pricing, slot auctions, and laissez-faire. To gain support for implementing a new paradigm for slot allocation, we must prove that it provides clear benefits over alternative methods.

Appendix A

Complete Set of Predictive Features

Table A.1: Complete set of predictive features

Feature	Type	Description
Historical Traffic		
Lagging Itinerary Passengers	Num	Historical passenger count for the itinerary; 12-month lagging dependent variable
Lagging Market Passengers	Num	Historical passenger count for the same city-pair (e.g., London-Singapore) and month in the prior year, irrespective of airline and airport
Annual Airport Traffic	Num	Geometric mean of the annual traffic at the origin and destination airports in the previous year
Itinerary Passengers Growth	Num	Percent growth of itinerary passengers over the two years prior to prediction (i.e., final two years of the training dataset)
Seasonality	Num	Percent deviation from the mean for the airport pair (o, d) in month m
New Itinerary	Bin	1 if the itinerary (same set of airports, airlines, and month) did not exist the year prior; 0 if the itinerary has continuity
Socioeconomic		
Population	Num	Geometric mean of the population in a 100km circular buffer around o and d
GDP	Num	Geometric mean of the GDP in a 100km circular buffer around o and d
GDP per Capita	Num	Geometric mean of the GDP per capita in a 100km circular buffer around o and d
GDP Difference	Num	Percent difference between the GDP at o and d
Itinerary Characteristics		
Connecting	Bin	1 if the itinerary is connecting; 0 if nonstop
Distance	Num	Flying distance $o \rightarrow h \rightarrow d$ (km)
Routing Factor	Num	Measure of efficiency for connecting itineraries; flying distance divided by nonstop distance
Airlines & Airports		
Alliance Membership	Cat	Alliance membership of the primary airline; Star Alliance, SkyTeam, Oneworld, and Other
Low Cost Carrier (LCC)	Bin	1 if the primary airline is considered a LCC according to OAG; 0 otherwise
Most Frequent Airlines (SQ, TR, MI)	Bin	Individual dummy variables for airlines with the most frequent operations at Singapore airport (Singapore Airlines, Scoot, and SilkAir)
Airport Continents	Cat	Categorical representation of airport locations (e.g., Europe-Asia for a London-Singapore itinerary)

Bibliography

[1] International Air Transport Association. *Worldwide Airport Slot Guidelines*. International Air Transport Association, Montreal, Quebec, 2020.

[2] Bilge Atasoy, Matteo Salani, and Michel Bierlaire. An Integrated Airline Scheduling, Fleeting, and Pricing Model for a Monopolized Market: An integrated airline scheduling, fleeting, and pricing model. *Computer-Aided Civil and Infrastructure Engineering*, 29(2):76–90, February 2014.

[3] M. Ball, C. Barnhart, M. Dresner, M. Hansen, K. Neels, A. Odoni, E. Peterson, L. Sherry, A. Trani, and Bo Zou. Total delay impact study: A comprehensive assessment of the costs and impacts of flight delay in the united states. 2010.

[4] Cynthia Barnhart, Douglas Fearing, and Vikrant Vaze. Modeling Passenger Travel and Delays in the National Air Transportation System. *Operations Research*, 62(3):580–601, June 2014.

[5] Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

[6] Sebastian Birolini, António Pais Antunes, Mattia Cattaneo, Paolo Malighetti, and Stefano Paleari. Integrated flight scheduling and fleet assignment with improved supply-demand interactions. *Transportation Research Part B: Methodological*, 149:162–180, July 2021.

[7] Sebastian Birolini, Mattia Cattaneo, Paolo Malighetti, and Chiara Morlotti. Integrated origin-based demand modeling for air transportation. *Transportation Research Part E: Logistics and Transportation Review*, 142:102050, October 2020.

[8] Stephane Bratu and Cynthia Barnhart. Flight operations recovery: New approaches considering passenger recovery. *Journal of Scheduling*, 9(3):279–298, June 2006.

[9] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, October 2001.

[10] Edward Crawley, Bruce Cameron, and Daniel Selva. *System Architecture: Strategy and Product Development for Complex Systems*. Pearson, Hoboken, New Jersey, 2016.

[11] Richard De Neufville and Stefan Scholtes. *Flexibility in Engineering Design*, chapter 2, pages 15–38. The MIT Press, Cambridge, Massachusetts, 2011.

[12] Zhang Dong, Yu Chuhang, and H.Y.K. Henry Lau. An integrated flight scheduling and fleet assignment method based on a discrete choice model. *Computers & Industrial Engineering*, 98:195–210, August 2016.

[13] EUROCONTROL. Coda digest annual report, 2014.

[14] EUROCONTROL. Coda digest annual report, 2015.

[15] EUROCONTROL. Coda digest annual report, 2016.

[16] EUROCONTROL. Coda digest annual report, 2017.

[17] EUROCONTROL. Coda digest annual report, 2018.

[18] EUROCONTROL. Coda digest annual report, 2019.

[19] Jamie Fairbrother and Konstantinos G. Zografos. Optimal scheduling of slots with season segmentation. *European Journal of Operational Research*, 291(3):961–982, June 2021.

[20] Arthur E. Hoerl and Robert W. Kennard. Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*, 42(1):80–86, 1970.

[21] IATA. Worldwide airport slots: Fact sheets, November 2020.

[22] Alexandre Jacquillat. Predictive and Prescriptive Analytics toward Passenger-centric Ground Delay Programs. *SSRN Electronic Journal*, 2020.

[23] Alexandre Jacquillat and Vikrant Vaze. Interairline Equity in Airport Scheduling Interventions. *Transportation Science*, 52(4):941–964, August 2018.

[24] Yu Jiang and Konstantinos G. Zografos. A decision making framework for incorporating fairness in allocating slots at capacity-constrained airports. *Transportation Research Part C: Emerging Technologies*, 126:103039, May 2021.

[25] Fotios A. Katsigiannis and Konstantinos G. Zografos. Optimising airport slot allocation considering flight-scheduling flexibility and total airport capacity constraints. *Transportation Research Part B: Methodological*, 146:50–87, April 2021.

[26] Fotios A. Katsigiannis, Konstantinos G. Zografos, and Jamie Fairbrother. Modelling and solving the airport slot-scheduling problem with multi-objective, multi-level considerations. *Transportation Research Part C: Emerging Technologies*, 124:102914, March 2021.

[27] Alix Lhéritier, Michael Bocamazo, Thierry Delahaye, and Rodrigo Acuna-Agost. Airline itinerary choice modeling using machine learning. *Journal of Choice Modelling*, 31:198–209, June 2019.

[28] Scott Lundberg and Su-In Lee. A unified approach to interpreting model predictions. December 2017.

[29] Lavanya Marla, Bo Vaaben, and Cynthia Barnhart. Integrated Disruption Management and Flight Planning to Trade Off Delays and Fuel Burn. *Transportation Science*, 51(1):88–111, February 2017.

[30] Daniel L. McFadden. Conditional Logit Analysis of Qualitative Choice Behavior. In Paul Zarembka, editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York, 1973.

[31] Thomas Morisset and Amedeo Odoni. Capacity, delay, and schedule reliability at major airports in europe and the united states. *Transportation Research Record*, 2214(1):85–93, 2011.

[32] Richard L. de Neufville, Amedeo R. Odoni, Peter Belobaba, and Tom G. Reynolds. *Airport Systems, Second Edition: Planning, Design and Management*. McGraw Hill Professional, March 2013.

[33] Stefano Paleari, Renato Redondi, and Paolo Malighetti. A comparative study of airport connectivity in china, europe and us: Which network provides the best service to passengers? *Transportation Research Part E: Logistics and Transportation Review*, 2:198–210, 03 2010.

[34] João P. Pita, Nicole Adler, and António P. Antunes. Socially-oriented flight scheduling and fleet assignment model with an application to Norway. *Transportation Research Part B: Methodological*, 61:17–32, March 2014.

[35] Jasenka Rapajic. *Beyond Airline Disruptions*. Routledge, New York, New York, second edition, 2019.

[36] Nuno Antunes Ribeiro, Alexandre Jacquillat, and António Pais Antunes. A Large-Scale Neighborhood Search Approach to Airport Slot Allocation. *Transportation Science*, 53(6):1772–1797, November 2019.

[37] Nuno Antunes Ribeiro, Alexandre Jacquillat, António Pais Antunes, Amedeo R. Odoni, and João P. Pita. An optimization approach for airport slot allocation under IATA guidelines. *Transportation Research Part B: Methodological*, 112:132–156, June 2018.

[38] Galit Shmueli, Peter C. Bruce, Inbal Yahav, Nitin R. Patel, and Kenneth C. Lichtendahl Jr. *Data Mining for Business Analytics: Concepts, Techniques, and Applications in R*. John Wiley and Sons, Hoboken, New Jersey, 2017.

[39] Robert Tibshirani. Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

[40] Keji Wei, Vikrant Vaze, and Alexandre Jacquillat. Airline Timetable Development and Fleet Assignment Incorporating Passenger Choice. *Transportation Science*, pages 139–163, December 2020.

[41] Chiwei Yan, Vikrant Vaze, Allison Vanderboll, and Cynthia Barnhart. Tarmac delay policies: A passenger-centric analysis. *Transportation Research Part A: Policy and Practice*, 83:42–62, January 2016.

[42] Konstantinos G. Zografos, Konstantinos N. Androutsopoulos, and Michael A. Madas. Minding the gap: Optimizing airport schedule displacement and acceptability. *Transportation Research Part A: Policy and Practice*, 114:203–221, August 2018.

[43] Konstantinos G. Zografos and Yu Jiang. A Bi-objective efficiency-fairness model for scheduling slots at congested airports. *Transportation Research Part C: Emerging Technologies*, 102:336–350, May 2019.

[44] Konstantinos G. Zografos, Michael A. Madas, and Konstantinos N. Androutsopoulos. Increasing airport capacity utilisation through optimum slot scheduling: review of current developments and identification of future needs. *Journal of Scheduling*, 20(1):3–24, February 2017.

[45] Konstantinos G. Zografos, Yiannis Salouras, and Michael A. Madas. Dealing with the efficient allocation of scarce resources at congested airports. *Transportation Research Part C: Emerging Technologies*, 21(1):244–256, April 2012.