Transportation Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Passenger-Centric Slot Allocation at Schedule-Coordinated Airports

Sebastian Birolini, Alexandre Jacquillat, Phillip Schmedeman, Nuno Ribeiro

To cite this article:
Sebastian Birolini, Alexandre Jacquillat, Phillip Schmedeman, Nuno Ribeiro (2022) Passenger-Centric Slot Allocation at Schedule-Coordinated Airports. Transportation Science

Published online in Articles in Advance 17 Aug 2022

. https://doi.org/10.1287/trsc.2022.1165

http://pubsonline.informs.org/journal/trsc

TRANSPORTATION SCIENCE
Articles in Advance, pp. 1–23
ISSN 0041-1655 (print), ISSN 1526-5447 (online)

Passenger-Centric Slot Allocation at Schedule-Coordinated Airports

Sebastian Birolini,[a] Alexandre Jacquillat,[b,*] Phillip Schmedeman,[c] Nuno Ribeiro[d]

[a] Department of Management, Information and Production Engineering, University of Bergamo, 24123 Bergamo, Italy; [b] Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142; [c] Department of Systems Engineering, United States Military Academy, West Point, New York 10996; [d] Aviation Studies Institute, Engineering Systems and Design, Singapore University of Technology and Design, Singapore 487372
*Corresponding author

Contact: sebastian.birolini@unibg.it, https://orcid.org/0000-0002-9397-7165 (SB); alexjacq@mit.edu, https://orcid.org/0000-0002-2352-7839 (AJ); philds@mit.edu, https://orcid.org/0000-0002-2463-3879 (PS); nuno_ribeiro@sutd.edu.sg (NR)

Received: July 26, 2021
Revised: March 26, 2022; June 25, 2022
Accepted: July 5, 2022
Published Online in Articles in Advance: August 17, 2022

https://doi.org/10.1287/trsc.2022.1165

Abstract. Schedule coordination is the primary form of demand management used at busy airports. At its core, slot allocation involves a highly complex combinatorial problem. In response, optimization models have been developed to minimize the displacement of flight schedules from airline requests, subject to physical and administrative constraints. Existing approaches, however, may not result in the best itineraries for passengers. This paper proposes an original passenger-centric approach to airport slot allocation to maximize available itineraries and minimize connecting times. Because of the uncertainty regarding passenger demand, the proposed approach combines predictive analytics to forecast passenger flows in flight networks from historical data and prescriptive analytics to optimize airport slot assignments in view of flight-centric and passenger-centric considerations. The problem is formulated as a mixed-integer nonconvex optimization model. To solve it, we propose an approximation scheme that alternates between flight-scheduling and passenger-accommodation modules and embed it into a large-scale neighborhood search algorithm. Using real-world data from the Singapore Changi and Lisbon Airports, we show that the proposed model and algorithm return solutions in acceptable computational times. Results suggest that slot-allocation outcomes can be made much more consistent with passenger flows at a relatively small cost in terms of flight displacement. Ultimately, this paper provides a new paradigm that can create more attractive flight schedules by bringing together airport-level considerations, airline-level considerations, and, for the first time, passenger-level considerations.

Funding: Financial support from the Civil Aviation Authority of Singapore [Project on Airfield Management and Economics] is gratefully acknowledged.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/trsc.2022.1165.

Keywords: slot allocation • multiobjective optimization • integer programming • analytics

1. Introduction

Following the era of uncertainty caused by COVID-19, the air transportation industry is slowly ramping operations back to prepandemic levels. Yet, air traffic management systems face strict capacity constraints, hindering flight operations globally and making the efficient use of scarce infrastructure resources even more critical. One of the main levers to manage demand-capacity imbalances in air transportation is known as schedule coordination, which coordinates the schedules of flights to prevent overcapacity scheduling and thereby manage congestion at busy airports.

Schedule coordination, however, involves hard combinatorial problems to create attractive schedules in interconnected networks of flights. Extensive research has developed dedicated optimization models, with the primary objective of scheduling flights as closely as possible to airlines' requested times. However, the "best" schedules of flights may not result in optimal outcomes from the perspective of passengers, by potentially restricting the number of available itineraries or making connections less attractive. Failing to incorporate passenger-level considerations can therefore increase the costs of schedule coordination for passengers, airlines, and other air transportation stakeholders.

In response, this paper develops an original passenger-centric approach to airport demand management, combining predictive and prescriptive analytics. Before detailing our contributions, we first provide some necessary background on demand management and passenger-centric operations.

1.1. Airport Slot Allocation

Demand-management interventions fall into two main categories. Economic approaches based on congestion pricing (Carlin and Park 1970, Daniel 1995, Brueckner 2002) and slot auctions (Rassenti et al. 1982, Ball et al. 2006, 2020) can achieve efficient outcomes, at least in theory. However, busy airports—especially outside the United States—rely on nonmonetary schedule coordination. As of 2019, schedule coordination was applied at 204 airports, representing 43% of global traffic.

The core component of schedule coordination is known as slot allocation, which assigns departure and arrival times to each flight over a season. Namely, airlines submit requests for series of slots, each comprising at least five flights on the same day of the week and specifying desired scheduled times. Whenever demand exceeds the airport's declared capacity, some slots need to be displaced to earlier or later times, following the Worldwide Airport Slot Guidelines (WASG) from the International Air Transport Association (IATA). These rules include, in particular, (i) connectivity requirements to maintain minimum turnaround times, (ii) regularity requirements to ensure consistency within each slot request, and (iii) grandfather requirements to abide by historical precedence (IATA 2020).

Slot allocation is a central problem in air transportation, but also a highly complex combinatorial problem—due, in part, to the interdependencies across all slot requests throughout the season. Zografos et al. (2012) proposed the first optimization model that minimizes schedule displacement (defined as the difference between flights' requested times and assigned slots), subject to connectivity and regularity constraints. The model solves in five minutes at a regional airport and achieves improvements of 14%–95% over the historical schedule. Ribeiro et al. (2018) developed the Priority-based Slot Allocation Model (PSAM), which captures the WASG grandfather requirements, and trades off several measures of displacement. Thanks to a strengthened formulation with valid inequalities, PSAM could be implemented at medium-sized airports. Nonetheless, slot allocation remains too complex to be solved with off-the-shelf methods at the largest schedule-coordinated airports. Therefore, Ribeiro et al. (2019) developed a large-scale neighborhood search algorithm to solve PSAM at the largest airports worldwide. This algorithm found the optimal solution in 10 hours at the Lisbon Airport, which operates more than 200,000 aircraft movements per year.

This research demonstrated that advanced optimization methods can match airlines' scheduling requests more closely than historical schedules. Additionally, optimization methods provide flexibility by augmenting primary objectives (e.g., schedule displacement) with secondary objectives, such as market preferences (Jorge et al. 2021), interairline fairness (Jacquillat and Vaze 2018, Zografos and Jiang 2019, Jiang and Zografos 2021, Katsigiannis et al. 2021), airline acceptability (Zografos et al. 2017, 2018), airline flexibility preferences (Katsigiannis and Zografos 2021), and scheduling flexibility (Fairbrother and Zografos 2021). The main principle guiding these formulations is that all flight displacements are not equally costly, and capturing such heterogeneity can yield more attractive slot-allocation outcomes.

Yet, these objectives focus on flight-centric and airline-centric considerations, but do not capture passenger itineraries. Existing approaches may thus result in (i) fewer available connecting itineraries; and (ii) longer connecting times and, hence, less attractive itineraries. In response, this paper develops a passenger-centric approach to make slot-allocation outcomes more consistent with passenger demand. A major challenge, however, is that passenger itineraries are unknown at the time of slot allocation—months before operations. Therefore, we use historical data and machine learning to *predict* passenger flows in flight networks, which we then leverage into a new multiobjective formulation to *optimize* slot allocation, trading off flight-centric and passenger-centric goals.

1.2. Demand Modeling and Passenger-Centric Air Transportation

On the predictive side, the empirical literature for estimating air-travel demand falls into three broad categories. The first one involves discrete choice models, using disaggregate passenger-level data (Cadarso et al. 2017, Lurkin et al. 2017, Lhéritier et al. 2019) or semiaggregate itinerary-level data (Birolini et al. 2021b). These models use schedule-related information (e.g., scheduled times, connecting times, airfares, aircraft types, and competitors) that comes from subsequent airline decisions and is thus not available at the time of slot allocation. The second category includes time-series models, using aggregated longitudinal data (e.g., at the airport or regional level over several years) (Marazzo et al. 2010, Tsui et al. 2014, Hakim and Merkert 2016, Phyoe et al. 2016, Jin et al. 2020). These models are mainly used to investigate causal relationships between air transport and economic growth, tourism, and trade. Moreover, these approaches are limited when it comes to predicting demand on new itineraries—a critical component for slot allocation. Our paper falls into the third category—aggregated gravity models—but exhibits three main novelties. First, gravity models involve service-related factors, which are not available at the time of slot allocation, which require new features in our model. Second, gravity models usually estimate market demand at the airport-pair or city-pair level (Adler and Hashai 2005, Grosche et al. 2007, Adler et al. 2018, Boonekamp et al. 2018) or at the airport-pair-airline level (Wei and Hansen 2006, Birolini et al. 2021a). To support more disaggregated slot-allocation decisions, however, we need to estimate demand at the

passenger-route level, defined as an airport route between city pairs and corresponding airlines. Viewed through this lens, our approach trades off the high level of detail required for operational decision making and the limited information availability at the strategic level, establishing a middle ground between discrete choice and gravity models. Finally, gravity models are generally estimated by using linear regression—ordinary least squares and two-stage least squares to cope with endogeneity issues. In contrast, our approach leverages machine-learning methods to capture nonlinearities.

On the optimization side, this paper builds upon adjacent literature that demonstrates the benefits of incorporating passenger considerations in air transportation. Early airline fleet assignment and scheduling models incorporated itinerary demand as a fixed input and leveraged split-and-recapture models to distribute demand among competing itineraries (Lohatepanont and Barnhart 2004, Sherali et al. 2010, Pita et al. 2013). More recent studies employed discrete-choice models to capture interactions between flight schedules and passenger demand (Atasoy et al. 2014, Dong et al. 2016, Wei et al. 2020, Birolini et al. 2021b). However, as discussed above, slot allocation occurs at a more upstream level with less information on competing itineraries, airfares, and market shares, which necessitates a new approach to demand prediction and schedule optimization.

Like our slot-allocation setting, airline disruption recovery is also characterized by complex and nonlinear relationships between flight delays and passenger delays (Barnhart et al. 2014). Whereas early models focused on minimizing airline operating costs, Bratu and Barnhart (2006) incorporated the impact of recovery actions on missed connections and passenger accommodations, showing that airlines could reduce passenger delay by nearly 20% without increasing operating costs significantly. Marla et al. (2017) extended this approach by incorporating flight planning into a model of airline disruption recovery that trades off airline operating costs and passenger-accommodation costs.

In air traffic flow management, Jacquillat (2020) used predictive analytics to infer passenger itineraries and incorporate them into the optimization of Ground Delay Programs at the tactical level (during the day of operations). Our paper applies a similar logic to slot allocation at the strategic level (months in advance of operations), with three implications. First, slot allocation can leverage considerably less information on passenger itineraries, requiring new analytics models to predict passenger demand at a passenger-route level. Second, the WASG create interdependencies between slot allocations across a full season, leading to much larger optimization problems. Third, slot allocation can result in large flight displacements, making the static approximation approach from Jacquillat (2020) to handle nonlinearities invalid. This requires, in turn, a new mixed-integer nonlinear programming formulation and dedicated solution algorithms.

1.3. Contributions and Outline

This paper develops a passenger-centric approach to airport slot allocation that combines (i) an analytics model to predict passenger demand at the passenger-route level, (ii) an optimization formulation to trade off schedule displacement and passenger connectivity in slot allocation, and (iii) a solution algorithm to solve the resulting mixed-integer nonconvex programming at the largest schedule coordinated airports. Specifically, this paper makes the following contributions:

1. *Proposing a new paradigm for passenger-centric slot allocation at schedule-coordinated airports (Section 2)*: We demonstrate, via simplified examples, that minimizing schedule displacement may not result in optimal outcomes for passengers. Accordingly, we propose a novel slot allocation that balances flight-level and passenger-level objectives.

2. *Developing machine-learning models to predict passenger demand at the passenger-route level (Section 3)*: Using a data set of 343,829 observations, we engineer features to reflect airport, airline, and itinerary characteristics. Using random forest models, we achieve excellent out-of-sample performance on existing routes, but also on new routes (which did not appear in historical observations). In particular, our approach outperforms gravity-based benchmarks from the literature, which omit attributes on airline characteristics and itinerary quality.

3. *Formulating a multiobjective optimization model for passenger-centric slot allocation: PSAM with Passengers considerations (PSAM-Pax) (Section 4)*: PSAM-Pax optimizes slot assignments based on one flight-centric metric—minimizing schedule displacement—and two passenger-centric objectives—minimizing the number of lost passengers due to infeasible connections and minimizing connecting times. Because of the flight-passenger interactions, PSAM-Pax is formulated as a mixed-integer nonconvex (bilinear) program.

4. *Developing algorithms to solve the passenger-centric model efficiently (Section 5)*: To solve PSAM-Pax, we propose an approximation scheme that iterates between a mixed-integer linear program for flight scheduling and passenger accommodation. We also embed the proposed approximation scheme into a large-scale neighborhood search procedure inspired by Ribeiro et al. (2019). This approach ultimately yields solution algorithms that can scale to the largest schedule coordinated airports, while accounting for the passenger-centric considerations.

5. *Characterizing the benefits of capturing passenger-level considerations in airport slot allocation (Section 6)*: We implement PSAM-Pax using real-world data from the Singapore Changi Airport—a high-growth airport and

Figure 1. (Color online) Eight Flights Requested by Airlines

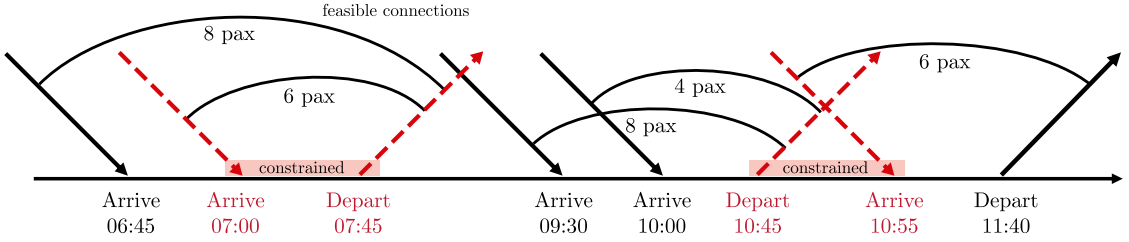

Notes. The predicted number of passengers (pax) is shown on each arc. Boxes indicate constrained blocks of time. Arrows indicate flight requests that must be displaced.

a major connecting hub. Results show that solutions can be obtained in manageable computation times, consistent with implementation requirements. From a practical standpoint, PSAM-Pax can provide significant improvements to slot-allocation outcomes, by trading small increases in schedule displacement (1%) for considerable improvements to passenger-level metrics (50% reduction in lost passengers and 16% decrease in connecting times). We also report additional results from the Lisbon Airport, which show the robustness and generalizability of our findings. Ultimately, the results from this paper lay out the way toward new practices and policies supporting airport schedule coordination that can create more attractive flight schedules from the passengers' perspective.

2. Modeling Framework: From Predictive to Prescriptive Analytics

2.1. Proof of Concept

Let us underscore the effects of slot-allocation optimization on passenger itineraries. We begin with a stylized example of flight requests in Figure 1. This example comprises eight slot requests, corresponding to five arrivals and three departures. Passenger demand leads to 32 predicted connections between five flight pairs, with a minimum connecting time of 45 minutes. We also assume that the airport is already scheduled at capacity during two time blocks (for instance, due to grandfathered slots), indicated with boxes. Therefore, four flight requests must be displaced.

The following figures illustrate different slot-allocation outcomes, with different objectives. Figure 2 shows the benchmark allocation, minimizing schedule displacement alone. This solution displaces each of the four flights requested during capacity-constrained periods to the closest available slot. The resulting schedule has a total displacement of 20 minutes. However, 10 passengers are lost, and, among feasible connections, the total connecting time increases by 100 passenger minutes.

Next, Figure 3 shows an allocation that balances schedule displacement and lost passengers. In this case, we displace conflicting itineraries to maintain feasible connections. This results in a schedule with arrival and departure waves separated by the minimal connecting time of 45 minutes. The resulting schedule has a total displacement of 40 minutes, but all connections remain feasible, and the connecting time is increased by 370 passenger minutes.

Last, Figure 4 balances schedule displacement, lost passengers, and the increase in connecting time. Because the 10:45 departure receives only four connecting passengers from the 10:00 arrival, we sacrifice this connection rather than lengthening the connecting time for the eight passengers arriving at 09:30. This strategy trades off itinerary feasibility (by ensuring sufficient connecting times) and attractiveness (by disincentivizing overly long connections). The resulting schedule has a total displacement of 30 minutes, four lost passengers, and an increase in connecting times by 190

Figure 2. (Color online) Minimal Schedule Displacement (Disrupted Itineraries in Dashed Lines)

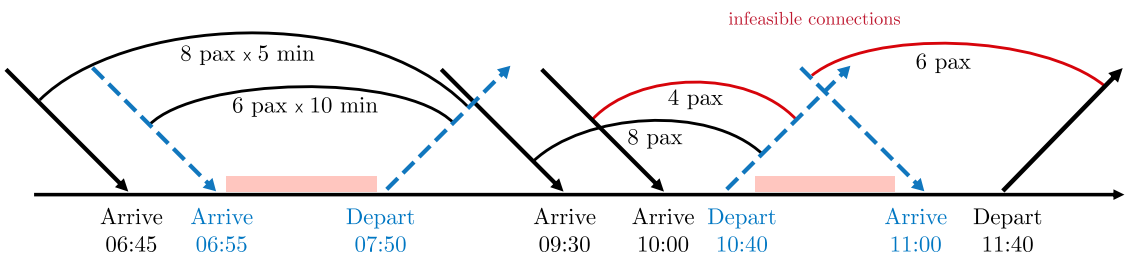

Note. Pax, passengers.

Figure 3. (Color online) Balancing Schedule Displacement and Lost Passengers

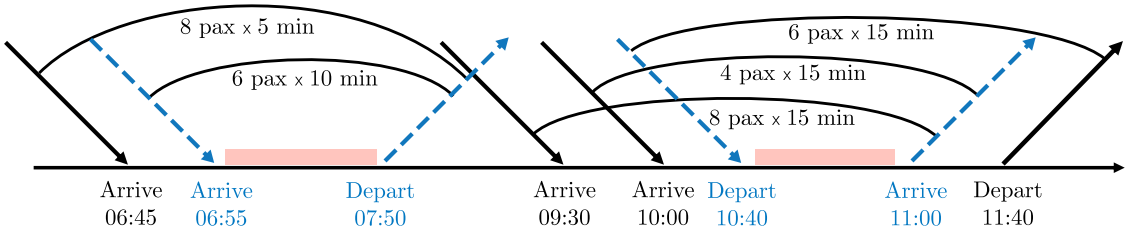

Note. Pax, passengers.

passenger minutes—striking a middle ground between the two previous cases.

This simplified example highlights that minimizing flight displacement does not necessarily translate into the most attractive itineraries for passengers. Instead, the impact of slot allocation can be mitigated by maximizing feasible itineraries and minimizing connecting times. Obviously, reality is more complex, due to the large sample of itineraries from competing airlines, aircraft-capacity constraints, uncertainty regarding passenger demand, etc. In response, this paper combines predictive and prescriptive analytics toward passenger-centric slot allocation in real-world networks.

2.2. Modeling Framework

Figure 5 outlines our modeling framework. At its core lies the PSAM-Pax optimization model, which takes as inputs airlines' slot requests and the airport's declared capacities. It optimizes slot allocation based on one flight-centric objective (schedule displacement) and two passenger-centric objectives (number of lost passengers and connecting time), while complying with the WASG.

As compared with flight-centric formulations, this model features three complexities. One is that linking slot-allocation decisions to passenger flows creates interdependencies between flight schedules and passenger accommodations. Capturing these interdependencies creates nonlinearities in the objective function, thereby considerably increasing the complexity of the optimization problem. To overcome this challenge, we propose an approximation scheme that optimizes flight-level and passenger-level decisions sequentially and integrate it into large-scale neighborhood search methods for implementation at the largest schedule-coordinated airports.

A second complexity is that slot allocation occurs months before flight operations, prior to passenger bookings. Therefore, accounting for passenger-level costs requires estimates of passenger demand. To this end, we leverage a data set from the Official Airline Guide (OAG) that reports historical passenger flows, as well as a supplemental socioeconomic data set providing demographic and economic information. We engineer features and fit machine-learning models to predict passenger flows. Ultimately, the proposed approach combines predictive analytics—to infer passenger flows in flight networks—and prescriptive analytics—to optimize slot-allocation decisions accordingly.

Finally, there is no one-to-one mapping between the outputs of the predictive model and the inputs of the prescriptive model. Indeed, the optimization model considers, as an input, a set of possible itineraries that can materialize from the airlines' slot requests, which necessitates a dedicated itinerary-construction procedure. Moreover, the data set used for prediction is aggregated by month (e.g., 1,500 passengers travel from London Heathrow Airport (LHR) to Sydney Airport (SYD) through Singapore Changi Airport (SIN) per month), whereas the optimization model treats itineraries at a much more disaggregated level (e.g., five passengers connect from the 8:00–9:30 flight from LHR to SIN and the 10:30–12:30 flight from SIN to SYD). Accordingly, we propose four procedures to disaggregate forecasted

Figure 4. (Color online) Balancing Schedule Displacement, Lost Passengers, and Increase in Connecting Times

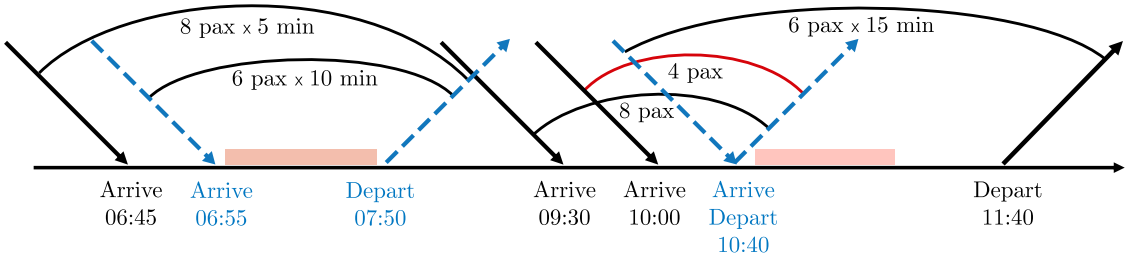

Note. Pax, passengers.

Figure 5. Overview of Modeling Framework, Combining Predictive and Prescriptive Analytics

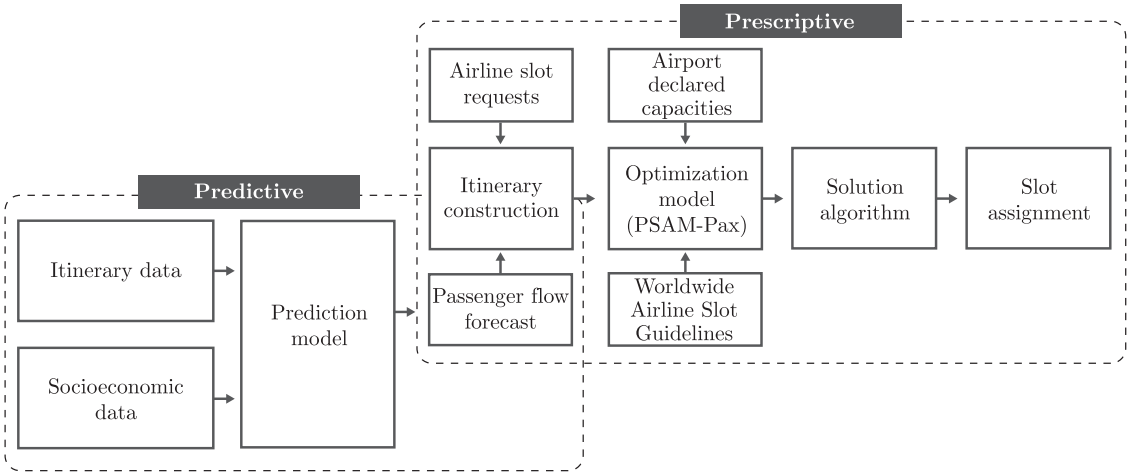

passenger flows across candidate itineraries. In Section 6.1.4, we conduct sensitivity analysis to show that the optimization results are robust to the design of our predictive model and these integration procedures.

3. Predictive Modeling: Passenger-Flow Forecasting

This section presents predictive analytics models for passenger-flow forecasting. We first introduce the data sources and features used for prediction. We then train prediction models, evaluate their out-of-sample performance, and select the best predictive model for slot allocation.

3.1. Data and Feature Engineering

We use data from the Official Airline Guide, which reports passenger flows on a monthly basis for each *passenger-route*. We define a passenger-route z as a sequence of airports and airlines from origin o to destination d. For each origin-destination pair (o, d), each nonstop passenger-route is fully characterized as the airline a operating a nonstop flight from o to d—that is, $z = (o, d, a)$; each one-stop passenger-route is characterized by the connecting airport h, the airline a_1 operating from o to h, and the airline a_2 operating from h to d—that is, $z = (o, h, d, a_1, a_2)$.

We consider the passenger-routes that depart, transit, and arrive at Singapore Changi Airport from 2016 to 2019. We remove all low-volume passenger-routes with less than 30 passengers per month—these represent sporadic connections rather than consistent travel patterns and would introduce noise in our predictions. Similarly, we remove all passenger-routes with two or more connections, which represent a small traffic share (Barnhart et al. 2014). The final data set includes 343,829 observations—that is, 23% of initial observations, but 97% of passenger traffic at Changi.

Each observation records the month, year, airports, airline(s), airline alliance (if any), whether the airline is a low-cost carrier (LCC), and the volume of passengers—our dependent variable. We use a geographic information system to supplement these data with socioeconomic information. We obtain the coordinates of each airport from the OpenFlights Airport Database, and we compute the population and gross domestic product (GDP) within catchment area with a 100-km radius.

Forecasting air-passenger flows is a challenging prediction problem. The demand potential on each passenger-route depends on its own attributes and those of competing ones. Demand is also subject to temporal fluctuations determined by trend and seasonality components. We define the following features to capture these dynamics (summary statistics are reported in Table 1 and a complete list is provided in Online Appendix EC.1).

Historical traffic. To capture temporal patterns, we consider one-year lagged variables for the number of monthly passengers on the passenger-route (the dependent variable), the number of monthly passengers on the origin-destination market, and the geometric mean of annual traffic at the origin and destination airports. For each airport pair, we define a year-over-year growth variable and a seasonality variable (percent deviation from annual mean).Note that historical traffic variables are missing for "new" passenger-routes, whenever the combination of airports and airlines was not offered during the same month in the previous year—these represent roughly 10% of passenger-routes at Changi. Because these values are missing for structural reasons (no one traveled on that route in the past because it did not exist), rather than data-collection issues, we impute the lagged dependent variable as zero and create a binary feature that identifies whether the passenger-route is new. We validate this modeling choice against other imputation

Table 1. Data Summary of Continuous Variables

Variables	Mean	Std. dev.	Min	Median	Max
Monthly Passengers (D.V.)	569	2,218	30	90	48,691
Lagging Monthly Passengers	518	2,175	0	72	47,601
Lagging Market Passengers	6,550	15,637	0	1,259	220,192
Annual Airport Traffic (millions)	25.826	14.373	0.447	23.671	76.862
Passenger Growth	5.950%	226.891	−1.00%	0.06%	8,479%
Seasonality	0.002	0.242	−0.924	−0.012	4.429
Population (millions)	6.876	4.720	0.099	5.911	32.657
GDP (billions USD)	209.275	155.562	2.304	175.046	933.848
GDP Per Capita	33,011	17,508	1,811	28,879	163,823
GDP Difference	0.553%	0.300	0.001%	0.579%	1.000%
Distance (km)	7,634	4,491	282	6,763	22,500
Routing Factor	1.144	0.257	1.00	1.055	11.320

Note. D.V., dependent variable; Std. dev., standard deviation.

procedures (see Online Appendix EC.3) and consider separate linear models for new and existing passenger routes. We impute the remaining lagged historical variables (market passengers, airport traffic, growth, and seasonality) using the other independent variables as predictors—we obtained the best results with *k*-nearest neighbors, both in terms of imputation quality and in terms of out-of-sample performance of the predictive model.

Passenger-route characteristics. These features capture the attractiveness of connecting routes (Paleari et al. 2010, Birolini et al. 2020). We consider a dummy variable indicating whether the passenger-route is nonstop or connecting. Similarly, to differentiate connecting passenger-routes, we consider the flying distance (measured as the total great circle distance across both flight legs, thus capturing travel impedance) and the routing factor (defined as the ratio of the flying distance to the corresponding nonstop distance from origin to destination, which serves as a proxy for the efficiency of a connecting airport for a given origin-destination pair).

Socioeconomic. We consider generative factors to model the origination of air travel: geometric mean of population, GDP, and GDP per capita of the catchment areas of the origin and destination airports (Grosche et al. 2007). We also consider economic distance, measured as the percent difference between the origin and destination GDP (Adler and Hashai 2005).

Airlines and airports. The Changi data set contains 893 and 234 unique airports and airlines, respectively. We group airlines by alliance membership (i.e., Star Alliance, SkyTeam, or Oneworld) and class (LCC versus full-service). We also create dummy variables for the most frequently operating airlines at Changi: Singapore Airlines (SQ), Scoot (TR), and SilkAir (MI). For airports, we consider their geographical location by grouping them by continents.

To assess predictive performance, we partition the data sequentially into training and test sets. Given our lagging variables, we discard the first year (2016) for modeling purposes. We use data from 2017–2018 for training, using 10-fold cross-validation to tune hyperparameters. We then evaluate out-of-sample performance on the test set, comprising observations from 2019.

3.2. Prediction Results

We use machine learning to predict demand on each passenger-route in each month. Given the problem size and context, we focus on simple and scalable techniques: stepwise regression, ridge regression (Hoerl and Kennard 1970), least absolute shrinkage and selection operator (LASSO) (Tibshirani 1996), *k*-nearest neighbors (*k*-NN), and random forests (Breiman 2001). We define a naïve benchmark with an intercept and one feature—the lagging dependent variable (LDV). We report out-of-sample performance in Table 2. Because of continuous changes in the transportation network, a key requirement is to predict flows on new passenger-routes—especially at a fast-growing airport such as Changi. Accordingly, we perform a stratified evaluation on existing versus new passenger-routes. For all predictive methods, we apply both a single-model approach (fitting one model for all itineraries) and a two-model approach (fitting one model for new itineraries and another one for two itineraries). These two approaches reflect an underlying bias-variance trade-off, and we leverage cross-validation to perform model selection.

Across all observations, the benchmark leads to 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. Not surprisingly, it performs poorly on new passenger-routes. Regularized linear regression models retain three features with nonzero coefficients (for stepwise regression and LASSO) or not-too-small coefficients (for ridge regression): the lagged dependent variable, nonstop versus connecting

Table 2. Final Model Performance on Test Data

	Entire population				Existing passenger-routes			New passenger-routes		
		o.o.s.			o.o.s.			o.o.s.		
Method	c.v. R^2	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE
Single model for all observations										
Naïve (LDV only)	0.956	0.960	137	436	0.974	129	392	0.119	206	709
Stepwise 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**
Separate models for new and existing passenger-routes										
Stepwise regression	0.967	0.967	131	401	0.971	120	390	0.589	222	484
Ridge regression	0.967	0.962	131	428	0.971	120	390	0.496	228	486
LASSO	0.967	0.962	131	428	0.971	120	390	0.496	227	486
k-NN	0.990	0.965	126	410	0.970	121	397	0.451	171	507
Random forests	0.997	0.968	120	390	0.972	118	388	0.661	137	440

Notes. Bold indicates the model with the strongest predictive performance. c.v., cross-validation; o.o.s., out-of-sample.

routes, and existing versus new passenger-routes. As with the naïve benchmark, these models perform well on the population as a whole, but performance deteriorates on new passenger-routes.

Note, next, that a two-model approach improves the performance of regularized linear regression methods, especially in predicting demand on new passenger-routes. For instance, the MAE and RMSE decrease to 222 and 484, respectively, for stepwise regression. This suggests that linear models are not flexible enough to capture the underlying differences between new and existing itineraries, warranting two separate fitting procedures to reduce model bias. In contrast, nonlinear models such as *k*-NN and random forests are mode-flexible to provide differentiated predictions for the two itinerary segments, justifying the selection of a one-model approach toward variance reduction.

Ultimately, we obtain the best performance with a one-model random-forests approach, which yields an out-of-sample R^2 of 0.969, an MAE of 121 passengers per month, and an RMSE of 388 passengers per month. Most noticeably, by capturing complex nonlinear relationships between continuous features and passenger flows, the random forests model performs remarkably well on new passenger-routes—performance drops only slightly, to an MAE and RMSE of 129 and 408 passengers per month, respectively, on new passenger-routes. As a data-driven, nonparametric method, *k*-NN also performs reasonably well; however, it falls short of random forests on new passenger-routes, with an RMSE of 482 passengers per month.

To underscore this point, Figure 6 plots predicted versus observed values across existing and new passenger-routes. For random forests (Figure 6(a)), the values for new itineraries are scattered appropriately

Figure 6. (Color online) Predicted vs. Observed Values for Existing and New Passenger Routes

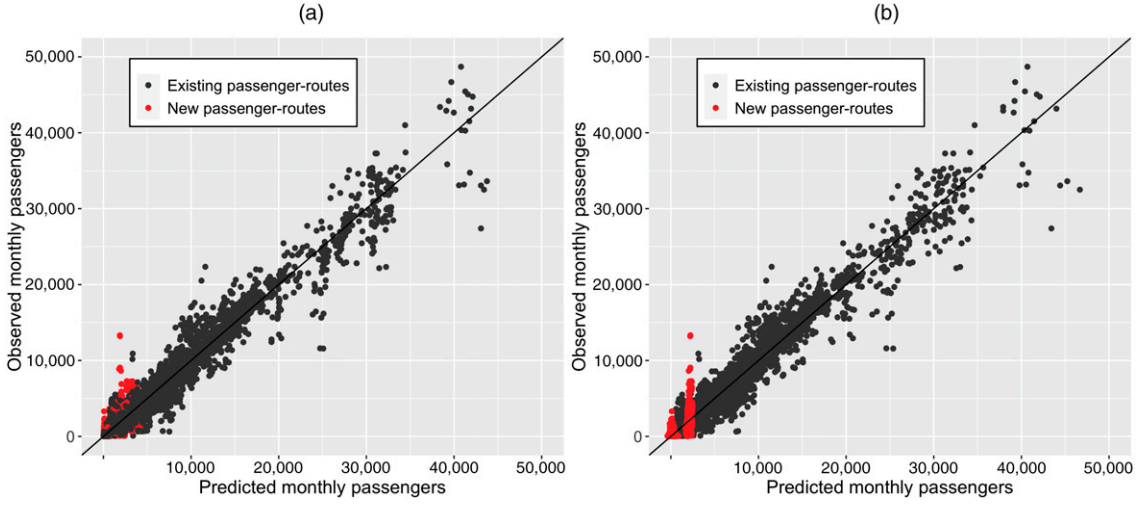

Notes. (a) Random forests (one-model approach). (b) Stepwise regression (two-model approach).

Figure 7. (Color online) Feature Importance of Random Forests

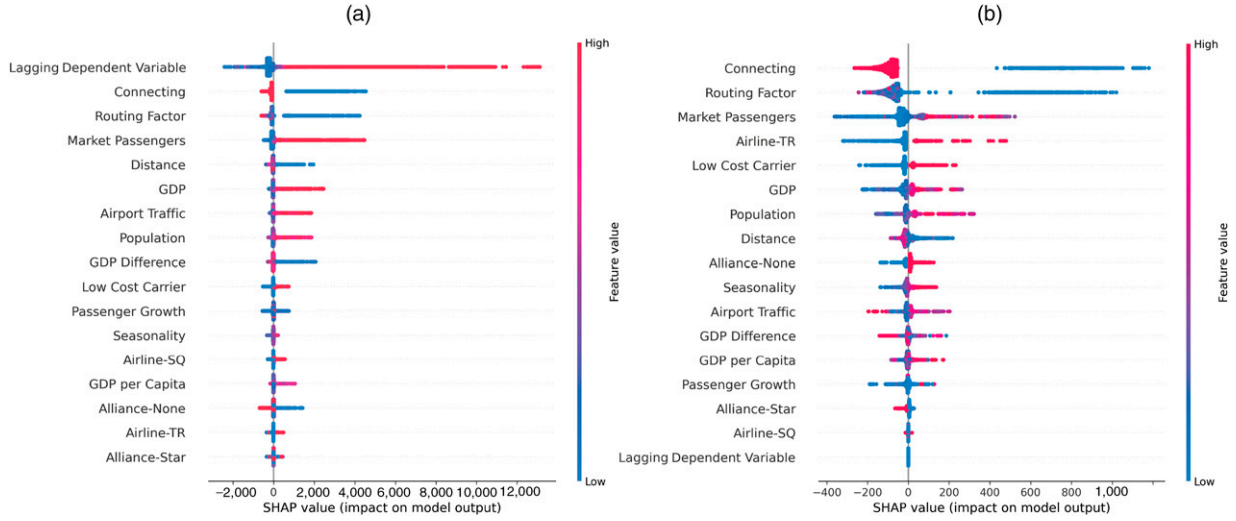

Notes. (a) Entire population. (b) New passenger routes.

along the 45-degree line. In contrast, the best linear model—that is, stepwise regression using a two-model approach (Figure 6(b))—underestimates demand on new passenger-routes, indicating that the simplicity of the linear models restricts their ability to make fine-grained predictions for new itineraries.

Finally, to understand the explanatory power of the various features, Figure 7 reports the Shapley Additive Explanation values for random forests (Lundberg and Lee 2017). When evaluating the model on the entire test set (Figure 7(a)), the lagging dependent variable is, by far, the most important feature. However, the (imputed) lagging variable becomes the least important feature for new passenger-routes (Figure 7(b)). This observation reinforces that models that rely upon historical passenger flows will perform poorly on new passenger-routes. In contrast, random forests provide more accurate predictions by leveraging features characterizing itinerary quality, airline characteristics, and airport characteristics. We quantitatively assess the impact of these variables by dropping them from our specification, which yields a significant degradation of model performance. Specifically, dropping the routing factor and connecting dummy results in an increase of MAE on new passenger-routes from 129 to 187 (+45%) and up to 208 (+61.2%) when airline and airport features are also omitted. Complete results for the comparison with these reduced gravity-based benchmarks are reported in Online Appendix EC.2. Collectively, these results underscore the contribution of our feature engineering and machine-learning approaches, as compared with existing ones.

4. Prescriptive Modeling: Passenger-Centric Slot Allocation

This section formalizes our passenger-centric slot-allocation optimization approach. This approach builds upon the PSAM from Ribeiro et al. (2018), but augments it by incorporating information on passenger itineraries and balancing flight-level objectives (schedule displacement) and passenger-centric objectives (number of lost passengers and connecting time).

4.1. Slot Requests and Slot-Allocation Constraints

The PSAM-Pax model optimizes the time assigned to each slot request while complying with airport capacity constraints and the WASG regulations. To this end, the model requires two primary inputs: (i) the airport's declared capacities, and (ii) the airlines' slot requests. Declared capacities are specified by the schedule-coordinated airport at the beginning of each year. Slot requests are submitted by the airlines five months prior to the beginning of each season (Summer or Winter). Following the code provided by chapter 6 of the Standard Schedules Information Manual (IATA 2014), slot requests include information about the type of operation (arrival versus departure), airline, date, time, and frequency (e.g., daily, weekly, or weekdays). Each slot request is assigned a priority: historic slots (slots operated the previous year and complying with use-it-or-lose-it rules), change-to-historic slots (historic slots with a requested change in schedule or in aircraft capacity), slots from a new-entrant airline, and other slots. An illustrative example is provided in Online Appendix EC.4.

The slot-allocation problem must adhere to the rules and priorities specified by the WASG:

1. *Declared capacity constraints*: Allocated slots must not exceed declared capacities. These may apply at different timescales (e.g., five minutes, 15 minutes, or 60 minutes) and may differ depending on the type of operation (e.g., capacity for arrivals, departures, or total number of movements).
2. *Schedule regularity*: Flights belonging to the same slot request must be allocated to the same time of the day throughout of the season.
3. *Slot priorities*: Historic slot requests are given the highest priority, and equal treatment is then given to slots from all other priority classes (change-to-historic slots, slots from new entrants, and other slots). However, 50% of available capacity must be allocated to new-entrant slots if enough requests are submitted, and change-to-historic slots must be allocated at their historic times if the new requested time is not available.
4. *Turnaround times*: A sufficient turnaround time must be applied to arrival-departure pairs.

4.2. Defining Passenger-Centric Inputs

The PSAM-Pax adopts a passenger-centric approach to make slot-allocation outcomes more consistent with passenger demand. Accordingly, PSAM-Pax optimizes a flight-centric objective (minimizing schedule displacement) and two passenger-centric objectives (minimizing the number of infeasible passenger connections and the increase in connecting times). As such, the PSAM-Pax leverages the passenger-flow forecasts obtained with our predictive model (Section 3). However, defining the relevant itineraries and corresponding cost parameters is complicated by two factors:

- *Unknown candidate itineraries*: Given the endogeneity of passenger itineraries with respect to slot allocation, we need to define the set of candidate itineraries on which passengers might travel. We propose a procedure for itinerary construction that pairs arrival- and departure-slot requests based on feasibility criteria, including sequentiality in time and space, minimum connecting time, and an upper bound on the maximum detour (proxied by the routing factor).
- *Unknown passenger demand on each candidate itinerary*: Combining the outputs of the predictive model and the inputs of the optimization model is not straightforward because the two models have different units of analysis: The predictive model estimates passenger flow for each *passenger-route* at the monthly level—consistently with the underlying OAG data—whereas the optimization model considers passenger flows for each specific *itinerary*—consistently with the fine-grained slot-allocation decisions. To circumvent this challenge, we distribute monthly demand predictions across all candidate itineraries.

To tackle these challenges, we proceed in three steps, outlined in Figure 8 and detailed below.

4.2.1. Itinerary Construction. Itineraries are uniquely identified by the origin airport, connecting airport (if any), destination airport, airline(s), and reference time. To define the set of itineraries, we begin with the set of airline slot requests $\mathcal{S}$. Each request represents a specific flight, characterized by arrival time (or departure time), airline, and origin airport (or destination airport). For nonstop itineraries, any single slot request constitutes a feasible itinerary. For connecting itineraries, we pair arrival and departure slot requests if (i) both flights are operated by the same airline alliance, (ii) the connecting time lies between 45 minutes and 6 hours, and (iii) the "routing factor" is less than 1.5 (i.e., the connecting airport increases flying distance by at most 50%).

Figure 9(a) illustrates how a single arrival slot request can generate many feasible connecting itineraries. In this example, $\mathcal{S}^C(i,j) = \{(i_1,j_2),(i_1,j_3),(i_1,j_4)\}$ includes three 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 because it allows less than 45 minutes for transfer. Each departure can similarly create multiple itineraries.

4.2.2. Identification of Alternative Itineraries. To account for passenger reaccommodations, we define the set of all feasible alternatives for each itinerary—that is, alternative slot requests for nonstop itineraries and alternative pairs of slot requests for connecting itineraries. We define these sets by assuming a time window of ± 4 hours; that is, for an itinerary with a reference time at time t, we consider alternatives itineraries with reference time $t' \in [t-4, t+4]$.

Figure 9(b) illustrates alternatives for itinerary p_1, constructed from slot request i_2 (to SYD). If, (i_2,j_6) becomes infeasible due to displacement decisions or capacitated flights, passengers can be recaptured by the following alternatives: $\mathcal{S}_p^C(i,j) = \{(i_1,j_5),(i_1,j_6),(i_1,j_7),(i_2,j_7),(i_3,j_7)\}$, as long as that they themselves meet the feasibility criteria defined above.

To account for the costs of reaccommodation, we define δ_{pd} as the minimum connecting time across the set of alternatives $(i,j) \in \mathcal{S}_p^C(i,j)$ for day d and itinerary p. This is written in Equation (1), where $\gamma_{i,j}$ denotes the connecting time between slot requests i and j, and B_{id} (B_{jd}) is equal to one if slot request i (j) is requested on day d, and zero otherwise.

$$\delta_{pd} = \min_{(i,j)\in\mathcal{S}_p^C(i,j)} \gamma_{i,j} B_{id} B_{jd}. \tag{1}$$

4.2.3. Passenger Distribution. We then distribute passenger flows (predicted at the monthly level on each passenger-route) across itineraries (corresponding to a

Figure 8. Itinerary Construction Schematic Representation

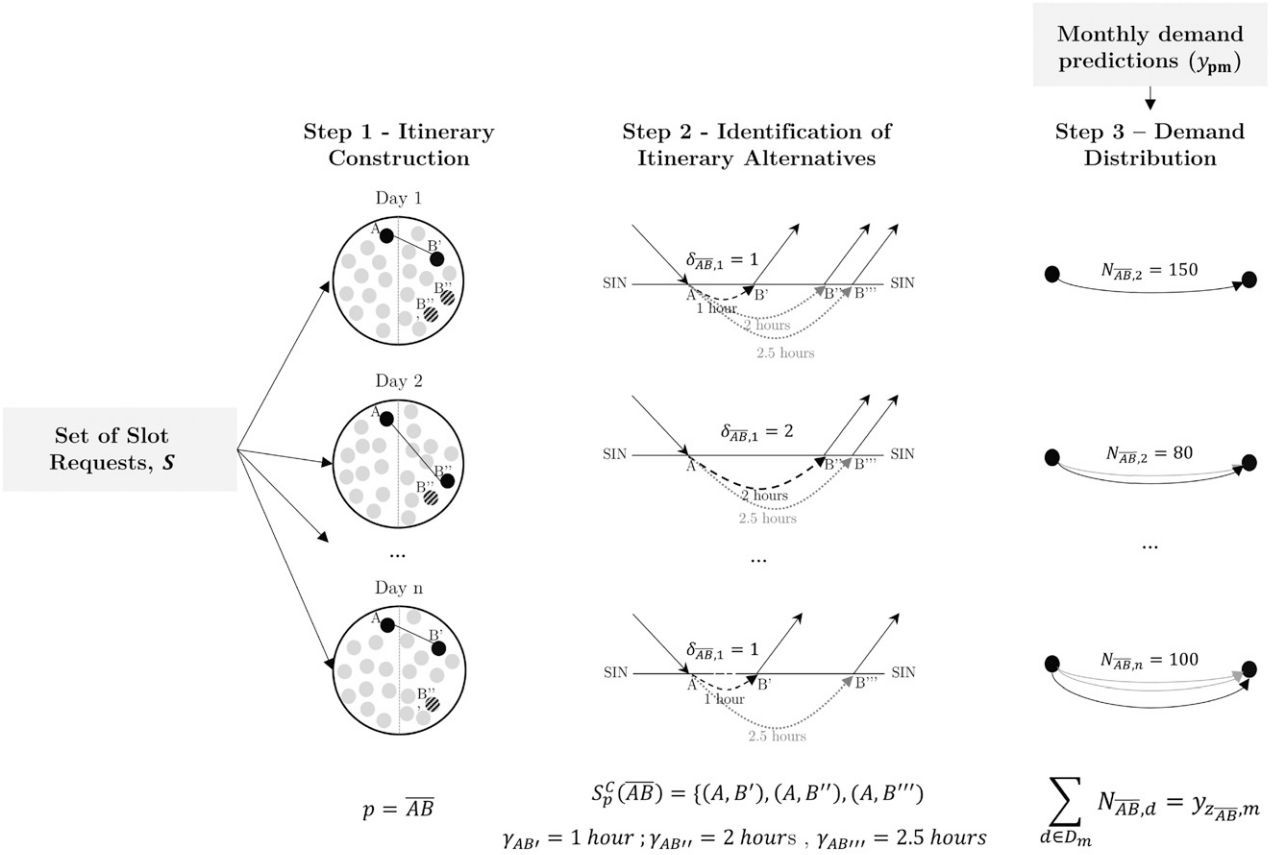

Step 1 – Itinerary Construction
The set of connecting itineraries $p \in P^C$ is constructed by pairing arrival and departure slots with unique origin-destination pairs, based on the following feasibility criteria: (i) arrival-departure slots operated by the same airline alliance; (ii) connecting time between 45 minutes and 6 hours; (iii) maximum routing factor of 1.5.

Step 2 – Identification of Itinerary Alternatives
To account for passenger reaccommodations, a set of feasible alternatives S_p^C for each itinerary $p \in P^C$ is defined. This involves identifying all slot pairs that can accommodate passengers in each itinerary within a time window of 4 hours in relation to the reference connecting time δ_{pd} (i.e. minimum itinerary connecting time) .

Step 3 – Passenger Distribution
We use the predictive model developed in Section 4 to predict the monthly passenger flows across all passenger routes, y_{zm}. For each passenger-route, we can have several alternative itineraries differentiated by their departure/arrival times and day of operations. Different methods are used to distribute the monthly demand predicted in each passenger route y_{zm} across all daily itineraries N_{pd}

specific day and time) proportionally to a weighting factor ζ_{pd} (equal to zero if itinerary p is not offered on day d):

$$N_{pd} = \frac{\zeta_{pd}}{\sum_{d \in D_m} \sum_{p \in P_z} \zeta_{pd}} y_{z_p m}. \quad (2)$$

Here, $y_{z_p m}$ denotes the passengers forecasted to travel on passenger route z_p in month m—subscript p is added to highlight that route z_p is the one from which itinerary p stems—and N_{pd} is the number of passengers forecasted to travel on itinerary p on day d.

Note that the weights cannot be estimated due to the different levels of data aggregation. As such, any distribution procedure is necessarily approximate. In Online Appendix EC.5, we perform sensitivity analysis showing the robustness of the model with respect to the choice of the weighting factors (based on aircraft capacity, time-of-the-day preferences, etc.).

Figure 9. (Color online) Illustration of Itinerary Preprocessing

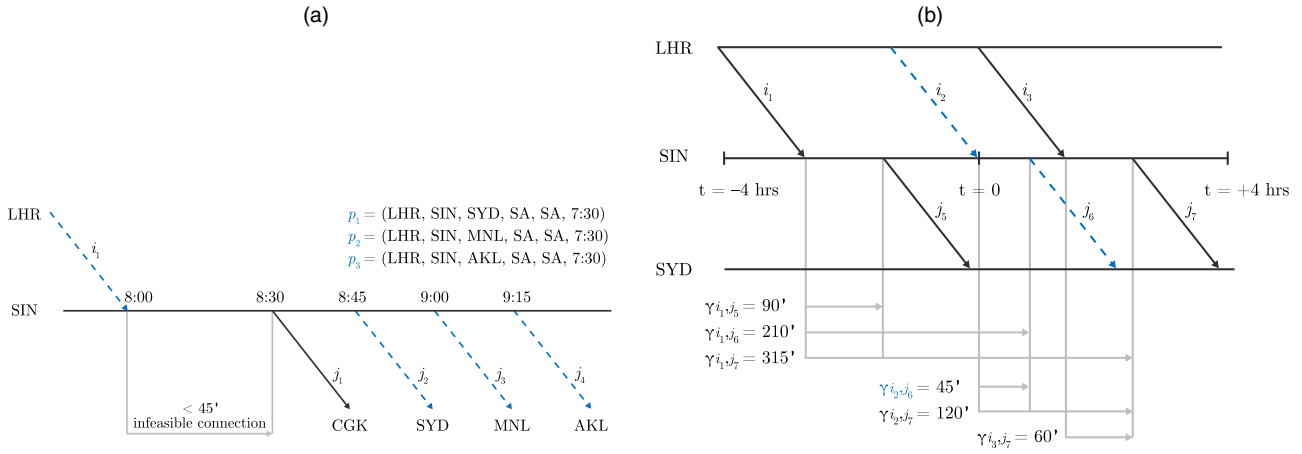

Notes. (a) Itinerary construction example. (b) Identification of alternative itineraries example.

4.3. PSAM-Pax Formulation

4.3.1. Inputs: Preliminaries.

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

4.3.2. Inputs: Slot Requests.

$\mathcal{S}$ = 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
$\mathcal{S}_H/\mathcal{S}_{CH}/\mathcal{S}_{NE}/\mathcal{S}_O \subset \mathcal{S}$ = subset of "historic," "change-to-historic," "new-entrant," and "other" slots
$\overline{\mathcal{T}}_i^{CH} \subset \mathcal{T}$ = subset of time periods during which a change-to-historic slot $i \in \mathcal{S}_{CH}$ can be scheduled
$\mathcal{N} \subset (\mathcal{S} \times \mathcal{S})$ = subset of slot pairs $(i \times j) \in \mathcal{S} \times \mathcal{S}$ with an *aircraft* connection
A_{it} = 1 if slot $i \in \mathcal{S}$ is requested to operate no earlier than period $t \in \mathcal{T}$; 0 otherwise
H_{it} = 1 if slot $i \in \mathcal{S}$ was operated in the previous year no earlier than period $t \in \mathcal{T}$; 0 otherwise
B_{id} = 1 if slot $i \in \mathcal{S}$ is requested to operate on day $d \in \mathcal{D}$; 0 otherwise
$T^{\max}/T^{\min}$ = maximum/minimum change in aircraft turnaround times
ΔH_{ij} = historic aircraft turnaround time between i and j, i.e., $\Delta H_{ij} = \sum_{t \in \mathcal{T}}(H_{jt} - H_{it})$, $\forall (i,j) \in \mathcal{N}$
ΔA_{ij} = requested aircraft turnaround time between$i$ and j, i.e., $\Delta A_{ij} = \sum_{t \in \mathcal{T}}(A_{jt} - A_{it})$, $\forall (i,j) \in \mathcal{N}$

The set of slot requests $\mathcal{S}$ is partitioned into departure and arrival slots to comply with the corresponding declared capacities. It is also partitioned into categories based on historical precedence: historic, change-to-historic, new entrant, and other slots. Change-to-historic slots are subject to additional scheduling rules: "CR" slots can be scheduled any time between the requested and historic slot times, and "CL" slots can only be scheduled at the requested or the historic slot times. We capture these rules by means of the subset of acceptable time periods $\overline{\mathcal{T}}_i^{CH}$ for slot request $i \in \mathcal{S}_{CH}$. The remaining parameters define slot requests and aircraft connections.

4.3.3. Inputs: Airport Declared Capacity.

$\mathcal{C}$ = set of capacity timescales, indexed by c (e.g. 15-minute, 60-minute periods), $\{1, \ldots, C\}$
$CR_{tdc}^{\text{dep}}/CR_{tdc}^{\text{arr}}/CR_{tdc}^{\text{tot}}$ = departure/arrival/total capacity 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)

4.3.4. Inputs: Passenger Itineraries.

$\mathcal{P}$ = 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 pairs of slot requests $(i,j) \in \mathcal{S}_p$ that can accommodate passengers on itinerary $p \in \mathcal{P}^N$
Q_i = number of aircraft 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 connecting time for itinerary $p \in \mathcal{P}_C$ on day $d \in \mathcal{D}$
σ_p = minimum feasible connecting time for passengers traveling on itinerary $p \in \mathcal{P}$,

As following the procedure outlined in Section 4.2 to define (i) the set of candidate itineraries, (ii) the set of reaccommodation opportunities, and (iii) estimates of passenger demand on each itinerary and each day. To estimate reaccommodation costs, we define a reference connecting time δ_{pd} as the minimum connecting time across all alternatives to itinerary p on day d. Finally, PSAM-Pax leverages the aircraft capacity from

each slot request i, denoted by Q_i, and the minimum connecting time for each itinerary p, denoted by σ_p (assumed to be 45 minutes in this paper).

4.3.5. Decision Variables.

Y_{it} = if slot $i \in \mathcal{S}$ is rescheduled to operate no earlier than period $t \in \mathcal{T}$; 0 otherwise
X_i^+ / X_i^- = displacement of slot $i \in \mathcal{S}$ if rescheduled to a later/earlier time
Z^C_{pijd} = number of passengers on itinerary $p \in \mathcal{P}^C$ assigned to arrival slot i and departure slot $j \in \mathcal{S}^C_p$ on day $d \in \mathcal{D}$
Z^N_{pid} = number of passengers on itinerary $p \in \mathcal{P}^N$ assigned to slot $i \in \mathcal{S}^N_p$ on day $d \in \mathcal{D}$
Z^0_{pd} = number of passengers on itinerary $p \in \mathcal{P}$ that are lost, or not accommodated
λ_{pij} = if the connection between slots $i, j \in \mathcal{S}^C_p$ is infeasible for itinerary p; 0 otherwise
τ_{pijd} = increase in connecting time of passengers for itinerary $p \in \mathcal{P}^C$ on day $d \in \mathcal{D}$, if accommodated to slots $i, j \in \mathcal{S}^C_p$.

The flight-scheduling variables Y_{it} comprise the primary decision variables. The variables X_i^+ and X_i^- define the displacement of each slot request. Subsequent variables estimate passenger accommodations, identify infeasible connections, and quantify the increase in connecting time. All of them are secondary variables in the sense that the slot-allocation model does not assign passengers to flights, but, rather, tracks the impact of slot assignments on passenger bookings.

4.3.6. Mathematical Formulation.

Objectives :

$$O_1 = \min \sum_{i\in\mathcal{S}} \sum_{d\in\mathcal{D}} (X_i^+ + X_i^-) B_{id}, \tag{3}$$

$$O_2 = \min \sum_{p\in\mathcal{P}} \sum_{d\in\mathcal{D}} Z^0_{pd}, \tag{4}$$

$$O_3 = \min \sum_{p\in\mathcal{P}^C} \sum_{i,j\in\mathcal{S}^C_p} \sum_{d\in\mathcal{D}} \tau_{pijd} Z^C_{pijd}, \tag{5}$$

Constraints :

$$\text{s.t.} \quad Y_{it} \geq Y_{i,t+1} \qquad \forall i \in \mathcal{S}, t \in \mathcal{T}, \tag{6}$$

$$\sum_{t\in\mathcal{T}} (1 - A_{it}) Y_{it} = X_i^+ \qquad \forall i \in \mathcal{S}, \tag{7}$$

$$\sum_{t\in\mathcal{T}} A_{it} (1 - Y_{it}) = X_i^- \qquad \forall i \in \mathcal{S}, \tag{8}$$

$$X_i^+ = X_i^- = 0 \qquad \forall i \in \mathcal{S}_H, \tag{9}$$

$$Y_{it} = Y_{i,t-1} \qquad \forall i \in \mathcal{S}_{CH}, \ \forall t \notin \overline{\mathcal{T}}_i^{CH}, \tag{10}$$

$$\sum_{i\in\mathcal{S}_{\text{arr}}} \sum_{t=s}^{s+L_c} (Y_{it} - Y_{i,t+1}) B_{id} \leq CR^{\text{arr}}_{sdc} \qquad \forall t \in \mathcal{T} \mid t < T - L_c + 1, d \in \mathcal{D}, c \in \mathcal{C}, \tag{11}$$

$$\sum_{i\in\mathcal{S}_{\text{dep}}} \sum_{t=s}^{s+L_c} (Y_{it} - Y_{i,t+1}) B_{id} \leq CR^{\text{dep}}_{sdc} \qquad \forall t \in \mathcal{T} \mid t < T - L_c + 1, d \in \mathcal{D}, c \in \mathcal{C}, \tag{12}$$

$$\sum_{i\in\mathcal{S}} \sum_{t=s}^{s+L_c} (Y_{it} - Y_{i,t+1}) B_{id} \leq CR^{\text{tot}}_{sdc} \qquad \forall t \in \mathcal{T} \mid t < T - L_c + 1, d \in \mathcal{D}, c \in \mathcal{C}, \tag{13}$$

$$\sum_{t\in\mathcal{T}} (Y_{jt} - Y_{it}) - \Delta A_{ij} \geq T^{\min} \qquad \forall (i,j) \in \mathcal{N}, \tag{14}$$

$$\sum_{t\in\mathcal{T}} (Y_{jt} - Y_{it}) - \Delta A_{ij} \leq T^{\max} \qquad \forall (i,j) \in \mathcal{N}, \tag{15}$$

$$\sum_{t\in\mathcal{T}} (Y_{jt} - Y_{it}) \geq \min(\Delta A_{ij}, \Delta H_{ij}) \qquad \forall (i,j) \in \mathcal{N} \cap (\mathcal{S}_{CH} \times \mathcal{S}_{CH}), \tag{16}$$

$$\sum_{t\in\mathcal{T}} (Y_{jt} - Y_{it}) \leq \max(\Delta A_{ij}, \Delta H_{ij}) \qquad \forall (i,j) \in \mathcal{N} \cap (\mathcal{S}_{CH} \times \mathcal{S}_{CH}), \tag{17}$$

$$\sum_{t\in\mathcal{T}} (Y_{jt} - Y_{it}) - \sigma_p + (T + \sigma_p) \lambda_{pij} \geq 0 \qquad \forall p \in \mathcal{P}^C, (i,j) \in \mathcal{S}^C_p, \tag{18}$$

$$Z^C_{pijd} \leq N_{pd} (1 - \lambda_{pij}) \qquad \forall p \in \mathcal{P}^C, (i,j) \in \mathcal{S}^C_p, d \in \mathcal{D}, \tag{19}$$

$$\sum_{i\in\mathcal{S}^N_p} Z^N_{pid} + \sum_{i,j\in\mathcal{S}^C_p} Z^C_{pijd} + Z^0_{pd} = N_{pd} \qquad \forall p \in \mathcal{P}, d \in \mathcal{D}, \tag{20}$$

$$\sum_{p\in\mathcal{P}^C} \sum_{j\in\mathcal{S}^C_p} Z^C_{pijd} + \sum_{p\in\mathcal{P}^D} Z^N_{pid} \leq Q_i \qquad \forall i \in \mathcal{S}_{\text{arr}}, d \in \mathcal{D}, \tag{21}$$

$$\sum_{p\in\mathcal{P}^C} \sum_{i\in\mathcal{S}^C_p} Z^C_{pijd} + \sum_{p\in\mathcal{P}^D} Z^N_{pjd} \leq Q_j \qquad \forall j \in \mathcal{S}_{dep}, d \in \mathcal{D}, \tag{22}$$

$$\tau_{pijd} \geq \sum_{t\in\mathcal{T}} (Y_{jt} - Y_{it}) - \delta_{pd} \qquad \forall p \in \mathcal{P}^C, (i,j) \in \mathcal{S}^C_p, d \in \mathcal{D}, \tag{23}$$

$$\boldsymbol{\tau} \geq 0, \ \boldsymbol{X}^+, \boldsymbol{X}^-, \boldsymbol{Z}^C, \boldsymbol{Z}^N, \boldsymbol{Z}^0 \ \text{integer}, \ \boldsymbol{Y} \ \text{binary}. \tag{24}$$

Expressions (3)–(5) formulate a triobjective model: (i) minimizing schedule displacement (3), (ii) minimizing the number of lost passengers (4), and (iii) minimizing the increase in connecting time (5). Following current practice and the literature, we treat schedule displacement as the primary objective and the two passenger-centric objectives as secondary and tertiary objectives. Yet, we explore the Pareto-optimal frontier to identify the trade-offs between these three objectives, by means of the ε-constrained method. Constraints (6)–(17) are adapted from the PSAM (Ribeiro et al. 2018). Constraint (6) ensures that Y_{it} is nonincreasing in t. Constraints (7) and (8) link the scheduling and displacement variables, also providing valid equalities with tight linear-programming relaxations. Constraint (9) ensures that historic slots are not displaced. Constraint (10) ensures that change-to-historic slots are scheduled at appropriate times. Constraints (11)–(13) ensure that the numbers of arrivals, departures, and movements do not exceed the corresponding capacities. Constraints (14) and (15) ensure that aircraft turnaround times do not increase or decrease by more than the allowable limits. Constraints (16) and (17) state

that the turnaround time between two change-to-historic slots remains between the requested and historic ones.

Constraints (18)–(23) incorporate passenger-level dynamics. Constraint (18) defines the feasibility of connections between slot pairs. The term $T + \sigma_p$ serves as a big-M parameter to ensure that the left-hand side is always positive. Constraint (19) ensures that passengers are only assigned to feasible connections. For example, if connection ij is infeasible—that is $\sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it}) - \sigma_p < 0$—then, $\lambda_{pij} = 1$ from Constraint (18) and $Z^C_{pijd} = 0$ from Constraint (19). Constraint (20) ensures that all passengers are accommodated to a nonstop itinerary, accommodated to a connecting itinerary, or not accommodated (due to insufficient aircraft capacities or unavailable itineraries). Constraints (21) and (22) ensure that passenger accommodations comply with aircraft capacities. Finally, Constraint (23) computes the increase in connecting times. The decision variable τ_{pijd} is given by the difference between the allocated connecting time $\sum_{t\in\mathcal{T}}(Y_{jt} - Y_{it})$ and the minimum alternative connecting time δ_{pd} for itinerary $p \in \mathcal{P}$, if positive (so the model is not incentivized to reduce connecting times beyond the reference time). Constraint (24) defines the domain of all variables.

4.4. Discussion

The PSAM-Pax formulation results in a mixed-integer bilinear programming model, due to the product of the increase in connecting time per itinerary and the number of passengers (Expression (5)). This bilinear structure results in a mixed-integer nonconvex quadratic program, thereby significantly increasing computational complexity. In Section 5, we propose a solution approach that approximates the nonlinear formulation with a sequence of mixed-integer programs.

Table 3 compares the size of PSAM-Pax to the baseline PSAM. With passenger-level considerations, the PSAM-Pax formulation is significantly larger than the PSAM one. Indeed, 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. 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 connections, according to the criteria in Section 4.2. As we shall see, empirically, PSAM-Pax remains tractable in real-world settings from the largest schedule-coordinated airports.

Table 3. 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. Solution Algorithm

We propose an approximation scheme to solve the mixed-integer bilinear program at the core of PSAM-Pax (Section 5.1). Section 5.2 then integrates PSAM-Pax into large-scale neighborhood search methods to solve real-world instances at the largest schedule-coordinated airports.

5.1. Solution Approximation

The objective function (Expression (5)) includes bilinear combination of decision variables $\tau_{pijd}Z^C_{pijd}$. We could, in theory, replace the bilinear term with a single variable and add a set of big-M constraints (see Online Appendix EC.6). However, this transformation significantly increases the size of the model—we were not able to find any feasible solution in realistic instances, even after days of computation. Instead, we propose an approximation approach to solve the bilinear model.

First, we construct a feasible solution in two steps by:

1. Optimizing flight-centric slot allocation:

$$\min\left(\sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}(X^+_i + X^-_i)B_{id}\right)$$

$$\text{s.t. : Equations (6)–(17), (24).}$$

2. Optimizing passenger's accommodations with the resulting slot assignment:

$$\min\left(M\sum_{p\in\mathcal{P}}\sum_{d\in\mathcal{D}}Z^0_{pd} + \sum_{p\in\mathcal{P}^C}\sum_{i,j\in\mathcal{S}^C_p}\sum_{d\in\mathcal{D}}\tau_{pijd}Z^C_{pijd}\right)$$

$$\text{s.t. : Equations (18)–(24).}$$

This initial solution can be viewed as a flight-centric benchmark, which first determines the slot schedule without considering passenger itineraries and then replicates passenger assignments for the resulting schedule of flights. To accommodate passengers in step 2, we first minimize the number of lost passengers (the first term in the objective function, weighted by a large scalar M). This provides a benchmark of the number of passenger itineraries supported by the flight-centric slot allocation. We then minimize, as a secondary objective, the connecting time.

As opposed to this sequential approach, our integrated model brings flight-centric and passenger-centric considerations together into a unified framework. To approximate the bilinear objective function, we could devise a coordinate descent approach that alternates

between updating passenger-level decisions by fixing τ_{pijd} and endogenizing Z^C_{pijd}, and updating slot schedules by fixing Z^C_{pijd} and endogenizing τ_{pijd}. Rather, we found it more effective to combine these two approaches into a single step, using a weight-based approach. Namely, we solve the following problem:

$$\min\left(\alpha \sum_{p\in P^C}\sum_{i,j\in S^C_p}\sum_{d\in D}\widehat{\tau}_{pijd}Z^C_{pijd} + (1-\alpha)\sum_{p\in P^C}\sum_{i,j\in S^C_p}\sum_{d\in D}\tau_{pijd}\widehat{Z}^C_{pijd}\right)$$

$$\text{s.t. : Equations (6)–(24),} \tag{25}$$

where $\widehat{\tau}_{pijd}$ and $\widehat{Z}^C_{pijd}$ are approximations obtained from the initial (fight-centric) solution, and α is a calibration parameter. When $\alpha = 1$, the approximation scheme reoptimizes slot allocation by fixing the increase in connecting times. This may provide adequate solutions, as long as the variations in displacement remain small. However, when the passenger-centric solution deviates significantly from the flight-centric one, the approximation of τ_{pijd} may become less accurate. Conversely, when $\alpha = 0$, the approximation scheme penalizes increases in connecting times, but fixes the number of passengers assigned to each itinerary to retain linearity. As such, a value of $\alpha = 0$ yields impractical solutions because passengers can be freely assigned to remote slots. With a positive, but very small, value of α (e.g., $\alpha = 0.001$), the approximation scheme still underestimates the impact of increases in connecting times on passenger accommodations. In-between, intermediate values of α strike a balance between these two approximation strategies. We perform sensitivity analyses in Section 6.1.4 to find appropriate values of α (leading to a value of $\alpha = 0.3$).

Finally, recall that we are dealing with a three-objective optimization problem. To isolate the linearized approximation of the third objective, we explore the solution space by means of the ε-constrained method. Specifically, let X^* and Z^* denote the objective values of the first two objectives from the flight-centric solution. We define θ as the allowable percent-wise increase in total displacement and φ as the allowable percent-wise decrease in the number of lost passengers. We implement the first two objectives through the following constraints:

$$\sum_{i\in S}\sum_{d\in D}(X^+_i + X^-_i)B_{id} \le (1+\theta)\cdot X^*, \tag{26}$$

$$\sum_{p\in P}\sum_{d\in D}Z^0_{pd} \le (1-\varphi)\cdot Z^*. \tag{27}$$

5.2. Integration with Large-Scale Neighborhood Search Algorithm

Because of the problem's complexity, slot-allocation models have only been implemented with exact methods at medium-sized airports. Ribeiro et al. (2019) developed a large-scale neighborhood search (LNS) algorithm, with a constructive heuristic and an improvement heuristic (Figure 10):

1. *Constructive heuristic*: It seeks a feasible solution by subdividing slot requests into equally sized groups by decreasing order of priority and frequency (i.e., the number of days on which the slot request will be scheduled). Slot requests with the highest priority and spanning more days are allocated first, whereas requests with the lowest priority and spanning fewer days are allocated last. The constructive heuristic is parametrized by the number of groups. If the number of groups is set too low, the constructive heuristic is equivalent to exact methods, with long computational times. At the other extreme, many groups lead to myopic optimization, without capturing coupling constraints across slot requests and time periods.

2. *Improvement heuristic*: It seeks improvements to the incumbent solution using a "destroy and repair" approach. At each iteration, it reoptimizes, within a maximum runtime t_o, the slot requests during a selected time window (e.g., from 8:00 to 10:00 or 9:00 to 11:00), while fixing all other slot assignments. It relies on an history-based procedure to orient the search toward neighborhoods where reductions in schedule displacement are more likely. At each iteration, a time window $w \in W$ is selected based on a probability distribution $PT^{(q)}_w$. If this leads to an improved solution, the probabilities $PT^{(q)}_w$ are not changed; otherwise, a penalty reduction is applied to the selected time window and distributed across all the other windows. The size of the time windows specifies the size of the neighborhood: Small windows tend to converge to local optima solutions, whereas longer windows require longer computational times. Whenever PSAM is not solved to optimality within the time limit t_o for a given time window $w \in W$ (hence, slot allocation is challenging "locally"), the algorithm will select a subset of slots in S_w for optimization whenever it returns to time window w. The number of slots to select n^s is governed by a factor ρ, a calibration parameter. This search procedure has been shown to considerably enhance the scalability of the slot-allocation model (Ribeiro et al. 2019). Further details can be found in Online Appendix EC.7

5.2.1. New LNS Framework to Solve PSAM-Pax.

The PSAM-Pax model cannot be readily integrated into this LNS framework due to (i) the multiobjective framework, which necessitates the integration of the ε-constraint method; and (ii) the interdependencies between slots due to the consideration of passenger itineraries, which require a different method to select the slots for optimization at each iteration. We therefore augment the LNS algorithm in three ways:

1. At each iteration, the procedure to select the time window for reoptimization is now guided by the third objective—increase in passenger connecting times, Γ. That is, X^*_i is replaced by Γ^*_i in Equation (EC.12) in the

Figure 10. Schematic Representation of the LNS Algorithm

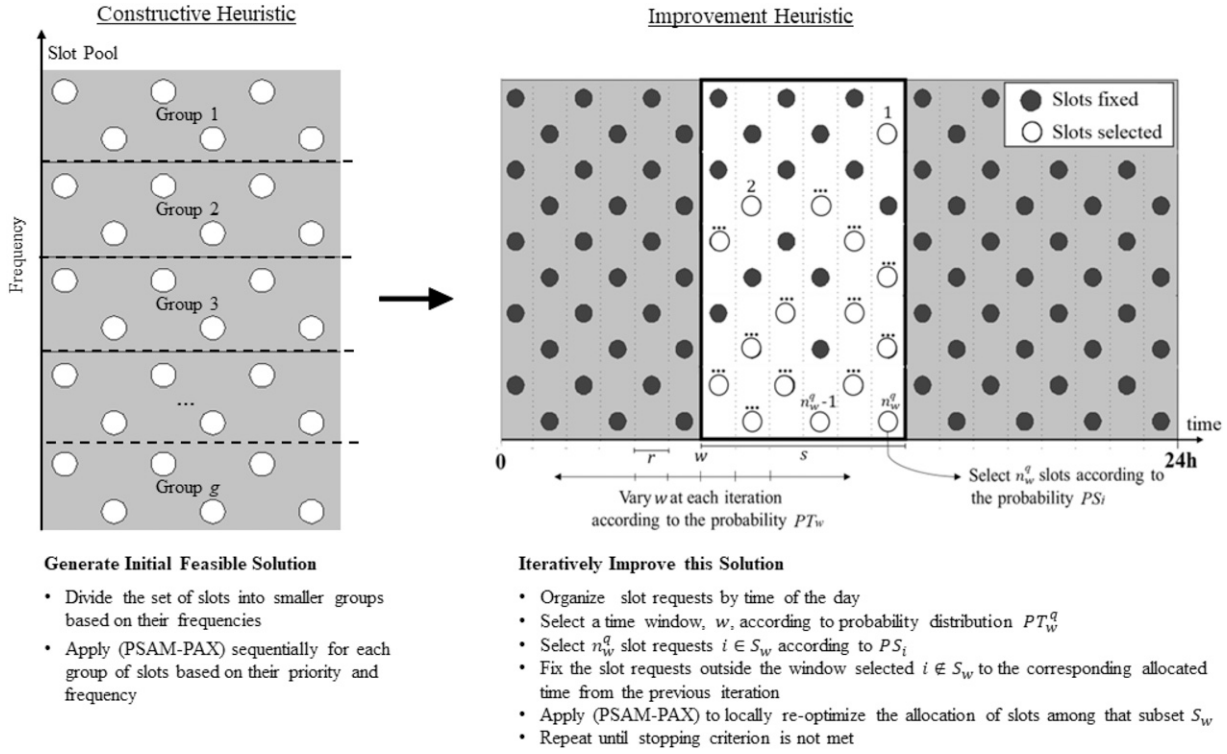

online appendix, and the probability PT_w is only updated if no improvement to the objective Γ is found. The other two objectives are formulated as ε-constraints (see Section 6.1).

2. To capture passenger accommodations, the slots included in the model are no longer restricted to those in the selected time window S_w, but include those that can involve a passenger connection with slots in time window w. The slots that fall outside the time window, however, are fixed to the previous allocated time, $Y_{it}^{(q)} = Y_{it}^{(q-1)}\ \forall i \notin \mathcal{S}_w$. This contrasts with the original LNS algorithm, where all slots $i \notin \mathcal{S}_w$ would be ignored (with a corresponding reduction in the effective declared capacities). Accordingly, the time window for reoptimization needs to be longer for PSAM-Pax than for PSAM, to avoid restricting the algorithm's flexibility.

3. We adjust the optimization problem solved at each iteration to guarantee feasibility with the ε-constraint method. Otherwise, the problem solved at any iteration may be infeasible, even though the overall problem is feasible, due to the random selection of a subset of slots for optimization. For instance, consider the case where $\theta = 0.01$ and $\varphi = 0.5$. In the first iterations, the local slot reoptimization at the core of the LNS algorithm may not be sufficient to meet the aggressive target of 50% reduction in the number of lost passengers. Yet, the overall problem may be feasible—that is, it may be possible to reduce the number of lost passengers by 50% if we increased schedule displacement by more than 1%. To cope with this issue, we introduce a new variable, R, as the deviation from the targeted number of lost passengers. At each iteration, the model minimizes this deviation as a primary objective and then minimizes the "actual" objective function as a secondary objective. As such, the LNS algorithm will prioritize finding a feasible solution and then search for an optimized solution when $R = 0$. If PSAM-Pax is truly infeasible, then the LNS algorithm will minimize the deviation in the number of lost passengers—treating the target in lost passengers as a soft constraint.

The full improvement heuristic is presented in Online Appendix EC.7. We use the same parameters as Ribeiro et al. (2019), except for the size of the time window s due to the stronger coupling constraints stemming from passenger itineraries. The sensitivity of the PSAM-Pax performance to the time windows is reported in Online Appendix EC.7. These results lead to a choice of time windows of five to six hours (which matches the maximum connecting time we are considering of four hours).

6. Experimental Results

We implement our passenger-centric slot-allocation approach using real-world data from the Singapore Changi Airport during the 2019 Summer season

(April–October), provided by the Civil Aviation Authority of Singapore. A total of 254,129 slots were requested by the airlines, distributed across 2,570 slot requests. Changi operates with time-varying runway-capacity limits that span different intervals (15, 30, and 60 minutes) and different types of movements (arrivals, departures, and total). For confidentiality reasons, we cannot report the declared capacities.

First, we consider in Section 6.1 a restricted version of the problem by relaxing some capacity constraints—that is, by only considering a capacity constraint on the total number of movements (departures plus arrivals) per hour. This version of the problem can be solved quickly without the LNS algorithm, facilitating experimentation and sensitivity analyses. We then apply the LNS algorithm to solve the full problem instance in Section 6.2. We also report results from the Lisbon Airport to establish the generalizability of our findings. All models are implemented by using CPLEX 12.7 on a computer with an i7 processor at 3.6 GHz and 8 Gb of RAM.

6.1. Restricted Problem Instance at Changi

6.1.1. Main Results. Table 4 compares the solutions of PSAM and PSAM-Pax. The first row corresponds to the flight-centric benchmark obtained with PSAM. By design, this solution achieves the smallest schedule displacement (157,390 minutes). However, 2,065 passenger connections become infeasible, and, among the feasible connections, the connecting time is increased by 1,826,660 passenger-minutes. The second solution constrains the schedule displacement to its optimal value and, by minimizing passenger-level costs, achieves a minor reduction in connecting times. In comparison, the subsequent solutions yield much more significant benefits for passenger accommodations, at the expense of a small increase in schedule displacement. For example, a 1% increase in schedule displacement reduces connecting time by 33% with the same number of feasible itineraries. Alternatively, we can reduce connecting times by 29% and the number of lost passengers by 40%—still at the cost of a 1% increase in schedule displacement. By allowing further increases in schedule

Table 4. Results for Restricted Problem: Added Connecting Times (in Minutes), as a Function of Schedule Displacement (Rows) and the Number of Lost Passengers vs. the Flight-Centric Baseline (Columns)

		Lost passengers							CPU (min)	
Model	Displacement	2,065 +0%	1,859 −10%	1,652 −20%	**1,239 −40%**	**826 −60%**	413 −80%	0 −100%	Min	Max
PSAM	157,390 +0%	1,826,660 +0%	—	—	—	—	—	—	120 +CPU (min)	120
PSAM-Pax	157,390 +0%	1,826,320 −0.19%	—	—	—	—	—	—	36	36
	157,783 +0.25%	1,711,620 −6.30%	1,715,445 −6.09%	—	—	—	—	—	38	39
	158,177 +0.5%	1,379,325 −24.49%	1,392,475 −23.77%	1,399,620 −23.38%	—	—	—	—	46	59
	158,570 +0.75%	1,305,880 −28.51%	1,314,015 −28.06%	1,335,550 −28.69%	**1,362,345 −25.42%**	—	—	—	43	104
	158,963 +1%	1,211,475 −33.68%	1,219,560 −33.24%	1,235,340 −32.37%	**1,301,575 −28.75%**	—	—	—	44	110
	160,537 +2%	1,082,340 −40.75%	1,088,990 −40.38%	1,111,370 −39.16%	**1,147,755 −37.17%**	**1,245,875 −31.79%**	1,753,740 −3.99%	—	47	201
	163,685 +4%	911,785 −50.08%	920,095 −49.63%	940,245 −48.53%	977,515 −46.49%	1,057,750 −42.09%	1,262,945 −30.86%	—	46	125
	166,833 +6%	774,920 −57.58%	778,150 −57.40%	785,540 −57.00%	862,315 −52.79%	947,960 −48.10%	1,093,670 −40.13%	—	51	105
	169,981 +8%	714,525 −60.88%	720,095 −60.58%	726,380 −60.23%	759,110 −58.44%	855,575 −53.16%	968,610 −46.97%	—	55	116
	173,129 +10%	626,830 −65.68%	663,605 −63.67%	663,610 −63.67%	713,585 −60.93%	791,990 −56.64%	890,020 −51.28%	2,051,855 +12.33%	57	338
	180,998 +15%	533,535 −70.79%	536,125 −70.65%	583,325 −68.07%	617,790 −66.18%	649,155 −64.46%	732,440 −59.90%	1,508,095 −17.44%	69	255
	204,607 +30%	262,425 −85.63%	295,455 −83.83%	282,565 −84.53%	339,325 −81.42%	385,120 −78.92%	461,975 −74.71%	914,850 −49.92%	84	453
	236,085 +50%	184,745 −89.89%	184,745 −89.89%	184,720 −89.89%	199,295 −89.09%	255,120 −86.03%	261,350 −85.69%	680,780 −62.73%	71	523
CPU (min)	Min	36	43	59	64	86	84	321		
	Max	88	84	92	119	453	338	523		

Note. Bold values indicate recommended solutions.

displacement (by 2% from the flight-centric optimum), we can further reduce the number of lost passengers by 80%, while still reducing connecting times. However, pushing the limit toward making all candidate itineraries feasible comes at the much higher cost of a 10% increase in schedule displacement and a 13% increase in connecting time. These results point to a middle ground with intermediate solutions, indicated in bold in the table and also shown in Figure 11.

In summary, our approach can provide substantial benefits by leveraging the highly nonlinear relationships between schedule displacement and the passenger-centric metrics. As such, the proposed solutions can make many more itineraries available and make available connecting itineraries much more attractive, while keeping the deviations from airlines' slot requests within 1%–2% of their optimum. In other words, the choice of *which* slots to displace from their requested times can make the outcome of schedule coordination much more consistent with passenger demand, ultimately creating more travel opportunities for passengers and more revenue opportunities for the airlines.

6.1.2. Computational Performance. Table 4 also reports computational time for the restricted instance. It takes two hours to solve PSAM at the initialization stage and an additional 30 minutes to eight hours to solve PSAM-Pax. Note that (PMAS-Pax) terminates considerably faster when the solution is constrained to remain closer to the flight-centric solution—that is, when the values of θ and φ are smaller. Vice versa, as the first two objectives (i.e., schedule displacement and lost passengers) can deviate more from their flight-centric values, the feasible region of the model becomes larger, increasing solution times.

Overall, despite increasing the size and complexity of the problem, the proposed passenger-centric approach remains tractable using the approximation scheme developed in Section 5. As such, the model and algorithm developed in this paper can provide a new decision tool in support of slot allocation, which returns solutions in acceptable computational times (in view of the strategic slot-allocation problem) and enhances slot-allocation outcomes (in view of passenger demand).

6.1.3. Itinerary-Level Results. To illustrate the aforementioned results at a lower level of granularity, we now focus on a handful of high-impact itineraries—namely, the 15 itineraries with the most passengers out of the all itineraries with an increase in connecting time greater than 100 minutes under the flight-centric solution. Table 5 reports the increase in connecting time for each of these 15 itineraries, as a function of

Figure 11. (Color online) Three-Dimensional Pareto Frontier Between Schedule Displacement (Horizontal Axis), Increase in Connecting Time (Vertical Axis), and Number of Lost Passengers

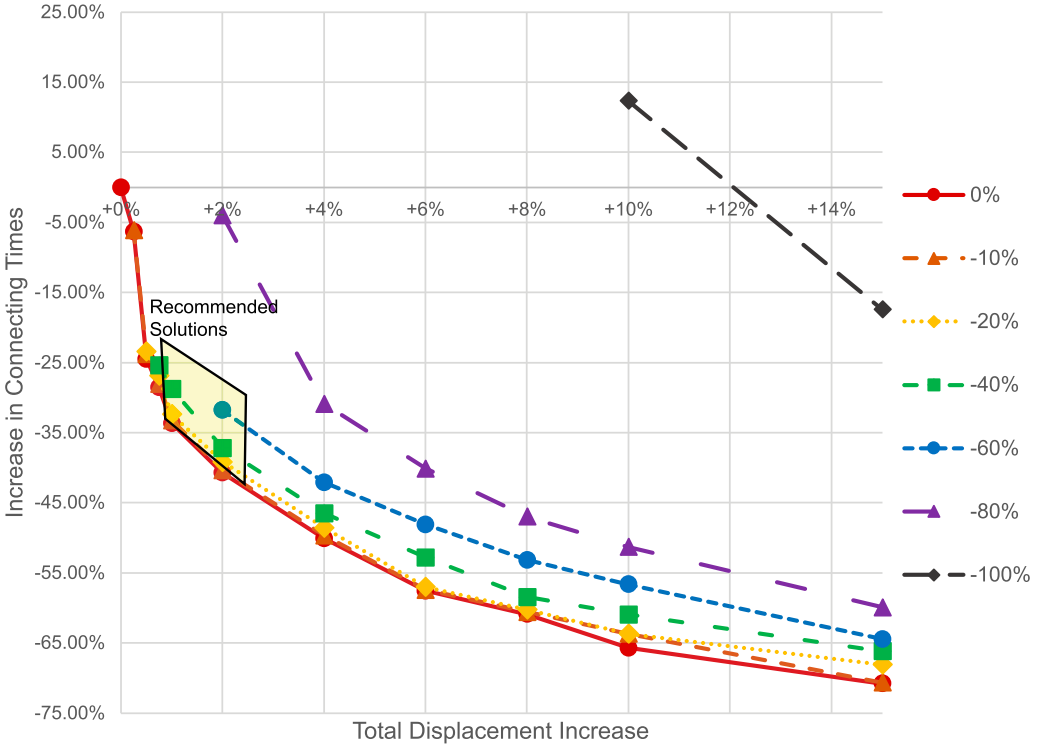

Note. The recommended solutions are shown in the shaded region.

Table 5. Results for High-Impact Itineraries

			Objective values						
		Displacement:	0%	2%	4%	6%	8%	10%	10%
		Lost pax.:	0%	0%	0%	0%	0%	0%	20%
Origin	Dest	# Pax	Increase in connecting time (min)						
EWR	CGK	738	175	0	0	0	0	0	0
EWR	DPS	626	—	0	0	0	0	0	0
EWR	BKK	609	125	0	0	0	0	0	0
PEN	SUB	477	105	45	—	—	—	—	0
EWR	BKK	438	80	0	0	0	0	0	0
SGN	SUB	385	100	40	40	40	0	0	0
MFM	PDG	147	—	105	105	45	105	0	105
HKT	SUB	140	105	45	—	45	0	—	0
SEA	CGK	138	190	—	—	—	—	—	190
IPH	SUB	137	—	45	45	45	0	0	0
MFM	SUB	137	—	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	—	—	—	—	—	105
TSN	SUB	102	105	45	45	45	0	0	0

Notes. The table reports the increase in connecting time for each itinerary, as a function of schedule displacement and lost passengers (Pax). — indicates infeasible itineraries. BKK, Suvarnabhumi International Airport (Bangkok); CGK, Soekarno-Hatta International Airport (Indonesia); DPS, Ngurah Rai (Bali) International Airport; EWR, Newark International Airport; HKT, Phuket International Airport (Thailand); IPH, Sultan Azlan Shah Airport (Malaysia); MFM, Macau International Airport; MNL, Ninoy Aquino International Airport (Philippines); PDG, Minangkabau International Airport (Indonesia; PEN, Penang International Airport (Malaysia); SEA, Seattle-Tacoma International Airport; SGN, Tan Son Nhat International Airport (Vietnam); SUB, Juanda International Airport (Indonesia); TAO, Qingdao Liuting International Airport (China); TSN, Tianjin Binhai International (China).

the increase in schedule displacement and in the number of lost passengers.

The first column reflects the flight-centric solution from PSAM. Under this solution, four itineraries become infeasible. Among the other 11 itineraries, the connecting times increase by one to three hours—at significant costs for passengers. Under the second scenario, a 2% increase in schedule displacement opens 13 of these 15 itineraries for passengers, as opposed to 11 itineraries under the flight-centric baseline. The increases in connecting times are also reduced—0.75 to two hours, as opposed to one to three hours. As the schedule displacement further increases, a handful of itineraries remain infeasible; the increase in connecting times decrease, but with marginal returns. Note that the relationship is not monotonic due to the restricted sample. In the final scenario, all 15 high-impact itineraries become feasible—albeit, at the expense of longer connecting times.

One interesting observation is that all itineraries originating at Newark (NJ) Liberty International Airport (EWR) remain feasible with a zero increase in connecting time. The EWR–SIN flight is a high-profile flight operated by Singapore Airlines—the longest commercial flight in the world. As such, this flight provides the only connection from North America to many locations in Southeast Asia. As a result, any displacement of this flight has significant downstream impact on passenger itineraries. This example, although simple, underscores the key drivers of passenger-centric slot-allocation decisions—maintaining attractive connections through the strategic selection of which slots to displace and which slots to schedule at the requested times.

To further illustrate the impact of passenger-centric slot allocation at the itinerary level, Figure 12 provides a scatterplot of the increase in connecting time and the number of passengers across all itineraries, under the flight-centric solution (Figure 12(a)) and under the PSAM-Pax solution with a 2% increase in schedule displacement

Figure 12. (Color online) PSAM vs. PSAM-Pax Solutions at the Itinerary Level

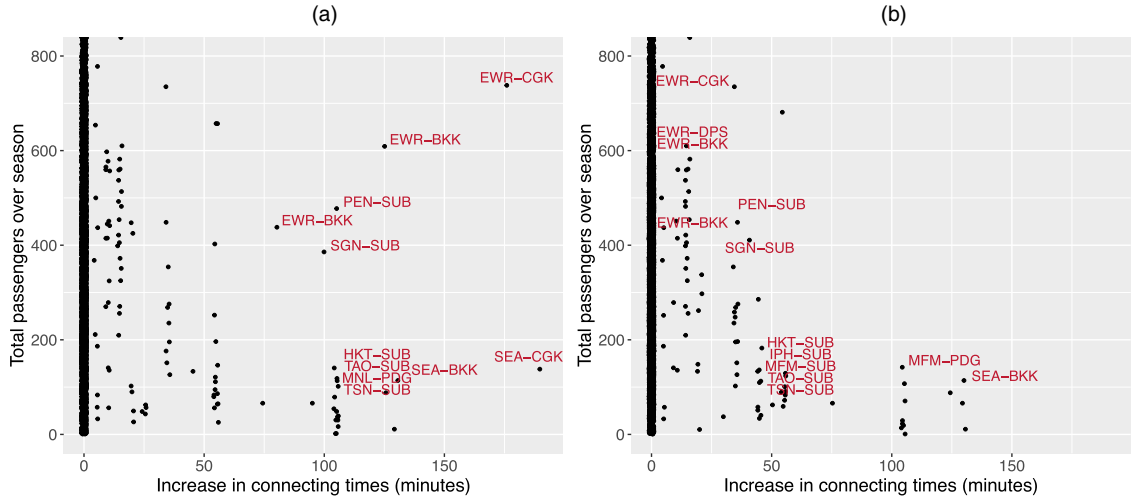

Notes. Labeled itineraries are reported in Table 5. (a) Flight-centric solution. (b) PSAM-Pax solution with 2% increase in displacement.

(Figure 12(b)). Note that all data points move to the left as the schedule displacement increases from 0% to 2%, confirming the significant impact on passenger utility of the proposed passenger-centric approach to slot allocation.

6.1.4. Sensitivity to the Model Parameters.

Recall that the approximation scheme from Section 5.1 is governed by the parameter α, which balances the updates in slot assignment and the updates in passenger accommodations. To evaluate the impact of α on PSAM-Pax performance, Table 6 reports the percent-wise increase in passenger connecting time as a function of the schedule displacement and the parameter α, when the number of lost passengers is fixed to the value from the flight-centric solution (first column of Table 4). The observations are twofold. First, the solutions are not very sensitive to α for values between 0.3 and 0.6, justifying our default choice of $\alpha = 0.3$. Second, solution quality decreases as α approaches the extreme values (i.e., zero and one). This validates our algorithmic choice of employing a weight-based approximation, as opposed to a coordinate-descent scheme that alternates between fixing each term of the bilinear expression. In between, the best solutions are achieved with an intermediate value of α, which balances the two terms in the approximation scheme (Equation (25)). Similar results are obtained in the other regions of the trade space.

6.2. Full Problem Instance at Changi and Lisbon

We now solve the full problem, with all declared capacities in place at the Singapore Changi Airport. As compared with the restricted problem, the full problem cannot be solved by using off-the-shelf methods and, thus, requires the LNS algorithm. To establish the robustness and generalizability of our findings, we conduct a second case study from the Lisbon Airport. We obtained slot-request data from 2015; because we only have access to demand data from 2016 onward, we reconstructed, albeit approximately, passenger demand in 2015 using 2016 data. This approach relies on the stability of passenger demand from one year to another, which is justified at a well-established airport like Lisbon. Table 7 reports the results of both case studies, using a similar nomenclature as Table 4. Each cell of the table was computed after 10 hours of computation.

The results confirm the main insight from Table 4—namely, that moderate adjustments in slot-allocation decisions can have significant impacts on passenger accommodations. For example, with a 1% increase in displacement, connecting times can be reduced by up to 48% in Changi and 36% in Lisbon. The total number of lost passengers can also be reduced considerably: With a 2% increase in schedule displacement, we obtain a solution that yields a 50% reduction of lost passengers, along with a 13%–33% reduction in added connecting times. Perhaps more surprisingly, the table underscores that significant benefits are now achieved without increasing the schedule displacement (19% reduction in added connecting times in Changi and 7% in Lisbon), thus demonstrating the even stronger impact of PSAM-Pax in complex slot-allocation instances.

These results also underscore the robustness of our insights across the two case studies, based on two airports that play very different roles in flight networks. In addition, they also indicate that the benefits of our passenger-centric slot-allocation approach—namely, the percent-wise reductions in lost passengers and connecting times for the same increase in schedule displacement—are stronger in Changi than in Lisbon. That is, capturing passenger-level considerations is even more important at major connecting hubs (like

Table 6. Increase in Passenger Connecting Time as a Function of α (Columns) and Schedule-Displacement Increase (Rows), for a 0% Decrease in Lost Passengers

	α value										
Displacement	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

Note. The largest improvements are highlighted in bold.

Table 7. Results for Full Problems at SIN and LIS: Added Connecting Times (in Minutes), as a Function of Schedule Displacement (Rows) and the Number of Lost Passengers vs. the Flight-Centric Baseline (Columns)

	Singapore Airport (SIN)					Lisbon Airport (LIS)				
		Lost passengers		CPU (mins.)			Lost passengers		CPU (mins.)	
		5,697	2,849				4,119	2,060		
Model	Displacement	+0%	−50%	Min	Max	Displacement	+0%	−50%	Min	Max
PSAM	188,850	2,496,110	—	600	600	248,835	652,200	—	600	600
	+0%	+0%				+0%	+0%			
PSAM-Pax	188,850	2,027,400	—	600	600	248,835	606,925	—	600	600
	+0%	−18.78%				+0%	−6.94%			
	189,795	1,472,325	—	600	600	250,080	451,590	—	600	600
	+0.5%	−41.02%				+0.5%	−30.76%			
	190,740	1,286,790	2,243,355	600	600	251,325	419,910	—	600	600
	+1%	−48.45%	−10.13%			+1%	−35.62%	—		
	192,625	1,214,330	1,673,790	600	600	253,810	384,825	567,231	600	600
	+2%	−51.35%	−32.94%			+2%	−41.00%	−13.03%		
	200,180	981,695	1,445,295	600	600	263,765	332,220	487,475	600	600
	+6%	−60.67%	−42.10%			+6%	−49.06%	−25.26%		
	207,735	928,315	1,180,230	600	600	273,720	299,870	447,390	600	600
	+10%	−62.81%	−52.72%			+10%	−54.02%	−31.40%		
CPU (min)	Min	600	600			Min	600	600		
	Max	600	600			Max	600	600		

Changi, ranked ninth in the world in terms of connectivity by the OAG Megahubs International Index) than at less connected airports (like Lisbon).

Let us conclude with the computational performance of the large-scale neighborhood search algorithm for the passenger-centric slot-allocation problem considered in this paper. Recall that the algorithm relies on a probabilistic destroy-and-repair approach. Accordingly, we implement the algorithm five times for each test instance. Table 8 reports the solution over time. Whereas the algorithm exhibits strong variability initially, performance stabilizes after 10 hours of computation. In fact, the final solutions vary by less than a percent upon termination of the algorithm, for

Table 8. Computational Performance of LNS Algorithm: Average Increase in Connecting Time and Range Thereof Over Five Runs of the Algorithm, After 1 Hour, 3 Hours, 6 Hours, and 10 Hours of Computations

		Computational time											
		1 hour			3 hours			6 hours			10 hours		
Instance	Displacement (%)	Avg. Inc. Con. Time (%)	% Range		Avg Inc. Con.Time (%)	% Range		Avg Inc. Con.Time (%)	% Range		Avg Inc. Con.Time (%)	% Range	
SIN	0	−0.5	0	−0.9	−4.7	−2.3	7.2	13.7	−12.5	−15.4	−18.8	−18.8	−18.8
(0%)	1	−4.3	−0.3	−11.3	−21.6	−18.2	−33.2	−41.4	−38.8	−43.7	−45.4	−44.6	−48.5
	2	−11.9	−1.8	−17.3	−26.3	−18.4	−34.7	−44.6	−39.0	−47.7	−51.2	−50.6	−51.4
	6	−18.8	−6.2	−22.5	−28.3	−23.9	−38.8	−49.6	−46.6	−53.6	−60.6	−60.3	−60.7
	10	−19.0	−11.2	−24.5	−38.8	−33.2	−47.2	−55.0	−46.8	−57.7	−62.6	−61.9	−62.8
SIN	1	4.9	+4.9	+4.9	1.2	+2.2	+0.2	−9.6	−9.5	−9.7	−9.6	−9.6	−9.7
(−50%)	2	0.1	+3.8	−1.8	−7.3	+0.4	−11.7	−13.0	−2.5	−26.7	−32.9	−32.9	−32.9
	6	−2.4	−1.2	−6.4	−11.4	−6.4	−21.8	−17.2	−9.6	−33.6	−42.0	−41.8	−42.1
	10	−7.8	−2.5	−14.5	−15.5	−8.9	−27.2	−28.9	−19.8	−37.7	−52.1	−51.5	−52.7
LIS	0	0	0	0	−4.7	−2.1	−6.5	−6.5	−5.8	−6.9	−6.9	−6.9	−6.9
(0%)	1	−1.4	−0.1	−5.8	−23.8	−16.8	−29.7	−33.9	−33.5	−35.1	−35.6	−35.6	−35.6
	2	−3.5	−2.5	−7.1	−27.1	−21.0	−33.6	−39.1	−38.5	−39.7	−41.0	−40.8	−41.0
	6	−4.7	−3.1	−11.0	−29.9	−25.8	−34.7	−44.1	−42.1	−48.7	−48.9	−48.7	−49.1
	10	−7.9	−6.2	−14.6	−33.0	−29.0	−37.0	−48.1	−45.4	−49.0	−53.7	−53.4	−54.0
LIS	2	+4.0	+6.3	+2.1	+1.7	+4.5	−1.2	−10.9	−10.4	−11.5	−13.0	−13.0	−13.0
(−50%)	6	+0.6	+1.8	−0.3	−4.5	−0.1	−8.4	−20.5	−18.9	−22.3	−25.2	−25.2	−25.3
	10	−1.4	−0.3	−2.4	−14.5	−7.1	−20.9	−24.7	−24.2	−25.4	−30.8	−30.6	−31.4

Notes. The instance is characterized by the airport (Lisbon (LIS) vs. SIN) and the number of lost passengers (0% vs. 50%). Avg. Inc. Con. Time, average increase in connecting time.

both airports. Ultimately, these results underscore the robustness of the algorithm, which returns consistently high-quality solutions to a very large-scale and complex problem, in reasonable times—therefore providing strong potential benefits in practice toward passenger-centric slot allocation.

7. Conclusion

This paper proposes an original approach to airport slot allocation that balances flight-level objectives (minimizing schedule displacement) and passenger-level objectives (maximizing available itineraries and minimizing connecting times). To this end, we developed an integrated approach that combines predictive analytics—to infer passenger flows across flight networks—and prescriptive analytics—to optimize passenger-centric slot allocation. To solve the resulting mixed-integer bilinear program, we employed a weight-based approximation scheme embedded into a large-scale neighborhood search. Despite the size of the slot-allocation problem and the added complexity due to passenger considerations, the proposed algorithm returns high-quality solutions in real-world instances at the largest slot-coordinated airports worldwide.

Results at the Singapore Changi Airport uncovered a highly nonlinear relationship between schedule displacement, the number of feasible itineraries, and connecting times. As a result, considerable improvements to passenger-level metrics can be achieved at a small cost in terms of flight displacement. For instance, we reduced the number of lost passengers by 50% and reduced connecting-time increases by 33% by increasing the schedule displacement by only 2%, as compared with a flight-centric baseline. Similar results were obtained at the Lisbon Airport, therefore establishing the generalizability and robustness of our findings. The modeling and algorithmic approach developed in this paper can thus make slot-allocation outcomes much more consistent with passenger demand at a small cost in terms of flight displacement, resulting ultimately in stronger itinerary opportunities for passengers and stronger revenue opportunities for airlines.

These encouraging results motivate further research toward passenger-centric transportation. First, this paper assumes a static model of passenger demand without considering the impact of flight scheduling on overall passenger demand and its distribution across itineraries. Future research can address this limitation by expanding the modeling framework to incorporate supply-demand interactions. Second, the algorithm relies on an approximation scheme and a large-scale neighborhood search heuristic to handle the mixed-integer nonconvex programming formulation. An interesting research opportunity consists of augmenting this solution approach with exact methods. Third, whereas this paper focuses on a single-airport problem, future research can expand the modeling and computational approach to multiple airports in order to capture interactions across schedule-coordinated airports. Finally, the COVID-19 pandemic has upended the air-travel industry, opening critical research opportunities toward designing new demand-management policies in the postpandemic world. This paper provides methodological foundations to explore these questions and underscores important opportunities to design inclusive transportation systems that bring together perspectives of infrastructure managers, service providers, and travelers.

Acknowledgment

Any opinions, findings, and conclusions expressed in this paper are those of the author(s) and do not reflect the views of the Civil Aviation Authority of Singapore.

References

Adler N, Hashai N (2005) Effect of open skies in the Middle East region. *Transportation Res. Part A Policy Practice* 39(10):878–894.

Adler N, Njoya ET, Volta N (2018) The multi-airline p-hub median problem applied to the African aviation market. *Transportation Res. Part A Policy Practice* 107:187–202.

Atasoy B, Salani M, Bierlaire M (2014) An integrated airline scheduling, fleeting, and pricing model for a monopolized market: An integrated airline scheduling, fleeting, and pricing model. *Comput.-Aided Civil Infrastructure Engrg.* 29(2):76–90.

Ball M, Donohue G, Hoffman K (2006) Auctions for the safe, efficient and equitable allocation of airspace system resources. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 507–538.

Ball MO, Estes AS, Hansen M, Liu Y (2020) Quantity-contingent auctions and allocation of airport slots. *Transportation. Sci.* 54(4): 858–881.

Barnhart C, Fearing D, Vaze V (2014) Modeling passenger travel and delays in the national air transportation system. *Oper. Res.* 62(3):580–601.

Birolini S, Cattaneo M, Malighetti P, Morlotti C (2020) Integrated origin-based demand modeling for air transportation. *Transportation Res. Part E Logistics Transportation Rev.* 142:102050.

Birolini S, Jacquillat A, Cattaneo M, Antunes AP (2021a) Airline network planning: Mixed-integer non-convex optimization with demand-supply interactions. *Transportation Res. Part B Methodological* 154:100–124.

Birolini S, Antunes AP, Cattaneo M, Malighetti P, Paleari S (2021b) Integrated flight scheduling and fleet assignment with improved supply-demand interactions. *Transportation Res. Part B Methodological* 149:162–180.

Boonekamp T, Zuidberg J, Burghouwt G (2018) Determinants of air travel demand: The role of low-cost carriers, ethnic links and aviation-dependent employment. *Transportation Res. Part A Policy Practice* 112:18–28.

Bratu S, Barnhart C (2006) Flight operations recovery: New approaches considering passenger recovery. *J. Scheduling* 9(3):279–298.

Breiman L (2001) Random forests. *Machine Learning* 45(1):5–32.

Brueckner JK (2002) Airport congestion when carriers have market power. *Amer. Econom. Rev.* 92(5):1357–1375.

Cadarso L, Vaze V, Barnhart C, Marín Á (2017) Integrated airline scheduling: Considering competition effects and the entry of the high speed rail. *Transportation Sci.* 51(1):132–154.

Carlin A, Park RE (1970) Marginal cost pricing of airport runway capacity. *Amer. Econom. Rev.* 60(3):310–319.

Daniel JI (1995) Congestion pricing and capacity of large hub airports: A bottleneck model with stochastic queues. *Econometrica* 63(2):327–370.

Dong Z, Chuhang Y, Lau HH (2016) An integrated flight scheduling and fleet assignment method based on a discrete choice model. *Comput. Indust. Engrg.* 98:195–210.

Fairbrother J, Zografos KG (2021) Optimal scheduling of slots with season segmentation. *Eur. J. Oper. Res.* 291(3):961–982.

Grosche T, Rothlauf F, Heinzl A (2007) Gravity models for airline passenger volume estimation. *J. Air Transport Management* 13(4): 175–183.

Hakim MM, Merkert R (2016) The causal relationship between air transport and economic growth: Empirical evidence from South Asia. *J. Transportation Geogr.* 56:120–127.

Hoerl AE, Kennard RW (1970) Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 42(1):80–86.

IATA (2014) *Standard Schedules Information Manual (SSIM)* (International Air Transport Association, Montreal).

IATA (2020) *Worldwide Airport Slot Guidelines* (International Air Transport Association, Montreal).

Jacquillat A (2020) Predictive and Prescriptive Analytics toward Passenger-centric Ground Delay Programs. Preprint, submitted November 20, https://dx.doi.org/10.2139/ssrn.3734008.

Jacquillat A, Vaze V (2018) Interairline equity in airport scheduling interventions. *Transportation Sci.* 52(4):941–964.

Jiang Y, Zografos KG (2021) A decision making framework for incorporating fairness in allocating slots at capacity-constrained airports. *Transportation Res. Part C Emerging Tech.* 126:103039.

Jin F, Li Y, Sun S, Li H (2020) Forecasting air passenger demand with a new hybrid ensemble approach. *J. Air Transportation Management* 83:101744.

Jorge D, Ribeiro NA, Antunes AP (2021) Towards a decision-support tool for airport slot allocation: Application to Guarulhos (Sao Paulo, Brazil). *J. Air Transportation Management* 93:102048.

Katsigiannis FA, Zografos KG (2021) Optimising airport slot allocation considering flight-scheduling flexibility and total airport capacity constraints. *Transportation Res. Part B Methodological* 146:50–87.

Katsigiannis FA, Zografos KG, Fairbrother J (2021) Modelling and solving the airport slot-scheduling problem with multi-objective, multi-level considerations. *Transportation Res. Part C Emerging Tech.* 124:102914.

Lhéritier A, Bocamazo M, Delahaye T, Acuna-Agost R (2019) Airline itinerary choice modeling using machine learning. *J. Choice Model.* 31:198–209.

Lohatepanont M, Barnhart C (2004) Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Sci.* 38(1):19–32.

Lundberg S, Lee SI (2017) A unified approach to interpreting model predictions. von Luxburg U, Guyon I, Bengio S, Wallach H, Fergus R, eds. *NIPS'17 Proc. 31st Internat. Conf. Neural Inform. Processing Systems* (Curran Associates, Red Hook, NY), 4768–4777.

Lurkin V, Garrow LA, Higgins MJ, Newman JP, Schyns M (2017) Accounting for price endogeneity in airline itinerary choice models: An application to continental US markets. *Transportation Res. Part A Policy Practice* 100:228–246.

Marazzo M, Scherre R, Fernandes E (2010) Air transport demand and economic growth in Brazil: A time series analysis. *Transportation Res. Part E Logist. Transportation Rev.* 46(2):261–269.

Marla L, Vaaben B, Barnhart C (2017) Integrated disruption management and flight planning to trade off delays and fuel burn. *Transportation Sci.* 51(1):88–111.

Paleari S, Redondi R, Malighetti P (2010) A comparative study of airport connectivity in China, Europe and US: Which network provides the best service to passengers? *Transportation Res. Part E Logistics Transportation Rev.* 2:198–210.

Phyoe S, Guo R, Zhong Z (2016) An air traffic forecasting study and simulation. *Internat. J. Sci. Tech.* 2(3):55–69.

Pita JP, Barnhart C, Antunes AP (2013) Integrated flight scheduling and fleet assignment under airport congestion. *Transportation Sci.* 47(4):477–492.

Rassenti SJ, Smith VL, Bulfin RL (1982) A combinatorial auction mechanism for airport time slot allocation. *Bell J. Econom.* 13(2):402–417.

Ribeiro NA, Jacquillat A, Antunes AP (2019) A large-scale neighborhood search approach to airport slot allocation. *Transportation Sci.* 53(6):1772–1797.

Ribeiro NA, Jacquillat A, Antunes AP, Odoni AR, Pita JP (2018) An optimization approach for airport slot allocation under IATA guidelines. *Transportation Res. Part B Methodological* 112:132–156.

Sherali HD, Bae KH, Haouari M (2010) Integrated airline schedule design and fleet assignment: Polyhedral analysis and benders' decomposition approach. *INFORMS J. Comput.* 22(4):500–513.

Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B Methodological* 58(1):267–288.

Tsui WHK, Balli HO, Gilbey A, Gow H (2014) Forecasting of Hong Kong Airport's passenger throughput. *Tourism Management* 42: 62–76.

Wei W, Hansen M (2006) An aggregate demand model for air passenger traffic in the hub-and-spoke network. *Transportation Res. Part A Policy Practice* 40(10):841–851.

Wei K, Vaze V, Jacquillat A (2020) Airline timetable development and fleet assignment incorporating passenger choice. *Transportation Sci.* 139–163.

Zografos KG, Jiang Y (2019) A bi-objective efficiency-fairness model for scheduling slots at congested airports. *Transportation Res. Part C Emerging Tech.* 102:336–350.

Zografos KG, Androutsopoulos KN, Madas MA (2018) Minding the gap: Optimizing airport schedule displacement and acceptability. *Transportation Res. Part A Policy Practice* 114:203–221.

Zografos KG, Madas MA, Androutsopoulos KN (2017) Increasing airport capacity utilisation through optimum slot scheduling: Review of current developments and identification of future needs. *J. Scheduling* 20(1):3–24.

Zografos KG, Salouras Y, Madas MA (2012) Dealing with the efficient allocation of scarce resources at congested airports. *Transportation Res. Part C Emerging Tech.* 21(1):244–256.