Where did the time go? On the increase in airline schedule padding over 21 years

Dennis J. Zhang
Olin Business School, Washington University in St. Louis, denniszhang@wustl.edu
Yuval Salant
Kellogg School of Management, Northwestern University, y-salant@northwestern.edu
Jan A. Van Mieghem
Kellogg School of Management, Northwestern University, vanmieghem@northwestern.edu

August 27, 2018

About 43 million US domestic flights operated on routes that were serviced consistently by the same airlines from 1997 to 2017. The scheduled duration of these flights, as posted on computerized reservation systems, increased on average by 8.1% over the 21-year span. Where did the time go? Building on Deshpande and Arıkan (2012), we develop a multi-period newsvendor model of how airlines decide their posted duration and show that the model can be reduced to a series of single-period newsvendor problems. The model explains more than 99% of the variation among the $43,062,886$ posted flight durations in our data. We use structural estimation and counterfactual analysis to establish that more than 45% of the increase in posted duration stems from airlines strategically padding their schedule to achieve higher on-time performance. We also provide evidence that decreased airline competition on a route is associated with increased strategic padding.

Key words: Time, empirical newsvendor, structural estimation, airline operations.

1. Introduction

Forty three million US domestic flights operated on routes that were serviced consistently by the same airlines from 1997 to 2017. These flights constitute about one-third of all US domestic flights over the same time span. The scheduled duration of these flights, as posted on computerized reservation systems, increased by 8.1% from 1997 to 2017. This 8.1% increase reflects a cumulative increase of 3.449 million hours in posted flight duration, or 341 million passenger hours assuming an average flight occupancy of 99 passengers.[1] Where did the time go? In other words, what drives the increasingly longer posted flight duration in the U.S. domestic airline industry?

This paper studies the drivers of this increase in posted flight duration. There are three takeaways from the analysis. First, more than 50% of the increase in posted duration stems from airlines

[1] The Bureau of Transportation Statistics (BTS) reports that there were 965 million passengers on the 9.7 million US domestic flights in 2017 implying an average occupancy of 99 passengers per-flight. (https://www.bts.gov/newsroom/2017-traffic-data-us-airlines-and-foreign-airlines-us-flights accessed on August 2018.)

increasing their schedule padding rather than from increased ground- or air-time. Second, more than 95% of the increase in padding is due to airlines strategically padding their schedule to achieve higher on-time performance. Put differently, the increase in posted duration stems to a large extent from changes in airlines' strategic decision making and not from physical constraints. Third, the amount of strategic padding is negatively associated with the intensity of competition on a route as measured by the number of airlines servicing that route.

Airlines typically schedule their flights and determine the posted flight duration several months ahead of the actual flight date. When deciding their posted flight durations, airlines may use the average historical flight duration on the same route as a baseline and add additional time on top of this baseline to improve their on-time performance, a practice called *schedule padding*.[2] Schedule padding may in turn be affected by both how much variance there is in the historical flight duration distribution, which we call *variability padding*, and how conservative the airline is in its targeted on-time performance, which we refer to as *strategic padding*. If the variance of historical flight duration increases, the airline needs to add more time on top of the baseline to ensure the same on-time performance. If the airline becomes more conservative, then it needs to add more time on top of the baseline to achieve higher on-time probability, even if the distribution of historical flight duration does not change.

In order to understand the contribution of historical flight duration, variability padding and strategic padding to the increase of posted flight duration, we consider a multi-period model in which airlines optimize posted flight duration to maximize their long-run average profit. The main trade-off in the model is between the overage and underage costs of posted flight duration. On one hand, if an airline posts a longer flight duration for a particular route, the airline may incur higher overage costs due to consumers' opportunity costs and the airline's duration-dependent cost (e.g., crew wages). On the other hand, if an airline posts a shorter flight duration, it may experience consumer backlash due to arrival delays which may affect the airline's future demand. We demonstrate that despite the multi-period nature of the model, its solution follows a classic single-period newsvendor logic: airlines post the flight duration corresponding to the critical fractile of the distribution of historical flight duration to balance underage and overage costs.

We use flight-level data from the Bureau of Transportation Statistics (BTS) to estimate the model parameters. The data includes about 43 million flights operated on routes that were serviced consistently by the same airlines from 1997 to 2017. These flights constitute about one-third of all US domestic flights in this time period. For each flight, we observe several characteristics of

[2] Some practitioners use schedule padding to refer to padding on top of the minimum (instead of the average) flight duration on a route. We refer to schedule padding as padding on top of the average flight duration on a route since it is more consistent with the newsvendor model.

its time-line including its posted departure and arrival times as well as its actual departure and arrival times.

The model performs well in describing airlines' posted flight duration. The in-sample R-square of the model is about 99.3% even though we use only two continuous parameters per-route, with each parameter taking about 1000 values, to explain about two million posted durations in any given year.

We use structural estimation and counterfactual analysis to decompose the 8.1% increase in posted flight duration from 1997 to 2017 into its four possible drivers according to the model: an increase in historical air-time, an increase in historical ground-time, an increase in variability padding, and an increase in strategic padding. We establish that the two main drivers of the increase in posted duration are the increase in strategic padding and the increase in ground time with each of them contributing roughly 45% to the increase in posted duration. Air-time and variability padding neither significantly increased, nor affected the increase of posted flight duration. Thus, a main force behind the increased posted flight duration is a change in airlines' strategic decision making: airlines aim to improve their on-time performance by increasing their posted duration rather than streamlining their operations.

A possible driver of changes in strategic padding is competition on a route. The number of airlines servicing a given route may influence airlines' strategic padding decision. If consumers care more about posted duration (i.e., how long the flight is) than about the likelihood that the flight arrives on-time (i.e., on-time performance), an increase in the number of airlines servicing a route may lead to a decrease in strategic padding. In contrast, if consumers care more about the likelihood of on-time arrival, then an increase in the number of airlines servicing a route may lead to an increase in strategic padding. We utilize an event-study framework to analyze how the entry or exit of an airline on a route affects other airlines' posted flight duration on the same route. We estimate that the entry of an airline is associated with a decrease in posted duration of other airlines on the same route, whereas the exit of an airline is associated with an increase in the posted flight duration of other airlines on the same route. These findings hint at the possibility that merger activity from 1997 to 2017 is a relevant driver of the increased strategic padding.

2. Literature Review

Our paper follows a literature that structurally estimates parameters in classic operations models using data. Olivares et al. (2008) started this literature by structurally estimating the overage and underage costs in a newsvendor model from operating room scheduling data. Subsequently, Deshpande and Arıkan (2012) structurally estimate a newsvendor model using airlines schedule data from 2005 to 2007 and conclude that the posted flight duration of the majority of airlines is longer than average historical duration. Arıkan et al. (2013) extend this analysis to consider multi-leg flights and identify the bottleneck airports in the US air-travel infrastructure. Akşin et al. (2013) and Yu et al. (2016) extend the structural estimation methods to queuing models where they model customers' decisions to join or leave a queue as a single-agent dynamic model, following Rust (1987). Bray et al. (2017) develop a dynamic discrete choice model to estimate the level of ration gaming and bullwhip effect using transaction-level data from a supermarket supply chain.

Our paper is mostly related to Olivares et al. (2008), Deshpande and Arıkan (2012) and Arıkan et al. (2013) since we also structurally estimate a newsvendor model. We contribute to this literature in four ways. First, we extend the single-period newsvendor model in Olivares et al. (2008) and Deshpande and Arıkan (2012) to a multi-period newsvendor model that links an airline's dynamic scheduling decisions with consumer demand. We show that these yearly scheduling decisions, despite their dynamic nature, still have closed-form newsvendor solutions. Second, we analyze how the structural parameters of airlines' scheduling decisions evolve over 21 years. Third, we decompose the increase in airlines' padding, which is responsible for the majority of the increase in posted duration, into an increase in variability padding and an increase in strategic padding. We use counterfactual analysis to show that more than 96% of the increased schedule padding over 21 years is due to increased strategic padding. Finally, we provide suggestive evidence on the relevance of decreased airline competition on a route to the increase in strategic padding.

Since we study the inter-temporal changes of airlines' posted flight duration, our work also relates to the large literature in economics and operations management focusing on airline scheduling and delays (Shumsky 1995, Hebert and Dietz 1997, Mueller and Chatterji 2002, Mayer and Sinai 2003a,b, Deshpande and Arıkan 2012, Deshpande et al. 2018). Shumsky (1995) is one of the early papers that proposes a forecasting algorithm to generate accurate forecasts of gate delays. Mayer and Sinai (2003b) demonstrate that a significant amount of travel time can be explained by network benefits due to hubbing since airlines try to cluster their flights in short time spans to provide passengers as many potential connections as possible with minimum waiting time. Based on airline time traffic data from 1988 to 2000, Mayer and Sinai (2003a) show that most airlines choose a posted flight duration that is very close to the minimum allowed under federal regulations to minimize wage costs. Based on the same data source but spanning 2005 to 2007, Deshpande and Arıkan (2012), however, demonstrate that the posted duration of more than 66% of all US domestic flights is longer than their average historical flight duration.

We contribute to this literature by providing a dynamic model of how an airline decides its posted flight duration and a corresponding structural estimation procedure. Moreover, our empirical analysis spans 1997 to 2017 to focus on the inter-temporal changes of airlines' posted flight duration decisions. We demonstrate that airlines post increasingly longer flight duration while the actual air-time in 2017 is virtually unchanged relative to 1997. This result reconciles and extends the findings from Mayer and Sinai (2003a) and Deshpande and Arıkan (2012): before 2000, most flight durations were not padded relative to historical durations; in 2007, about two thirds (66%) of durations were padded; and in 2017, most (i.e., 88%) durations were padded.

Last, we contribute to a growing literature in empirical operations management focusing on the airline industry (Ramdas and Williams 2006, Deshpande and Arıkan 2012, Ramdas et al. 2013, Nicolae et al. 2016, Cui et al. 2018). Ramdas and Williams (2006) is the first paper that empirically studies the trade-off between aircrafts' capacity utilization and their on-time performance. Nicolae et al. (2016) uses publicly available data to demonstrate that the implementation of checked bag fees from 2007 to 2009 significantly increases the on-time departure performance of airlines and has a positive spillover on airlines serving the same route. Cui et al. (2018) uses proprietary data to study the impact of introducing additional cabin classes on airlines' price distribution and revenue.

3. Data

This section first provides background on airlines' schedule development process. It then describes our data source and how we group flights in the analysis. It concludes with providing preliminary evidence about the evolution of posted flight duration from 1997 to 2017.

3.1. Background

Airlines' schedule development process consists of four stages: (1) service planning, (2) schedule generation, (3) resource allocation, and (4) execution scheduling. The service planning phase involves forming a service plan that specifies the frequency of flights offered in each market and the desired time window (e.g., 5pm–6pm) and aircraft type (e.g., wide or narrow body) for each flight. In the schedule generation stage, this service plan is transformed to an actual passenger schedule taking into account constraints such as the total number of available aircraft and flight crews. The resulting passenger schedule includes the exact departure and arrival times of each flight. The resource allocation stage involves various resource allocation decisions such as assigning aircraft with specific tail numbers to appropriate aircraft rotations taking into account constraints such as maintenance and other operational requirements. Finally, the execution scheduling stage involves implementing the developed schedule by taking exceptions into account.

In the United States, airlines are required to report their passenger schedule, including scheduled departure and arrival times, to the Bureau of Transportation Statistics (BTS) and the Federal Aviation Administration (FAA) several months prior to the departure date. After a flight lands, airlines report several additional statistics to the BTS and the FAA which we describe next.

Figure 1 The BTS data contains the scheduled and actual time-line of each flight.

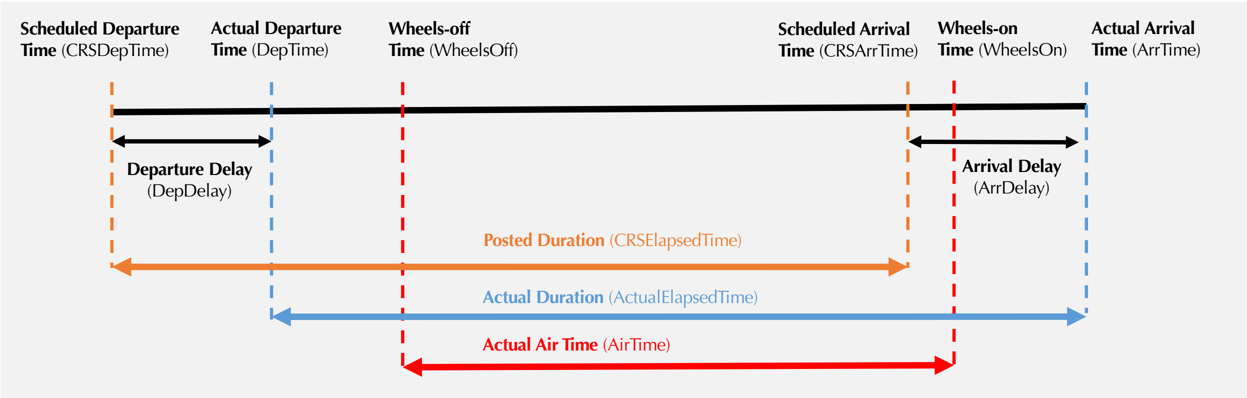

3.2. Data Source

We obtain flight data from the Airline On-Time Performance data set, which is downloadable from the BTS website (henceforth, the BTS data). The BTS data includes more than 129 million domestic flights between 1995 and 2017, with 385 origin airports, 388 destination airports, 9,334 routes (origin-destination pairs), and 26 airlines.

For every flight i, we use the following variables from the BTS data in the analysis (see Figure 1 and Table 1 for a summary of these and other variables):

- CRSDepTime_i and CRSArrTime_i are the Computerized Reservation System (CRS)'s scheduled departure (push-off from the departure gate) and arrival (push-in to the arrival gate) times of flight i. Airlines report these times to the BTS several months prior to flight i's departure date.
- DepTime_i and ArrTime_i are the realized departure and arrival times of flight i, where departure and arrival are defined as in the previous bullet point.
- WheelsOff_i and WheelsOn_i correspond to when flight i's wheels leave the ground during takeoff and when flight i's wheels touch the ground during landing, respectively.
- CRSElapsedTime_i is the *posted duration* of flight i:

$$CRSElapsedTime_i = CRSArrTime_i - CRSDepTime_i,$$

which is the posted gate-to-gate duration that passengers observe on computerized reservation systems.

- $\text{ActualElapsedTime}_i$ is the actual gate-to-gate duration of flight i:

$$ActualElapsedTime_i = ArrTime_i - DepTime_i.$$

- AirTime_i is the time between takeoff and landing:

$$AirTime_i = WheelsOn_i - WheelsOff_i.$$

Table 1 Variables contained in the dataset

Flight Characteristics		**Flight Timeline**	
Variable Name	*Variable Description*	*Variable Name*	*Variable Description*
FlightDate	Flight date	CRSDepTime	Scheduled departure time
DayofMonth	Day of month	ArrTime	Actual arrival gate-push-in time
DayOfWeek	Day of week	DepTime	Actual gate-push-off time
AirlineID	An identification number to identify a unique airline	ArrTime	Actual arrival gate-push-in time
OriginAirportID	Origin airport identification number	WheelsOff	Actual departure wheels-off time
DestAirportID	Destination airport identification number	WheelsOn	Actual arrival wheels-on time
Distance	Distance between airports	CRSElapsedTime	Scheduled gate-to-gate time
		ActualElapsedTime	Actual gate-to-gate time
		AirTime	Difference between actual wheels-on and wheels-off time

From these variables, we compute two additional variables relevant to our analysis—DepGroundTime_i and ArrGroundTime_i—which are, respectively, the ground time of flight i from pushing-off the gate until taking-off during departure and from landing until pushing-in to the gate during arrival:

$$DepGroundTime_i = WheelsOff_i - DepTime_i$$

$$ArrGroundTime_i = ArrTime_i - WheelsOn_i.$$

3.3. Persistent ODAs

When studying the inter-temporal evolution of posted flight durations, we first assign the flights in the data to cohorts, called persistent ODAs. We then analyze the evolution of posted durations within each persistent ODA, and aggregate the results across the ODAs.

An *ODA* stands for Origin-Destination-Airline, and includes all flights with the same origin and destination airports operated by the same airline. Fixing the origin and destination airports is important because distance and airport identities determine to a large extent the air- and ground-time of any flight. Fixing the airline is important because the airline decides how to map air- and ground-time to posted duration.

A *persistent ODA* is an ODA that has at least one flight per-year from 1995 to 2017. Focusing on persistent ODAs enables us to track the evolution of posted duration from 1997 to 2017. The data from 1995 and 1996 is used to compute the historical distributions of actual flight duration that airlines may consider in determining their posted duration in 1997.

When two airlines merge, we have to decide which ODAs will be considered persistent. For example, United Airlines (UA) and Continental Airlines (CO) started to report jointly under UA in January 2012. Before the merger, there were two ODAs corresponding to flights from Los

Figure 2 **Annual number of US domestic flights (total and from persistent ODAs) and number of US operating airlines in the data from 1995 to 2017**

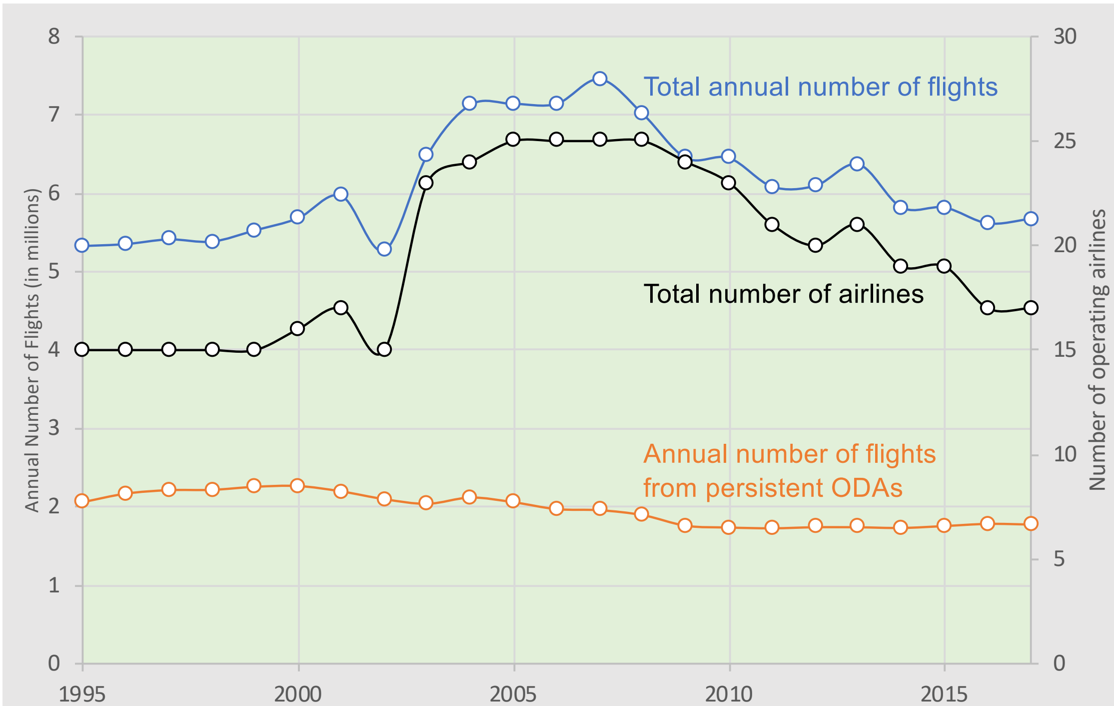

Notes. This figure presents the change in the number of airlines and flights from 1997 to 2017 for all ODAs and persistent ODAs.

Angeles (LAX) to San Francisco (SFO), denoted by LAX-SFO-UA and LAX-SFO-CO. After the merger, the data has only one ODA denoted by LAX-SFO-UA. In principle, either LAX-SFO-CO, LAX-SFO-UA, or both could be considered as persistent ODAs in this case because it is unclear which scheduling system is used post-merger. The analysis treats the post-merger ODA (LAX-SFO-UA in the example) as persistent and the other one (LAX-SFO-CO) as not. Our results remain qualitatively similar if we consider both ODAs as persistent.

3.4. Initial evidence

There are $1,038$ persistent ODAs in the data, and they come from five airlines: American Airlines, Alaska Airlines, Delta Airlines, Southwest Airlines and United Airlines. Persistent ODAs account for $43,062,886$ million flights, or about one-third of the flights in the data.

Figure 2 illustrates that there is a minor inter-temporal decrease in number of flights in persistent ODAs, while the total number of flights fluctuates across years. The figure also illustrates that the fluctuation in the total number of flights co-moves with the number of reporting airlines in the data. For example, there was a change in the BTS reporting rules in the early 2000s that led to an increase in the number of reporting airlines in the data, which in turn led to an increase in the total number of flights around that time. In contrast, merger activity leads to a reduction in

Table 2 **Summary Statistics for the 1,038 Persistent ODAs (All times reported in minutes)**

Statistic	Mean	St. Dev.	Min	Max
Panel A: Year 1997				
Number of Flights	2,136	1,638	40	13,522
Average Air Time	120.629	72.481	11.298	447.979
Average Posted Duration	142.459	77.141	18.689	496.417
Average Actual Duration	140.852	75.868	20.317	474.837
Average Departure Delay	10.007	3.800	0.640	33.900
Average Arrival Delay	11.653	3.909	2.745	33.500
Average Departure Ground Time	14.466	4.779	5.820	36.019
Average Arrival Ground Time	5.765	2.523	1.957	14.525
Average Distance (miles)	902.623	627.550	31	3,784
Panel B: Year 2017				
Number of Flights	1,710	1,182	31	6,729
Average Air Time	121.916	74.295	10.740	476.904
Average Posted Duration	150.588	79.061	22.416	515.138
Average Actual Duration	145.240	77.584	23.237	507.812
Average Departure Delay	11.263	4.286	2.540	37.784
Average Arrival Delay	11.238	4.384	2.364	38.081
Average Departure Ground Time	15.894	4.224	7.282	34.280
Average Arrival Ground Time	7.439	3.166	2.803	22.586
Average Distance (miles)	902.606	627.592	31.000	3,784.000

Figure 3 **Posted flight duration has increased more than actual flight duration from 1997 to 2017**

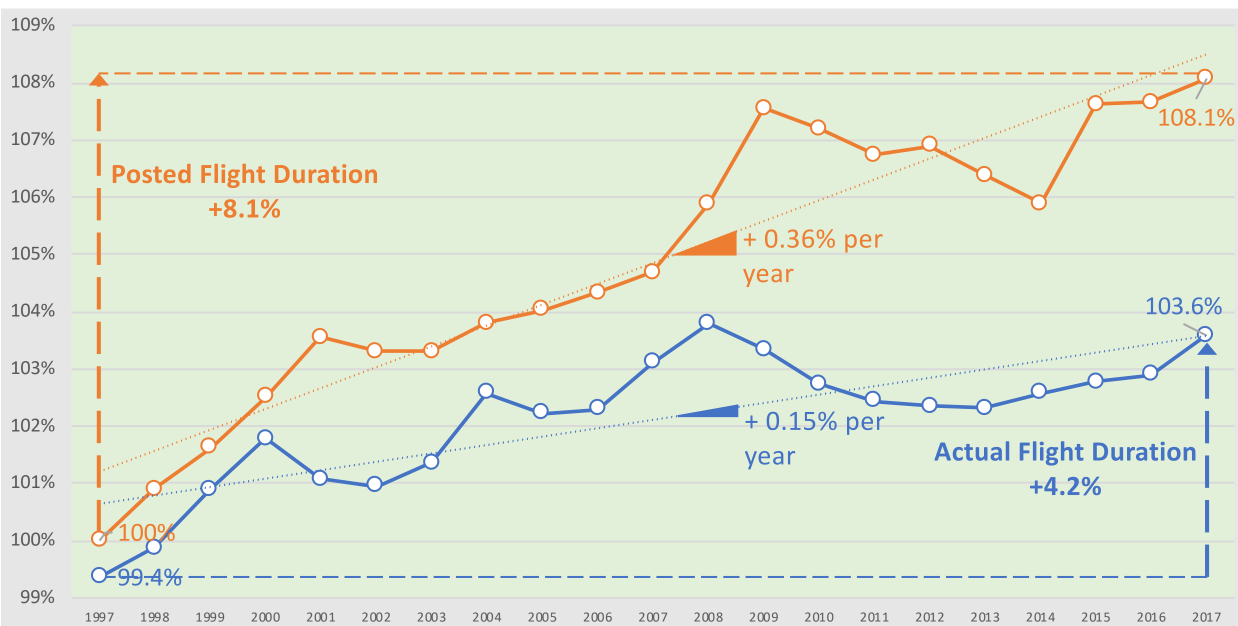

Notes. This figure presents the average posted flight duration and average actual flight duration from 1997 and 2017 across all flights belonging to persistent ODAs, as a percentage of that ODA's average posted flight duration in 1997.

the number of reporting airlines and often to a reduction in the number of flights for the merged airline.

Table 2 compares several statistics of the 1,038 persistent ODAs between 1997 and 2017. An average persistent ODA has $2,136$ flights per year in 1997 relative to $1,710$ in 2017. On average, actual flight duration has increased by less than five minutes from 141 minutes in 1997 to 145 minutes in 2017, whereas posted flight duration has increased by more than 8 minutes from 142 minutes to 151 minutes. To give a concrete example of the changes in duration, consider the ODA ABQ-ARL-DL, i.e., the Delta route from Albuquerque to Atlanta. In 1997, this ODA had 942 flights, average posted flight duration was 167 minutes, and average actual flight duration was 170 minutes. In 2017, the number of flights in this ODA went down to 865, actual flight duration increased by 4 minutes to 174 minutes, and the posted duration increased by 16 minutes to 183 minutes.

Figure 3 provides a more detailed description of the evolution of the average posted duration and the average actual duration from 1997 to 2017 across persistent ODAs, as a fraction of the relevant ODA's posted duration in 1997. For each year and each type of duration, the duration of each flight is first normalized by the ODA's average posted duration in 1997. The figure then reports the average of these normalized flight durations. For example, suppose there was only one ODA with average posted duration in 1997 of 120 minutes, and it had two flights in 2008 with actual durations 125 and 127 minutes. Then, the normalized flight durations are $\frac{125}{120} = 104.2\%$ and $\frac{127}{120} = 105.8\%$, and the figure would report the average $(104.2 + 105.8)/2 = 105\%$.

Two patterns emerge from Figure 3. First, posted flight duration has increased by 8.1% from 1997 to 2017. Second, this increase cannot be explained only by the increase in actual flight duration given that the increase in actual flight duration is about half of the increase in posted duration.

To appreciate the magnitude of the 8.1% increase in posted duration, Figure 4 plots the increase in posted duration, measured in hours, of all flights in a given year from 1998 to 2017. Summing over all years, the cumulative increase in hours of posted duration is 3.449 million hours. Because the BTS data does not provide the actual number of passengers on each flight, we cannot translate this number directly to total passenger hours. The BTS, however, reports that in 2017, there were 965 million passengers on 9.7 million US domestic flights representing an average occupancy of 99 passengers per domestic flight. Thus, as a rough approximation, the 3.449 million hours translate to $99 \times 3.449 \approx 341$ million passenger hours. Recall that this analysis considers only persistent ODAs, which represent one-third of all domestic flights, so that the 341 million hours estimate is probably conservative.

4. Model and Validation

To identify the drivers of the increase in posted flight duration, we consider a dynamic newsvendor model, and estimate its structural parameters. The model and the estimation build on the contributions of Olivares et al. (2008) and Deshpande and Arıkan (2012).

Figure 4 **The 8.1% increase in posted flight duration represents a cumulative increase of 3.449 million hours of posted flight duration from 1997 to 2017**

Notes. This figure reports the annual increase in posted duration, relative to the posted duration in 1997, of all flights in a given year belonging to persistent ODAs.

4.1. Model

There are N periods. In the beginning of period 1, an airline decides on the posted durations $T_1, \ldots, T_N$ of one given flight for the next N periods taking into account the effect of the posted duration on consumer demand for the flight and the airline's operations costs.

Consumer demand. There is a continuum of consumers and their per-period valuations for the flight are distributed uniformly on the $[0,1]$ interval. The utility of a consumer with valuation θ from taking the flight in period $n > 1$ is:

$$u_n(\theta, T_n, T_{n-1}, t_{n-1}, p_n) = \theta - c_n T_n - l_n (t_{n-1} - T_{n-1})^+ - p_n, \tag{1}$$

where c_n is the consumer's opportunity cost at period n, t_{n-1} is the actual flight duration in prior period $n-1$, l_n reflects the aversion of consumers to delays in period n, and p_n is the flight's price.

Thus, the posted duration influences consumer utility through two channels. First, consumers experience a disutility c_n from any additional unit (i.e. minute) of posted duration at period n. This is because consumers plan their schedule around the posted duration and a longer duration implies less time for other activities. Second, consumers experience a disutility l_n from any unit of delay relative to the posted duration in period $n-1$. This disutility aims to capture the idea that consumers take the past performance of the airline into account in their decision making.[3]

In the first period $n=1$, customers do not have information on prior airline performance so that their utility becomes $u_1 = \theta - c_1 T_1 - p_1$.

Consumers purchase a flight ticket if their utility is larger than their outside option, which is assumed to be 0. Thus, the demand for the flight in period n is:

$$D_n(T_n, T_{n-1}, t_{n-1}, p_n) = \int \mathbb{1}\{u_n(\theta, T_n, T_{n-1}, t_{n-1}, p_n) > 0\} d\theta = 1 - c_n T_n - l_n(t_{n-1} - T_{n-1})^+ - p_n.$$

Operations costs. In addition to consumer demand, the posted flight duration T_n also affects airlines' operations cost $C_n(T_n)$ in period n. We model this cost as follows:

$$C_n(T_n) = c_n^o T_n + l_n^o (t_n - T_n)^+,$$

where c_n^o is the operations cost of another minute of posted duration in period n, which includes for example pay to crew, and l_n^o is the per-minute operations cost of flight delay, which includes for example rescheduling costs.

The Airline's Optimization Problem. Based on the demand and operations cost, the airline chooses a posted flight duration vector $\vec{T} = (T_1, T_2, ..., T_N)$ and a price vector $\vec{p} = (p_1, p_2, ..., p_N)$ to optimize its expected N-periods average profit:

$$\pi(\vec{T}, \vec{p}) = \mathbb{E}_{\vec{t}}[\frac{1}{N} \sum_{n=1}^{N} D_n(T_n, T_{n-1}, t_{n-1}, p_n) p_n - C_n(T_n)]$$

where $\vec{t} = (t_1, t_2, ..., t_N)$ is the vector of actual flight durations.[4]

We assume that the airline does not jointly optimize posted durations and prices. Rather, it optimizes durations with respect to some distribution of prices with expected price p, and then, at a later time, optimizes prices. This seems to be the case in practice. Airlines schedule their flights and report the schedule to the BTS and the FAA several months prior to the flight date, while they change their prices dynamically until the day of the flight based on various factors such as seat availability.

Thus, the airline wishes to optimize:

$$\max_{\vec{T}} \pi(\vec{T}, p). \tag{2}$$

We denote a solution to this optimization problem by $\vec{T}^* = (T_1^*, \ldots, T_N^*)$.

[3] There are, of course, other ways to measure the past performance of the airline. For example, consumers may use several past periods to evaluate performance.

[4] The profit specification assumes that the marginal cost of serving a consumer is 0.

Optimal Posted Duration. Even though the optimization problem in (2) is an N-period dynamic programming problem, it can be decomposed into N single-period newsvendor problems by "shifting" the consumers' past experience expression, $l_n(t_{n-1} - T_{n-1})^+$, from period n to period $n-1$. This is the content of the following lemma.

LEMMA 1. *[Decomposition] For any period $n < N$, the optimal posted duration in this period, T_n^*, solves the following newsvendor problem:*

$$\min_{T_n} \mathbb{E}_{t_n}[(c_n + \frac{c_n^o}{p})T_n + (l_{n+1} + \frac{l_n^o}{p})(t_n - T_n)^+].$$

We can now use standard newsvendor logic to find the optimal posted duration T_n^* in every period n as a function of the period-dependent cost parameters and probability distribution:

PROPOSITION 1. *[Optimal Posted Duration] The optimal posted flight duration T_n^* in period $n < N$ solves:*

$$\Pr(t_n < T_n^*) = \min\{0, \frac{l_{n+1} + \frac{l_n^o}{p} - c_n - \frac{c_n^o}{p}}{l_{n+1} + \frac{l_n^o}{p}}\}. \tag{3}$$

The left-hand-side of equation (3) specifies the targeted on-time probability, which we will denote by f_n^* and refer to as the *targeted on-time performance* because it corresponds to the fraction of flights that are on time. Note that this targeted on-time performance has a simple closed-form expression in terms of the cost parameters and the expected price. In addition, when the actual flight duration t_n is normally distributed with mean μ_n and variance σ_n^2 the optimal posted duration also has a closed-form solution:

$$T_n^* = \mu_n + z_n^* \sigma_n \text{ where } z_n^* = \Phi^{-1}(f_n^*) \tag{4}$$

where $\Phi(\cdot)$ denotes the CDF of the standard normal distribution.

Proposition 1 demonstrates that it is optimal for an airline to pad its schedule when it perceives the distribution of actual flight duration as normal and the targeted on-time performance is greater than 50%. That is, the optimal posted duration in this case augments the mean flight duration μ_n by an "optimal padding" term $z_n^* \sigma_n$.

The padding term itself can be decomposed into (1) *variability padding* due to the unpredictability of actual flight duration, which is measured by the standard deviation σ_n of the actual flight duration, and (2) *strategic padding* due to the airline's targeted on-time performance f_n^*. Targeted on-time performance is a strategic objective of the airline, hence the name strategic padding. The targeted on-time performance is optimally connected to the cost parameters of consumers and the airline. In what follows, we will say that an airline is more "conservative" when it increases its targeted on-time performance.

4.2. Estimation

Assuming airlines perceive the flight duration as normally distributed, we can use Equation (4) to recover the targeted on-time performance by: (1) estimating the mean and standard deviation of the flight duration in an ODA from historical data, and then (2) using the estimated mean and standard deviation, and the posted duration to recover the targeted on-time performance.[5]

Specifically, fixing a persistent ODA j and year t, let $\bar{\mu}_{jt}$ denote the average historical flight duration of this ODA in years $t-1$ and $t-2$:

$$\bar{\mu}_{jt} = \frac{1}{N_{jt-2} + N_{jt-1}} \left(\sum_{i=1}^{N_{jt-1}} \text{ActualElapsedTime}_{ijt-1} + \sum_{i=1}^{N_{jt-2}} \text{ActualElapsedTime}_{ijt-2} \right),$$

where N_{jt} denotes the number of flights in ODA j at year t, and $\text{ActualElapsedTime}_{ijt}$ is the actual flight duration of flight i in ODA j at year t.

Similarly, let $\bar{\sigma}_{jt}$ denote the standard deviation of historical flight duration in ODA j in year t:

$$\bar{\sigma}_{jt} = \sqrt{\frac{1}{N_{jt-2} + N_{jt-1}} \left[\sum_{i=1}^{N_{jt-1}} (\text{ActualElapsedTime}_{ijt-1} - \bar{\mu}_{jt})^2 + \sum_{i=1}^{N_{jt-2}} (\text{ActualElapsedTime}_{ijt-2} - \bar{\mu}_{jt})^2 \right]}.$$

Given the values of $\bar{\mu}_{jt}$ and $\bar{\sigma}_{jt}$, we estimate the targeted on-time performance of flight i in ODA j at year t, $\hat{f}_{ijt}$, by using the formula:

$$\hat{f}_{ijt} = \Phi\left(\frac{\text{CRSElapsedTime}_{ijt} - \bar{\mu}_{jt}}{\bar{\sigma}_{jt}}\right), \tag{5}$$

where $\text{CRSElapsedTime}_{ijt}$ is the posted flight duration of flight i in ODA j at year t.

4.3. Model Validation

Before proceeding to study the drivers for the increase in posted duration over time, we examine how well our model, and in particular Equation (4), fits the data. Specifically, we regress posted flight duration on the estimated mean and standard deviation for the relevant ODA, and examine the in-sample R-square.

Consider the following linear regression:

$$\text{CRSElapsedTime}_{ijt} = \beta_t^0 \bar{\mu}_{jt} + \beta_t^1 \bar{\sigma}_{jt} + \epsilon_{ijt}.$$

According to the model, the estimated β_t^0 should be around 1, and β_t^1 should correspond to the estimated strategic padding term. Note that this regression model is restricted to use the same strategic padding term for all ODAs in a given year.

[5] It is also possible to estimate f_n^* without making assumptions on the distribution of actual flight duration by using the empirical distribution of historical actual flight duration. Our main results do not change qualitatively if we use this alternative estimation approach.

Table 3 In-sample Validation of the Newsvendor Model

Panel A: Model Validation without Fixed Effects

	β^0	β^1	Num of Observations	R-square	Adjusted R-square
1997	1.018^{***}	-0.055^{***}	2209635	0.993	0.993
1998	1.017^{***}	-0.016^{***}	2215525	0.993	0.993
1999	1.017^{***}	-0.01^{***}	2250152	0.993	0.993
2000	1.019^{***}	-0.007^{***}	2266218	0.993	0.993
2001	1.018^{***}	-0.023^{***}	2198063	0.993	0.993
2002	1.015^{***}	-0.054^{***}	2098559	0.993	0.993
2003	1.015^{***}	-0.002^{***}	2039793	0.993	0.993
2004	1.016^{***}	0.038^{***}	2119071	0.993	0.993
2005	1.022^{***}	-0.087^{***}	2059561	0.994	0.994
2006	1.021^{***}	-0.084^{***}	1971466	0.993	0.993
2007	1.019^{***}	-0.031^{***}	1957842	0.993	0.993
2008	1.019^{***}	0.004^{***}	1894371	0.992	0.992
2009	1.019^{***}	0.077^{***}	1759449	0.992	0.992
2010	1.021^{***}	0.058^{***}	1736754	0.993	0.993
2011	1.018^{***}	0.077^{***}	1727082	0.993	0.993
2012	1.018^{***}	0.166^{***}	1740537	0.993	0.993
2013	1.015^{***}	0.205^{***}	1739950	0.993	0.993
2014	1.016^{***}	0.164^{***}	1731232	0.994	0.994
2015	1.014^{***}	0.385^{***}	1755878	0.994	0.994
2016	1.01^{***}	0.386^{***}	1783375	0.994	0.994
2017	1.012^{***}	0.383^{***}	1775095	0.993	0.993
Average	1.017	0.08	1953791	0.993	0.993

Panel B: Model Validation with Fixed Effects

	β^0	β^1	Num of Observations	R-square	Adjusted R-square
1997	1.017^{***}	-0.153^{***}	2209635	0.994	0.994
1998	1.02^{***}	-0.079^{***}	2215525	0.993	0.993
1999	1.019^{***}	-0.098^{***}	2250152	0.994	0.994
2000	1.017^{***}	-0.127^{***}	2266218	0.994	0.994
2001	1.015^{***}	-0.178^{***}	2198063	0.994	0.994
2002	1.013^{***}	0.019^{***}	2098559	0.994	0.994
2003	1.015^{***}	0.084^{***}	2039793	0.994	0.994
2004	1.02^{***}	0.011^{***}	2119071	0.994	0.994
2005	1.019^{***}	-0.072^{***}	2059561	0.995	0.995
2006	1.016^{***}	-0.127^{***}	1971466	0.994	0.994
2007	1.016^{***}	-0.074^{***}	1957842	0.994	0.994
2008	1.013^{***}	0.081^{***}	1894371	0.994	0.994
2009	1.022^{***}	-0.106^{***}	1759449	0.994	0.994
2010	1.018^{***}	-0.035^{***}	1736754	0.994	0.994
2011	1.016^{***}	0.068^{***}	1727082	0.994	0.994
2012	1.013^{***}	0.171^{***}	1740537	0.994	0.994
2013	1.01^{***}	0.169^{***}	1739950	0.995	0.995
2014	1.018^{***}	0.013^{***}	1731232	0.995	0.995
2015	1.018^{***}	0.122^{***}	1755878	0.994	0.994
2016	1.017^{***}	0.143^{***}	1783375	0.995	0.995
2017	1.017^{***}	0.12^{***}	1775095	0.994	0.994
Average	1.016	0.00	1953791	0.994	0.994

Note: $*p < 0.10; **p < 0.05; ***p < 0.01$. Standard errors are robust and clustered at the ODA level.

Note: This table shows the regression results of posted flight duration on historical average and standard deviation of actual flight duration from the same ODA. Panel A does not control for any additional fixed effects. Panel B controls for the airline fixed effect, the origin airport fixed effect, the destination airport fixed effect and the date fixed effect.

Table 3 reports the estimates of β_t^0 and β_t^1 as well as the R-square for every year from 1997 to 2017 for two different specifications: Panel (A) in the table reports the estimates without controlling for any additional co-variates; Panel (B) augments this specification with airline fixed effect, origin and destination airports fixed effects, and date fixed effects.

Three patterns emerge from the table. First, the restricted two-variable specification in Panel (A) fits the data remarkably well: the R-square and the adjusted R-square are between 99.3% and 99.4% despite the fact that we use about $1,000$ different levels of historical averages and standard deviations every year to explain about two million posted durations. Second, adding another 626 variables (i.e, 5 unique airlines, 128 unique origin airports, 128 unique destination airports, and 365 unique dates) in Panel (B) only improves the R-square by 0.1%. Third, the estimate of β_t^0 in both specifications is close to 1 for each year, in conformance with the model. These patterns suggest that the restricted and simple two-variable model in Equation (4) fits the data well.

5. Results

The estimation equation for flight i in ODA j at year t is:

$$\text{CRSElapsedTime}_{ijt} = \bar{\mu}_{jt} + \text{Padding}_{ijt}, \tag{6}$$

where $\bar{\mu}_{jt}$ is the average historical flight duration in ODA j at years $t-1$ and $t-2$ and

$$\text{Padding}_{ijt} = \bar{\sigma}_{jt} \times \Phi^{-1}(\hat{f}_{ijt})$$

is the estimated padding term, which is the product of variability padding (as captured by $\bar{\sigma}_{jt}$, the historical standard deviation in ODA j at years $t-1$ and $t-2$) and strategic padding (as captured by the targeted on-time performance $\hat{f}_{ijt}$ computed using Equation (5)).

Historical flight duration can be further decomposed into historical air- and ground-time. Hence, letting $\bar{\mu}_{jt}^A$ and $\bar{\mu}_{jt}^G$ denote the average historical air- and ground-time in ODA j at years $t-1$ and $t-2$, we can rewrite equation (6) as

$$\text{CRSElapsedTime}_{ijt} = \bar{\mu}_{jt}^A + \bar{\mu}_{jt}^G + \text{Padding}_{ijt}. \tag{7}$$

In this section, we first quantify the contribution of each of the three factors in Equation (7) to the increase in posted duration. We then decompose the increase in padding into the increase in strategic padding and variability padding, and study how each of these sources of padding have evolved over the 21 year data span.

5.1. Padding Relative to Historical Duration

To quantify the contribution across ODAs of historical air-time, historical ground-time, and padding to the increase in posted duration, we transform Equation (7) into an equation about differences normalized by 1997 posted duration at the ODA level.

Specifically, let $\overline{\text{CRSElapsedTime}}_{j1997}$ and $\overline{\text{Padding}}_{j1997}$ denote the average posted duration and the average padding in ODA j at year 1997. Then, by Equation (7):

$$\overline{\text{CRSElapsedTime}}_{j1997} = \bar{\mu}^A_{j1997} + \bar{\mu}^G_{j1997} + \overline{\text{Padding}}_{j1997}. \tag{8}$$

Subtracting equation (8) from equation (7), we obtain the equation about differences:

$$\begin{aligned}
\Delta\text{CRSElapsedTime}_{ijt} &\equiv \text{CRSElapsedTime}_{ijt} - \overline{\text{CRSElapsedTime}}_{j1997} \\
&= (\bar{\mu}^A_{jt} - \bar{\mu}^A_{j1997}) + (\bar{\mu}^G_{jt} - \bar{\mu}^G_{j1997}) + (\text{Padding}_{ijt} - \overline{\text{Padding}}_{j1997}) \\
&= \Delta\bar{\mu}^A_{jt} + \Delta\bar{\mu}^G_{jt} + \Delta\text{Padding}_{ijt}.
\end{aligned}$$

Dividing by $\overline{\text{CRSElapsedTime}}_{j1997}$, we obtain the breakdown of the percentage change of posted duration of flight i relative to 1997 into a percentage change of air-time, ground-time, and padding:

$$\begin{aligned}
\%\Delta\text{CRSElapsedTime}_{ijt} &= \%\Delta\bar{\mu}^A_{jt} + \%\Delta\bar{\mu}^G_{jt} + \%\Delta\text{Padding}_{ijt} \\
&= \frac{\Delta\bar{\mu}^A_{jt} + \Delta\bar{\mu}^G_{jt} + \Delta\text{Padding}_{ijt}}{\overline{\text{CRSElapsedTime}}_{j1997}}
\end{aligned}$$

Figure 5 decomposes the average percentage increase in posted flight duration across all flights in persistent ODAs over a 21-year span into the average percentage increase in historical air-time, historical ground-time, and padding, where changes are evaluated relative to posted duration in 1997. Two patterns emerge from the figure. First, the increase in padding is responsible for slightly more than half of the increase in posted duration in 2017 relative to 1997. Second, the increase in average historical flight duration, which is responsible for slightly less than half of the increase in posted duration, is mostly due to an increase in ground-time but not air-time. Put differently, the two main drivers of the increase in posted duration are increased padding and increased ground-time while airtime has been virtually unchanged.

Padding, as a fraction of posted duration, has increased significantly from 1997 to 2017. To see this, we divide equation (6) by $\text{CRSElapsedTime}_{ijt}$ and re-arrange terms to obtain the percentage contribution of padding to the posted duration of flight i in ODA j at year t:

$$\frac{\text{Padding}_{ijt}}{\text{CRSElapsedTime}_{ijt}} = 1 - \frac{\bar{\mu}_{jt}}{\text{CRSElapsedTime}_{ijt}}. \tag{9}$$

Averaging the contribution of padding over all flights in persistent ODAs in a given year yields Figure 6. The figure illustrates that the contribution of padding to posted duration has more than tripled from 1997 to 2017.

Figure 5 **About half of the increased posted flight duration stems from an increase in padding while less than 10% is due to increased historical airtime $\bar{\mu}^A$.**

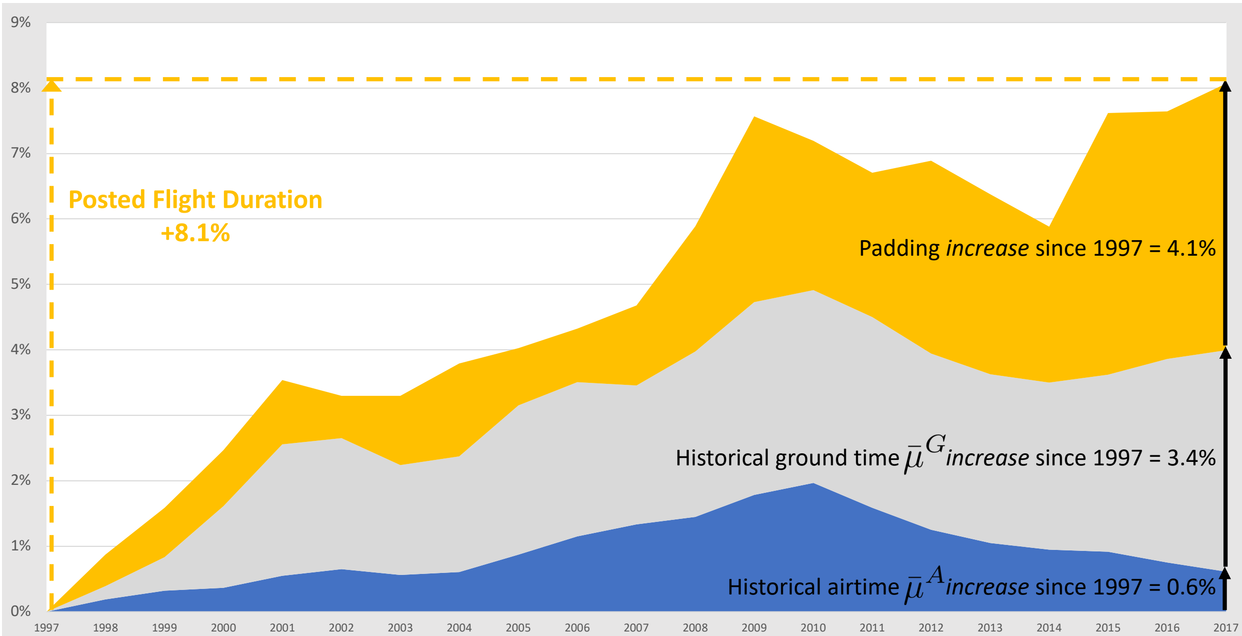

Notes. Using our model, the graph decomposes the increase in average posted flight duration (relative to 1997) into three causes and shows the fraction of the posted duration increase that is due to an increase in (1) historical airtime $\bar{\mu}^A$ (blue); (2) historical ground time $\bar{\mu}^G$ (grey); and (3) padding (yellow).

5.2. Strategic padding relative to variability padding

The significant increase in padding may result from an increase in strategic padding (as captured by the targeted on-time performance $\hat{f}_{ijt}$) or variability padding (as captured by the historical variability $\bar{\sigma}_{jt}$). To quantify the contribution of these two types of padding to the overall increase in padding, we consider two counterfactual scenarios. In the first, we compare the evolution of padding in the data to how padding would have evolved if targeted on-time performance had not changed from 1997 to 2017 and historical variability had evolved according to the data. In the second, we compare the evolution of padding in the data to how padding would have evolved if targeted on-time performance had changed from 1997 to 2017 according to our estimates and historical variability remained constant.

For the first scenario, let $\hat{f}_{j1997}$ denote the average estimated targeted on-time performance across all flights in ODA j in year 1997 and let

$$\Delta\text{Padding}_{ijt} = \text{Padding}_{ijt} - \bar{\sigma}_{j1997} \times \Phi^{-1}(\hat{f}_{j1997})$$

measure the increase in padding of flight i in ODA j at year t relative to the 1997 padding term evaluated with respect to $\bar{\sigma}_{j1997}$ and $\hat{f}_{j1997}$.

Figure 6 Padding, as a fraction of posted flight duration, has tripled over 20 years.

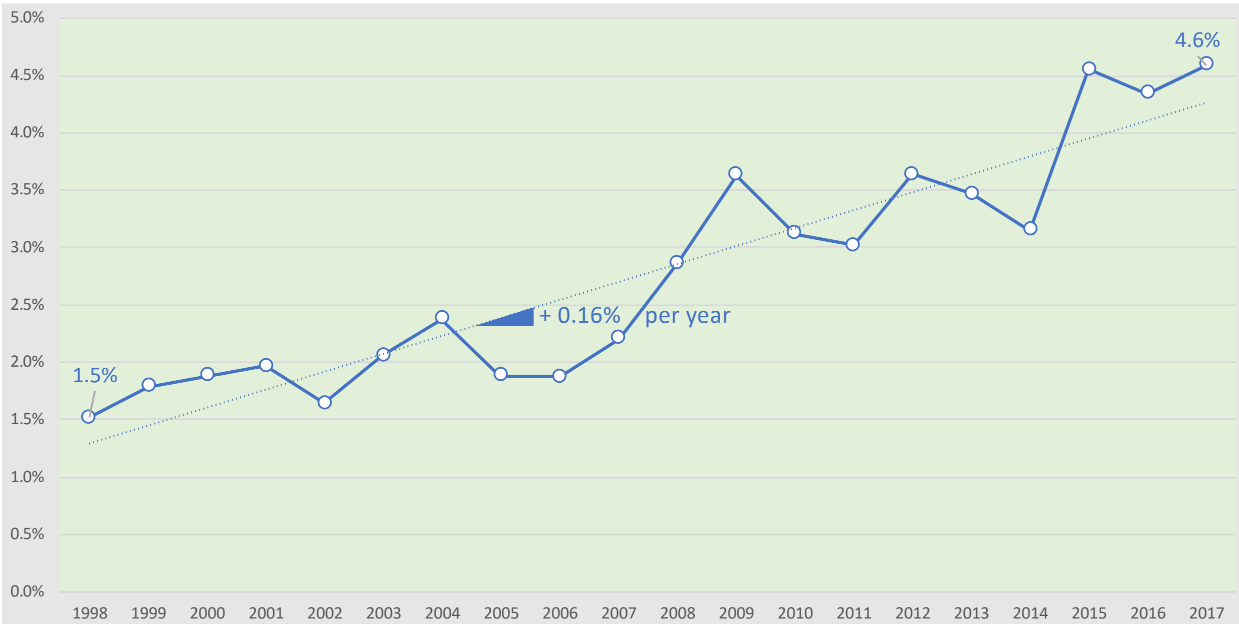

Notes. Using our model, the graph shows the average over all flights i in ODA j of the fraction of Padding_{ijt} relative to the posted duration $\text{CRSElapsedTime}_{ijt}$.

The counterfactual padding $\overline{\text{Padding}}^1_{ijt}$ of flight i in ODA j at year t is:

$$\overline{\text{Padding}}^1_{ijt} = \bar{\sigma}_{jt} \times \Phi^{-1}(\hat{f}_{j1997}), \tag{10}$$

which is the product of the historical variability in year t and the padding term of year 1997. Note that this counterfactual padding does not depend on the flight but only on the flight's ODA.

The percentage contribution of strategic padding to the increase in padding, denoted by $\text{StrategicPaddingPerc}^1_{ijt}$, is the difference between the estimated padding for flight i at year t and the counterfactual padding for flight i at year t divided by $\Delta\text{Padding}_{ijt}$:

$$\text{StrategicPaddingPerc}^1_{ijt} = \frac{\text{Padding}_{ijt} - \overline{\text{Padding}}^1_{ijt}}{\Delta\text{Padding}_{ijt}}.$$

The percentage contribution of variability padding, denoted by $\text{VariabilityPaddingPerc}^1_{ijt}$, is then:

$$\text{VariabilityPaddingPerc}^1_{ijt} = 1 - \text{StrategicPaddingPerc}^1_{ijt}.$$

For the second scenario, the counterfactual padding $\overline{\text{Padding}}^2_{ijt}$ of flight i in ODA j in year t is:

$$\overline{\text{Padding}}^2_{ijt} = \bar{\sigma}_{j1997} \times \Phi^{-1}(\hat{f}_{ijt}),$$

Figure 7 About 90% of the increase in padding is due to an increase in airline conservatism.

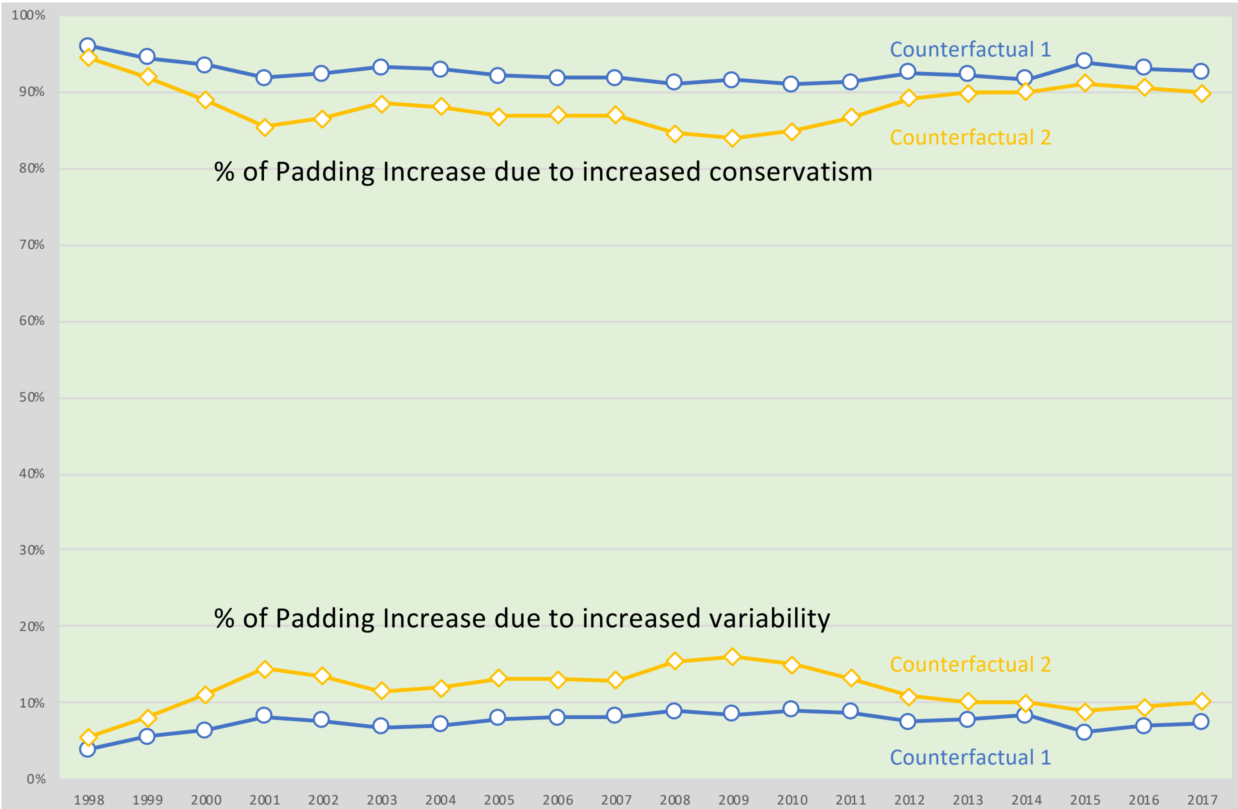

Notes. The graph shows the two counterfactual analyses to estimate the increase in padding due to (1) an increase in travel time variability (i.e., assuming the airline kept its critical fractile constant over 20 years) and (2) an increase in airline conservatism (i.e., an increase in critical fractile assuming flight time variability was constant over 20 years).

which is the product of 1997 historical variability and the estimated targeted on-time performance of flight i at year t. The percentage contribution of variability and strategic padding to the increase in padding are then:

$$\text{VariabilityPaddingPerc}^2_{ijt} = \frac{\text{Padding}_{ijt} - \overline{\text{Padding}}^2_{ijt}}{\Delta\text{Padding}_{ijt}}$$
$$\text{StrategicPaddingPerc}^2_{ijt} = 1 - \text{VariabilityPaddingPerc}^2_{ijt}.$$

Figure 7 presents the results of the two counterfactual scenarios. For each scenario and for each year, we plot the average contribution of strategic and variability padding of all flights in all persistent ODAs. It is evident that in both counterfactual scenarios, the increase in strategic padding is responsible for about 90% of the increase in padding. The estimate of the strategic padding contribution is systematically lower in all years under the second counterfactual scenario. This is because the first scenario uses the historical standard deviation in years $t > 1997$ to evaluate the contribution of strategic padding, while the second scenario uses the standard deviation in 1997, and the historical standard deviation in any year $t > 1997$ is on average slightly larger than that in 1997, which slightly increases the contribution estimate.

6. Strategic Padding and Competition

Section 5 demonstrates that more than half of the increase in posted duration over the 21-year data span is due to an increase in padding, and that most of the increase in padding is due to an increase in strategic padding. Strategic padding in turn decreases—according to our model—in the cost ratio

$$\frac{c_n + \frac{c_n^o}{p}}{l_{n+1} + \frac{l_n^o}{p}}.$$

Thus, the increase in strategic padding is driven by either an increase in the l parameters or a decrease in the c parameters.

Another possible driver of changes in targeted on-time performance that is not captured by model is competition. The number of airlines serving a given route (an origin-destination (OD) pair) may influence airlines' strategic padding decision. If consumers care more about posted duration (how long the flight is) than about the likelihood that the flight arrives on-time (on-time performance), an increase in the number of airlines serving a route may lead to an decrease in strategic padding. In contrast, if consumers care more about the likelihood of on-time arrival, then an increase in the number of airlines serving a route may lead to an increase in strategic padding. This section investigates the association between the intensity of competition on a route and strategic padding.

To do so, we consider routes that experienced an increase or a decrease in the number of airlines operating on that route. We refer to an increase as an (airline) entry and a decrease as an (airline) exit, and examine the association between entry and/or exit in year $t-2$ and airlines' targeted on-time performance in year t on the same route.[6] Specifically, we use the following empirical event-study specification:

$$\hat{f}_{jkt} = \gamma_0 + \gamma_1 \text{Entry}_{k,t-2} + \gamma_2 \text{Exit}_{k,t-2} + \gamma_3 X_{jkt} + \epsilon_{jkt}, \tag{11}$$

where j denotes the ODA; k denotes the route (OD) that corresponds to ODA j; t denotes the year; $\hat{f}_{jkt} = \hat{f}_{jt}$; $\text{Entry}_{kt} = 1$ if route k has an entry in year t, and is zero otherwise; $\text{Exit}_{kt} = 1$ if route k has an exit in year t, and is zero otherwise; and X_{jkt} consists of year, origin airport, destination airport, and airline fixed effects as well as the number of airlines serving the route at year t.

[6] If airline entry and exit decisions are independent of their strategic padding decisions, then our analysis provides causal evidence. This independence, however, is not testable with the data we have.

Table 4 The effect of entry and exit on targeted on-time performance

	Dependent variable:				
	Targeted on-time performance: $\hat{f}_{jkt}$ (%)				
	(1)	(2)	(3)	(4)	(5)
Entry	-4.497^{***}		-3.935^{***}	-3.057^{***}	
	(0.474)		(0.482)	(0.481)	
Exit		3.850^{***}	3.093^{***}	2.728^{***}	5.650^{***}
		(0.488)	(0.496)	(0.494)	(0.725)
Distance	0.010^{***}	0.010^{***}	0.010^{***}	0.010^{***}	0.012^{***}
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.001)
Airline Fixed Effect	Yes	Yes	Yes	Yes	Yes
Origin Fixed Effect	Yes	Yes	Yes	Yes	Yes
Destination Fixed Effect	Yes	Yes	Yes	Yes	Yes
Number of Airlines Fixed Effect	Yes	Yes	Yes	Yes	Yes
Observations	22,800	22,800	22,800	19,484	5,509
Adjusted R^2	0.444	0.443	0.445	0.497	0.373

Note: *p<0.1; **p<0.05; ***p<0.01. Standard errors are robust and clustered at the ODA level. The first two columns consider the effect of only entry (column 1) or only exit (column 2) by one or more airlines on the on-time performance of other airlines operating on the same route. Column (3) considers both entry and exit. Column (4) focuses on ODAs with daily flights in prior two years. Column (5) focuses on ODA-and-year pairs with at least one entry or exit.

We control for airline, origin airport and destination airport fixed effects to eliminate the time homogeneous effects at the airline, origin and destination levels that may be correlated with entry and exit decisions as well as with strategic padding decisions. We control for the number of airlines on a route to eliminate the effect that exit is potentially more likely on more competitive routes while entry is potentially more likely on less competitive routes.

Table 4 reports the results with different control specifications. The first specification estimates entry decisions only (i.e., $\gamma_2 = 0$ in regression equation (11)) and compares routes with entry to routes without entry. Column (1) of Table 4 shows that an entry to a route is associated with a 4.5 percentage point decrease in the average estimated targeted on-time performance of other airlines on that route. The second specification considers only exit decisions (i.e., $\gamma_1 = 0$ in regression equation (11)) and compares routes with exits to routes without exits. Column (2) shows that an exit of airlines on a route is associated with a 3.9 percentage point increase in the average estimated targeted on-time performance of other airlines on that route. The third specification compares routes with entry and/or exit to routes that have neither entry or exit. It shows that an entry (exit) on a route is associated with a 3.9 percentage point decrease (3.1 percentage point increase) in the average estimated targeted on-time performance of other airlines on that route.

Our estimated targeted on-time performance in a given year depends on airlines' historical flight durations in the prior two years. This estimation may have a large sampling error when the number of historical flights is small. Therefore, Column (4) of Table 4 replicates the specification from Column (3) but focuses only on ODAs that have at least daily flights in each of the prior two years. Focusing on these more frequent ODAs, our estimates remain qualitatively the same. The fifth specification restricts attention to routes with entry and routes with exit. Column (5) of Table 4 shows that the average estimated targeted on-time performance on routes with an exit is 5.650 percentage points higher than the on-time performance on routes with an entry. This result is consistent with the combination of estimates in column (4) $(2.728 + 3.057 = 5.785 \approx 5.650)$ and shows that the various specifications provide qualitatively similar results.

So far we coded entry and exit as binary events, while in practice multiple airlines may enter or exit a route in a given year. Indeed, in the data, the largest number of airlines entering a route is 4 and so is the largest number of airlines exiting a route. In order to capture the effect of different number of airlines entering or exiting a route, we use the following specification:

$$\hat{f}_{jkt} = \alpha_0 + \alpha_1 \text{NumAirlineDecrease}_{k,t-2} + \alpha_2 X_{jkt} + \epsilon_{jkt},$$

where $\text{NumAirlineDecrease}_{k,t-2}$ is the decrease in the number of airlines serving route k from Year $t-3$ to Year $t-2$.

Column (1) of Table 5 demonstrates that, without any controls, one less airline on a route is associated with a 1.7 percentage point increase in the average estimated targeted on-time performance of airlines on the same route. Column (2) shows that, controlling for airline, origin and destination fixed effects, the coefficient changes to 3.2 percentage points. Column (3) shows that the effect does not change qualitatively if we only focus on ODAs with daily flights, while Column (4) shows that the effect does not change qualitatively if we only focus on routes with at least one airline entering or exiting.

7. Conclusion

This paper considered four potential drivers of the increase in posted flight duration over a 21-year span: an increase in average air-time, an increase in average ground-time, an increase in variability padding, and an increase in strategic padding. We established that the increase in strategic padding and in average ground time are the two key drivers of the increase in posted duration. Air-time and variability padding neither significantly increased, nor affected the increase of posted flight duration, over the 21-year span.

While the increase in ground time can be attributed to increased airport congestion, it is less clear what drives the increase in strategic padding. We made modest progress in addressing this question by studying the association between competition on a route and targeted on-time performance. We established that an increase (decrease) in the number of airlines on a route is associated with a

Table 5 The effect of the number of airlines entering and exiting on targeted on-time performance

	Dependent variable:			
	Targeted on-time performance: $\hat{f}_{jkt}$ (%)			
	(1)	(2)	(3)	(4)
NumAirlineDecrease	1.728***	3.194***	2.660***	2.644***
	(0.333)	(0.254)	(0.254)	(0.291)
Distance	0.0004	0.011***	0.010***	0.012***
	(0.0003)	(0.0004)	(0.0004)	(0.001)
Airline Fixed Effect	No	Yes	Yes	Yes
Origin Fixed Effect	No	Yes	Yes	Yes
Destination Fixed Effect	No	Yes	Yes	Yes
Number of Airlines	No	Yes	Yes	Yes
Observations	22,800	22,800	19,484	5,509
Adjusted R^2	0.008	0.446	0.498	0.376

Note: *p<0.1; **p<0.05; ***p<0.01. Standard errors are robust and clustered at the ODA level. Column (1) considers the effect of decreasing by one the number of airlines serving a route on the remaining airlines' targeted on-time performance without any fixed effects. Column (2) adds fixed effect. Column (3) focuses on ODAs with daily flights in the prior two years, and Column (4) focuses on ODAs that have at least one entry or exit in a year.

decrease (increase) in targeted on-time performance. This finding hints at the potentially important role of mergers, acquisitions, and bankruptcies in the evolution of strategic padding.

Studying other drivers of the increase in strategic padding is a task for future research. One promising direction is to investigate the effect on strategic padding of the availability of information about airlines' historic on-time performance. The US department of Transportation publishes monthly and yearly reports, entitled "Air Travel Consumer Report" ranking airlines according to their historical on-time performance.[7] These reports are often featured in the popular press. A possible effect of the availability of this information to consumers is that airlines will streamline their operations to decrease the mean and/or variance of actual flight duration, so that they improve their on-time performance. Such a response by airlines is expected to increase consumer welfare. However, the availability of such information may also have unintended consequences: rather than streamlining their operations, airlines may strategically increase their padding to improve their ranking. Such a response by airlines would potentially reduce consumer welfare.

References

Akşin, Zeynep, Barış Ata, Seyed Morteza Emadi, Che-Lin Su. 2013. Structural estimation of callers' delay sensitivity in call centers. *Management Science* **59**(12) 2727–2746.

Arıkan, Mazhar, Vinayak Deshpande, Milind Sohoni. 2013. Building reliable air-travel infrastructure using empirical data and stochastic models of airline networks. *Operations Research* **61**(1) 45–64.

[7] See https://www.transportation.gov/airconsumer/air-travel-consumer-report-archive (accessed on August 22, 2018) for the archive of reports.

Bray, Robert, Yuliang Oliver Yao, Yongrui Duan, Jiazhen Huo. 2017. Ration gaming and the bullwhip effect. *working paper* .

Cui, Yao, A Yesim Orhun, Izak Duenyas. 2018. How price dispersion changes when upgrades are introduced: Theory and empirical evidence from airline industry. *Management Science (forthcoming)* .

Deshpande, Vinayak, Mazhar Arıkan. 2012. The impact of airline flight schedules on flight delays. *Manufacturing & Service Operations Management* **14**(3) 423–440.

Deshpande, Vinayak, Chandrasekhar Manchiraju, Milind Sohoni. 2018. Neither luck nor slack! what drives performance is improving operations: Evidence from the us airline industry. *working paper* .

Hebert, Joseph E, Dennis C Dietz. 1997. Modeling and analysis of an airport departure process. *Journal of Aircraft* **34**(1) 43–47.

Mayer, Chris, Todd Sinai. 2003a. Why do airlines systematically schedule their flights to arrive late? *Working paper, The Wharton School, University of Pennsylvania* .

Mayer, Christopher, Todd Sinai. 2003b. Network effects, congestion externalities, and air traffic delays: Or why not all delays are evil. *American Economic Review* **93**(4) 1194–1215.

Mueller, Eric, Gano Chatterji. 2002. Analysis of aircraft arrival and departure delay characteristics. *AIAA's Aircraft Technology, Integration, and Operations (ATIO) 2002 Technical Forum*. 5866.

Nicolae, Mariana, Mazhar Arıkan, Vinayak Deshpande, Mark Ferguson. 2016. Do bags fly free? an empirical analysis of the operational implications of airline baggage fees. *Management Science* **63**(10) 3187–3206.

Olivares, Marcelo, Christian Terwiesch, Lydia Cassorla. 2008. Structural estimation of the newsvendor model: an application to reserving operating room time. *Management Science* **54**(1) 41–55.

Petruzzi, Nicholas C, Maqbool Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations research* **47**(2) 183–194.

Ramdas, Kamalini, Jonathan Williams. 2006. An empirical investigation into the tradeoffs that impact on-time performance in the airline industry. *Working paper* 1–32.

Ramdas, Kamalini, Jonathan Williams, Marc Lipson. 2013. Can financial markets inform operational improvement efforts? evidence from the airline industry. *Manufacturing & Service Operations Management* **15**(3) 405–422.

Rust, John. 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society* 999–1033.

Shumsky, Robert Arthur. 1995. Dynamic statistical models for the prediction of aircraft take-off times. Ph.D. thesis, Massachusetts Institute of Technology.

Yu, Qiuping, Gad Allon, Achal Bassamboo. 2016. How do delay announcements shape customer behavior? an empirical study. *Management Science* **63**(1) 1–20.

Appendix

A. Proof of Lemma 1

Proof. To prove this lemma, let us write down Period 1 and Period 2's optimization problem in Equation 2:

$$\begin{aligned}
\text{Period } 1 &= E_{t_1}[(\int(\theta - p - c_1T_1)d\theta)p - c_1^oT_1 - l_1^o(t_1-T_1)^+] \\
&= \left(1 - p - (c_1 + \frac{c_1^o}{p})T_1 - \mathbb{E}_{t_1}[\frac{l_1^o}{p}(t_1-T_1)^+]\right)p \\
\text{Period } 2 &= E_{t_2}[(\int(\theta - p - c_2T_2 - l_2(t_1-T_1)^+)d\theta)p - c_2^oT_2 - l_2^o(t_2-T_2)^+] \\
&= \left(1 - p - (c_2 + \frac{c_2^o}{p})T_2 - \mathbb{E}_{t_1}[l_2(t_1-T_1)^+] - \mathbb{E}_{t_2}[\frac{l_2^o}{p}(t_2-T_2)^+]\right)p
\end{aligned}$$

Moving $-\left((c_1+\frac{c_1^o}{p})T_1 + \mathbb{E}_{t_1}[\frac{l_1^o}{p}(t_1-T_1)^+]\right)p$ from Period 1 to Period 2, we obtain

$$\begin{aligned}
&\text{Period } 1 = (1-p)p \\
&\text{Period } 2 = \left(1 - p - (c_1 + \frac{c_1^o}{p})T_1 - \mathbb{E}_{t_1}[(l_2 + \frac{l_1^o}{p})(t_1-T_1)^+]\right)p - \left((c_2+\frac{c_2^o}{p})T_2 + \mathbb{E}_{t_2}[\frac{l_2^o}{p}(t_2-T_2)^+]\right)p
\end{aligned}$$

Similarly, moving $-\left((c_2+\frac{c_2^o}{p})T_2 + \mathbb{E}_{t_2}[\frac{l_2^o}{p}(t_2-T_2)^+]\right)p$ from Period 2 to Period 3 yields:

$$\begin{aligned}
\text{Period } 2 &= \left(1 - p - (c_1 + \frac{c_1^o}{p})T_1 - \mathbb{E}_{t_1}[(l_2 + \frac{l_1^o}{p})(t_1-T_1)^+]\right)p \\
&= \left(1 - p - \mathbb{E}_{t_1}[(c_1 + \frac{c_1^o}{p})T_1 + (l_2 + \frac{l_1^o}{p})(t_1-T_1)^+]\right)p,
\end{aligned}$$

Since T_1 only exists in Period 2 after the transformation, the optimal solution of T_1 solves:

$$\max_{T_1}\left(1 - p - \mathbb{E}_{t_1}[(c_1 + \frac{c_1^o}{p})T_1 + (l_2 + \frac{l_1^o}{p})(t_1-T_1)^+]\right)p,$$

which is equivalent to solving:

$$\min_{T_1}\mathbb{E}_{t_1}[(c_1 + \frac{c_1^o}{p})T_1 + (l_2 + \frac{l_1^o}{p})(t_1-T_1)^+].$$

Following a similar transformation (i.e., moving $-\left((c_n+\frac{c_n^o}{p})T_n - \mathbb{E}_{t_n}[\frac{l_n^o}{p}(t_n-T_n)^+]\right)p$ from Period n to Period $n+1$) and induction, we can show that (a) each T_n only appears at Period $n+1$ for all $n<N$ after the transformation and (b) each T_n solves:

$$\min_{T_n}\mathbb{E}_{t_n}[(c_n + \frac{c_n^o}{p})T_n + (l_{n+1} + \frac{l_n^o}{p})(t_n-T_n)^+].$$

□

B. Proof of Proposition 1

Proof. The optimization in Lemma (1) can be re-written as:

$$\begin{aligned}
&\min_{T_n}\mathbb{E}_{t_n}[(c_n + \frac{c_n^o}{p})T + (l_{n+1} + \frac{l_n^o}{p})(t_n-T)^+] \\
=&\min_{T_n}\mathbb{E}_{t_n}[(c_n + \frac{c_n^o}{p})(T-t_n)^+ + (l_{n+1} + \frac{l_n^o}{p})(t_n-T)^+ + (c_n + \frac{c_n^o}{p})t_n] \\
=&(c_n + \frac{c_n^o}{p})t_n + \min_{T_n}\mathbb{E}_{t_n}[(c_n + \frac{c_n^o}{p})(T-t_n)^+ + (l_{n+1} + \frac{l_n^o}{p})(t_n-T)^+]
\end{aligned}$$

This problem can be viewed as a classic newsvendor problem with overage cost $(c_n + \frac{c_n^o}{p})$ and underage cost $(l_{n+1} + \frac{l_n^o}{p})$. Therefore, adopting the classic newsvendor solution (Petruzzi and Dada 1999), the optimal solution is:

$$f_n^* = \Pr(T_n^* > t_n) = \min\{0, \frac{l_{n+1} + \frac{l_n^o}{p} - c_n - \frac{c_n^o}{p}}{l_{n+1} + \frac{l_n^o}{p}}\}.$$

And, if the flight time t_n is normally distributed with mean μ_n and standard deviation σ_n, the optimal solution has a closed-form solution in terms of the distribution's moments:

$$T_n^* = \mu_n + \sigma_n \Phi^{-1}(f_n^*)$$

□