

# Linking Algorithms to Static and Dynamic pricing: theory and evidence\*

Claudio A. Piga<sup>†</sup>  
University of Genova, IT

Marco Alderighi  
Università della Valle d'Aosta, IT  
Università Bocconi, IT

Alberto A. Gaggero  
University of Pavia, IT

December 18, 2020

## Abstract

We study the main features of a pricing algorithm in the presence of organizational frictions, which impede continuous price adjustments. If prices cannot be continuously and instantaneously updated, and advance sale of fixed capacity is possible, the pricing algorithm optimally assigns a price to each unit in a (non-strictly) monotonically increasing sequence based on the order of sale; prices remain “static” until the algorithm induces dynamic pricing by modifying the sequence. Data from a large European carrier support the assumption that airlines do not adjust fares constantly, as static spells last 35 hours on average but may be longer depending on a flight’s load factor, its selling rate and the time to departure. Dynamic pricing may occur in two forms, with seats being either marked-down or marked-up. The latter tends to occur predominantly during the day, suggesting that human decision-making often complements the algorithm’s operativity.

**JEL Classification:** D22, L11, L93.

**Keywords:** Algorithm; dynamic pricing; competition policy; machine learning.

---

\*We wish to thank James Dana, Tomaso Duso, Philippe Gagnepain, Marc Möller, Alessandro Oliveira, Alessandro Sembenelli, Chris Tsoukis, Hannes Ulrich, Lorenzo Zirulia; the participants of the 2018 HIOC Conference in Indianapolis, the 2018 EARIE Conference in Athens, and the seminar participants at ITEA (online), Oporto University, Bern University, Copenhagen Business School and Berlin ZEW.

<sup>†</sup>Corresponding Author: Department of Economics, University of Genova, Via Vivaldi 5, 16126 Genova, IT. email: claudio.piga@unige.it

# 1 Introduction

The concept of Dynamic Pricing (DP) captures the idea that firms can make instantaneous price offerings based on a set of identifiable features of an online query (e.g., geolocation, device, stock availability, cookie history, etc.). In the representation of DP offered by non-specialized media, algorithms are assumed to play a crucial role as they can promptly process all relevant information, whereas human intervention is generally considered unnecessary and inadequate to the task. Such generalizations may derive from the actual reality of specific markets, such as the nano-seconds trading in financial transactions. There are, however, several reasons why such an operative model may not be directly applicable to other marketplaces, where the pricing algorithm may meet with several potential constraints.

Algorithms are considered to be extremely potent, a perspective that combines well with the economic models where pricing is a frictionless activity that always adjusts prices optimally and at the exact and wanted time (Deneckere and Peck, 2012; Lazarev, 2013; Williams, 2017).<sup>1</sup> However, this view fails to account for at least four complementary perspectives in the literature. First, pricing is a costly marketing activity shaped by a firm's resources, routines and organizational processes (Zbaracki et al., 2004; Bloom and Van Reenen, 2007; Nijs et al., 2014; Ellison et al., 2018). Changing a price may involve the approval of managers from different functional areas with possibly conflicting interests: the ensuing negotiation slows the price changing process down. Second, frictionless pricing is incompatible with the application of often observed uniform pricing; e.g., in DellaVigna and Gentzkow (2019), U.S. retail chains price the same product identically across their outlets, despite they face highly heterogenous demand conditions. Third, price rigidity is found to be prevalent even in online markets where menu costs are negligible (Cavallo, 2017). Lastly, human decision-making, although it takes longer to implement, may turn out to be more effective than an automatized system. Cho et al. (2018) illustrate how the pricing manager of an hotel in an U.S. city often correctly overrode the suggestion made by the computerized system.

This study challenges the notion that firms can instantaneously revise their prices using algorithms to produce tailor-made offerings. The choice of the airline industry as

---

<sup>1</sup>Assad et al. (2020) classify as adopters of Artificial Intelligence pricing agents those petrol stations that exhibit a structural break in the form of a higher frequency of price changes.

the testing ground may appear counter-intuitive, as it is often considered a prime market for the application of algorithmic DP. However, according to the sector-specific press, “... moving into a world of dynamic price offerings has proven technologically difficult for the airline sector, in large part due to the legacy distribution system that was put in place after deregulation ([Travel Weekly, 2018](#))”. That is, historically-determined technological and organizational constraints may hinder how airlines adjust their fares.

In the theoretical part, we derive the optimal pricing algorithm for a firm that has to sell its perishable inventory under demand uncertainty. Our key differentiating assumption is that the firm commits to a predetermined pricing rule for a fixed period of time. The commitment is introduced to capture the firm’s impossibility to adjust prices instantaneously. Another distinguishing feature is that each period is sufficiently long, so that multiple units may be sold. This is not possible in existing models; for instance, [Gershkov and Moldovanu \(2009\)](#) does not allow for the sale of multiple units of goods of the same quality; similarly, in [Gallego and van Ryzin \(1994\)](#) the price jumps after each sale, implying an instantaneous reprogramming of prices. The model focuses on two components of the pricing algorithm: one, a “pricing rule” for all the units of capacity in each period; two, an “adjustment process”, which oversees the revision of the pricing rule’s outcome over subsequent periods.

The pricing rule defines a “price sequence”, where a specific price is set for each unit of capacity; the optimal sequence may be either non-strictly or strictly monotonically increasing in the units’ order of sale, depending on the characteristics of the distribution of customers’ willingness to pay. Overall, we extend the result in [Dana \(1999\)](#), who was the first to describe airline pricing in terms of an equilibrium in distribution, to the multi-period case. Our analysis adds an insurance rationale for the increasing profile of the sequence: if instantaneous adjustments are not possible, then an increasing sequence automatically protects the firm against underselling additional units if a positive shock in demand occurs.

At the beginning of every period, the adjustment process revises the sequence so that, relative to the previous period, some unit prices may either stay the same or change, i.e., decrease or increase. A price decrease relates to the product’s perishable nature, as argued in [McAfee and te Velde \(2007\)](#) and shown in [Sweeting \(2012\)](#). In the empirical part, such decreases in the sequence identify a form of DP denoted as “S-down”. Instead,

the increase of some prices in the sequence, denoted as “S-up”, may be due to either a revision of demand expectations or an inter-temporal discriminatory motive. The latter usually applies closer to the departure date when, for instance, a larger proportion of business-people enters the market (Möller and Watanabe, 2010; Lazarev, 2013; Alderighi et al., 2020; Williams, 2017; Escobari et al., 2019).

In the empirical part of the study, the underlying assumptions and some key results of the theoretical analysis are put to a test or extended using data from the airline industry. First, we assess whether the pricing rule is used in the industry, i.e., whether “fare sequences” characterize airline pricing. Our data, retrieved from the website of easyJet, the second largest European Low-Cost Carrier (LCC), indicate that such sequences are unambiguously used in all the flights at all times.<sup>2</sup> We consider this as support for our result that price sequences provide an optimal strategy when instantaneous price adjustment is not feasible or extremely costly due to, e.g., organizational frictions and technological bottlenecks.<sup>3</sup>

Second, we provide further support to the previous result by measuring the duration of the periods during which the fare sequence remains unchanged. We call such periods “static pricing spells”. To this purpose, we let a web crawler collect data for each flight’s fare sequence continuously over the last fortnight of the booking period, resulting in an average of seven daily observations per flight. A static pricing spell lasts about 35 hours on average, although its duration may extend to over 80 hours. Overall, the evidence is compelling in showing that the airline does not constantly and instantaneously update its fares.

In the remaining empirical analysis, we relax some key assumptions of the theoretical model, which were introduced for the sake of tractability. Periods of fixed duration imply the impossibility, i.e., an infinite cost, to alter prices within the period. Under the more realistic assumption that the cost of price adjustments is finite, we can endogenize the duration of the static pricing spell. Indeed, it is reasonable to assume that airlines retain some degree of control over the management of the pricing rule, which can be changed as long as the expected benefit from doing so exceeds its cost.

Therefore, building on the literature that has emphasized the role of both capacity and

---

<sup>2</sup>The analysis in (Koenigsberg et al., 2008) does not allow for fare sequences.

<sup>3</sup>Alderighi et al. (2020) illustrate that a qualitatively similar pricing rule is used by other large LCCs, namely Ryanair and Southwest, and provides the foundation for the more complex pricing by legacy carriers.

inter-temporal pricing in the industry (Escobari, 2012; Alderighi et al., 2015; Williams, 2017), we hypothesize that the spell duration depends on the number of days to the departure date, the capacity still available on the flight, and the number of seats sold during the spell. The econometric models suggest that the higher the remaining capacity, the shorter the spell. Conversely, more seats sold during the spell lead to longer spells, with the marginal effect of one extra sold seat varying between 6.30 and 6.67 hours.

Furthermore, and relatedly, because in the data the end of a static spell is identified by a DP intervention either of the form S-down or S-up, we investigate whether these interventions are entirely automatic and autonomously managed by the algorithm, or whether they may involve a more direct human discretionary decision. Any evidence of a possible human role would further support the assumption of costly price adjustments. Belobaba (2009, p.90) reports that “*Revenue Management (RM) systems compare actual flight bookings to an expected or ‘threshold’ booking curve for the flight, and then issue ‘exception reports’ to RM analysts whenever a flight’s bookings deviates from the expected booking profile. These exception reports identify flights that might require a RM analyst’s attention*”.<sup>4</sup> As this quote suggests, in the airline industry the RM systems generate predictions to be used as inputs in human decision-making (Agrawal et al., 2019).

To test the relevance of a human intervention to generate DP, we evaluate the probability that a S-up is more likely to occur during the day (and less overnight when RM analysts do not work). The evidence overwhelmingly supports this hypothesis; all else equal, there is on average a 35% lower probability to observe a DP case with fare increases overnight. That is, fare reductions do not seem to require any human assent as they tend to be more driven by the declining option value of a seat and are therefore less discretionary in nature (McAfee and te Velde, 2007; Sweeting, 2012). Instead, fare increases are more often treated as exception reports of strategic importance that need further human approval, as in the case discussed in Cho et al. (2018). This finding suggests that even in an industry with a high reliance on automatized systems, price setting is embedded in organizational processes where the human intervention plays a key role, especially when such strategic decisions as price increases have to be made.<sup>5</sup>

This study contributes to various strands of the economic research. Firstly, with refer-

---

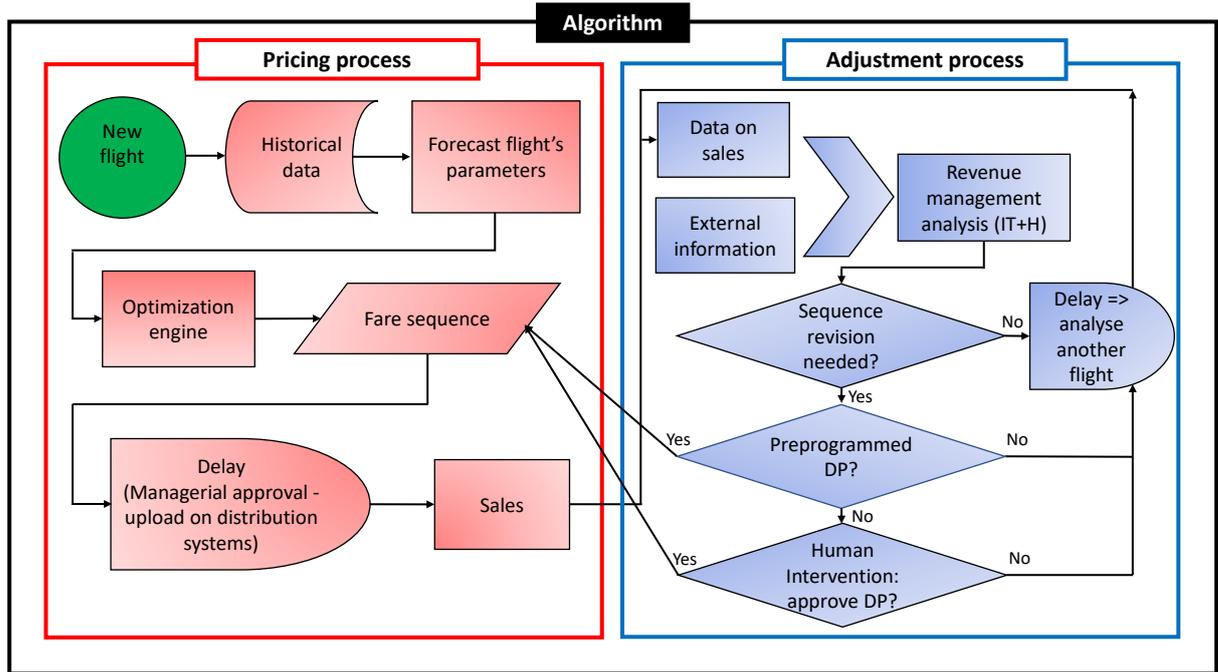
<sup>4</sup>The “booking curve” is defined as the plot of accumulated bookings over time (Swan, 2002).

<sup>5</sup>A fare increase may be more strategically sensitive because it may weaken the firm’s competitive position if its rivals then lower their prices, or whether it increases the risk of lowering the flight’s load factor (Cho et al., 2018).

ence to the recent literature on the use of Artificial Intelligence in organizations (Kleinberg et al., 2018; Agrawal et al., 2019), we analyze a company where the computerized systems are used to improve, as well as complement, human decision making by selecting specific flights for possible intervention. Secondly, by focussing on static pricing spells as an optimal response to technological and managerial frictions, this study fills a gap in the literature, which has exclusively considered continuous and instantaneous fare adjustments. Relatedly, no previous article studies both forms of dynamic pricing (FS-up and FS-down) jointly. On the one hand, a strand of literature predicts a temporally declining trajectory of fares (McAfee and te Velde, 2007), driven by a falling option value of retaining a seat for future sale, as in Sweeting (2012) for the case of secondary market tickets. On the other, fares may go up as the plane fills up (Alderighi et al., 2015; Dana, 1999; Escobari, 2012) or due to a discriminatory intent (Gerardi and Shapiro, 2009; Lazarev, 2013; Williams, 2017). Thirdly, no previous paper relates airline pricing to fare sequences. One implication is that the recent literature that assumes a single fare for all the seats, as in Williams (2017) and Lazarev (2013), needs to be reconciled with the existence of factors hindering instantaneous fares' updates and the ensuing use of sequences in the industry. Finally, whereas the nascent literature on the role of algorithms for competition policy has primarily focussed on theoretical models of collusive pricing (Harrington, 2019; Miklós-Thal and Tucker, 2019; Calvano et al., 2020), in the conclusive remarks we briefly mention some implications of our findings for the analysis of mergers and possible abuse of dominant position.

The organization of the paper matches the components of the pricing algorithm. The next Section depicts the pricing process (schematically illustrated on the left hand side of Figure 1) and lays out the modeling assumptions driving the optimization engine that generates fare sequences. Simulating the model (Section 2.1) highlights the inter-relationship between the pricing and the adjustment processes (schematically shown in left hand side of Figure 1), assuming no human intervention is feasible: it provides insights into how the revision of sequences can be identified in the empirical analysis. The experimental design underlying the data collection strategy is described in Section 3, with Section 4 providing *prima facie* evidence of the extensive use of fare sequences in the industry. The remainder of the paper delves deep into the working of the adjustment process. Section 5 endogenizes the duration of static pricing spell, that is, it investigates whether the flight's fare sequence

Figure 1: Components of the pricing algorithm



Note - IT=Information Technology; H=Human; DP=Dynamic Pricing.

should be revised to reflect the new information gathered during the selling period. For instance, is the flight selling rate adequate? Is the occupancy rate in line with forecasted levels? Has a drastic shock (a pandemic, or sport events such as a qualification for an NBA playoff) occurred? The RM system may choose to do nothing if all is as expected, thus prolonging the static pricing spell. When it intervenes, it may be via a procedure which may be either pre-programmed, i.e., managed by the IT system according to some routine, or mediated by a human decision. In Section 6 we analyze the factors triggering the selection of a particular form of DP, focussing also on which form is more likely to require a RM analyst's intervention.

## 2 A stylized theoretical set-up

In this Section a multi-period economic model for pricing multiple units of capacity under uncertainty is developed based on the following key components of RM. First, the product is perishable. Second, there is uncertainty in consumers' willingness to pay (WTP) and demand. Third, within a booking period, the carrier can sell any number of seats without triggering the start of the next period. This assumption is consistent with the existence of organizational frictions and technological bottlenecks that prevent fares from being constantly manipulated. For instance, DP may be obstructed by the airlines' use of

multiple distribution channels, whose access may not be available continuously. In the U.S., ATPCO is the airline-owned corporation that collects and distributes fare data. According to [Travel Weekly \(2018\)](#), "...airlines assign prices and restrictions to each fare class, then file those classes with ATPCO for dissemination to Global Distribution Systems. At present, carriers are able to update prices in each *fare class four times per day, but only once on weekends*". Other technological bottlenecks may result from the sheer complexity of having to manage, simultaneously, hundreds of thousands of flights over a long time span, covering markets with several competitors. Furthermore, managerial processes take time. According to [Belobaba \(2009\)](#) "*RM systems revise their forecasts and booking limits at regular intervals during the flight booking process, as often as daily in some cases* (p.91)".<sup>6</sup> In other words, the constant reprogramming of the pricing for all these flights would necessitate an unrealistic amount of computing and decision-making power. ([Misra et al., 2019](#)) also assume periods of finite duration during which a machine-learning algorithm is used for simulating demand curves through price experimentation. In their model the monopolist firm faces no capacity constraints and sells non-perishable products, and thus the main objective of the algorithm is to converge in real time to the profit maximizing single price of a very wide range of products.

A monopolistic carrier has to set the fares for a flight with capacity of  $N > 1$  homogeneous seats.<sup>7</sup> The booking process lasts  $T \geq 1$  periods; let  $t$  denote the remaining number of periods before departure, with  $T \geq t \geq 1$ , and  $t = 0$  the departure date. For tractability, each period has a fixed duration, at the start of which the firm commits to a pricing rule. In the empirical part, after assessing that a commitment to a pricing rule corresponds to standard industry practice, we endogenize the period's duration.

At each  $t$ , consumers  $h = 1, 2, \dots, \infty$  arrive sequentially. The probability that the first consumer arrives at  $t$  is  $\varphi_{1,t} \in (0, 1)$ , and that consumer  $h + 1$  arrives conditional on the fact that consumer  $h$  has already appeared is  $\varphi_{h+1,t} \in (0, 1)$ . Consumer  $(h, t)$  is myopic and her willingness to pay is a random variable  $\theta_{h,t}$ , with (right-continuous) cumulative distribution  $F_{h,t}$  on the support  $\Theta$ , with  $\underline{\theta} = \inf \Theta > 0$  and  $\bar{\theta} = \sup \Theta < \infty$ .<sup>8</sup> We assume that the arrival process is memoryless and consumers have the same ex-ante evaluation;

---

<sup>6</sup>The fact that the revision is made "at regular intervals" suggests in itself the impossibility to update fares instantaneously.

<sup>7</sup>Our model shares many features with that developed in [Alderighi et al. \(2020\)](#), which also show that the main implications of the model extend easily to less concentrated market structures.

<sup>8</sup>This guarantees the existence of a solution.

i.e., for any  $h$  and  $t$ ,  $\varphi_{h,t} = \varphi_{h+1,t} = \varphi \in (0, 1)$ . Thus, the probability of selling the first available seat at the fare  $p$  is:

$$q(p) = \varphi(1 - F(p)) \sum_{h=0}^{\infty} (\varphi F(p))^h = \frac{\varphi(1 - F(p))}{1 - \varphi F(p)} \in [0, 1], \quad (1)$$

where  $\varphi(1 - F(p))$  is the probability that consumer  $h$  arrives and buys at fare  $p$  provided that consumers  $1, \dots, h-1$  have previously refused to buy at the same fare; and  $(\varphi F(p))^h$  is the probability that consumers from 1 to  $h$  arrived and did not buy.

Having defined the flight's parameters, the carrier's maximization problem, schematically denoted in Figure 1 as the "optimization engine", is represented by the following Bellman equation:

$$V(t, M) = \max_{p \in \Theta} \{q(p) [p + V(t, M - 1)] + (1 - q(p)) V(t - 1, M)\}, \quad (2)$$

with boundary conditions  $V(t, 0) = 0$  and  $V(0, M) = 0$ , for any  $t \in \{0, \dots, T\}$  and  $M \in \{0, \dots, N\}$ . Moreover, eq. (2) entails a trade-off between selling now at least one seat, gaining  $p$  and the revenue flow coming from the remaining seats,  $V(t, M - 1)$ , and keeping the capacity intact and postpone the sale to the next period, gaining  $V(t - 1, M)$ .

Although other papers have used similar assumptions in terms of customers' heterogeneity and arrival rates, (see, inter alios, [Gallego and van Ryzin \(1994\)](#), [Gershkov and Moldovanu \(2009\)](#) and [Pang et al. \(2015\)](#)), the novel approach in (2), embedded in the part in squared brackets, is to allow for more than one seat to be sold within each  $t$ : this assumption implies the need to set always a (possibly different) fare for all the seats on an aircraft, which remains fixed for the duration of the period. Seats are sold sequentially according to the position of the seat in the distribution. That is, if seat  $m = M, \dots, 2, 1$  is the first in the sequence, then each traveler  $h$  presenting in period  $t$  faces fare  $p(t, m)$ . Within  $t$ , after seat  $m$  is sold, then the next fare on offer becomes  $p(t, m - 1)$ . In the empirical part, we show that this is exactly what the airline does on its website and show evidence that more than one seat is indeed sold without triggering a revision of the pricing rule. At the end of the period  $t$ , the unsold seats are offered in the next period,  $t - 1$ , until  $t = 1$ . Seats available at the end of the last selling period remain unsold.

The model assumes stationarity of parameters; its solution denotes therefore a pricing strategy set "once and for all" at time  $T$ . The main objective of this theoretical set-up

is to gauge whether it can generate prices that closely resemble those that can be found in the airlines’ computer reservation systems. Therefore, we do not delve deep into the general properties of the optimal solution, which are more fully developed in [Alderighi et al. \(2020\)](#). Instead, we simulate the model first by assuming that the distribution of the WTP is invariant over time, i.e.,  $F_{h,t} = F_{h+1,t-1} = F$ . We then relax this assumption, by shifting the weight from the left to the right tail of the distribution in the last three periods, to model the possibility that “business-people” generally enter the market in the later part of the booking process. We suppose  $F(\theta)$ , the distribution of consumers’ WTP, is discrete with its probability density functions reported in Table 1. Such a modeling choice is also dictated by the need to facilitate the comparison between the model’s main results with the data. It is without loss of generality: [Alderighi et al. \(2020\)](#) simulate the model by assuming a continuous, exponential distribution of the WTP and obtain qualitatively similar results.

Table 1: Willingness to pay distributions in the simulation

Time invariant	WTP ( $\theta$ )	50	65	80	95	110	130	150	175	200
	Weight		14/64	12/64	8/64	6/64	6/64	5/64	5/64	4/64
Time variant, last 3 periods	Weight	13/64	11/64	7/64	6/64	6/64	5/64	6/64	5/64	5/64
	Weight	12/64	10/64	6/64	6/64	6/64	5/64	7/64	6/64	6/64
	Weight	11/64	9/64	6/64	6/64	6/64	5/64	7/64	7/64	7/64

In the simulation, we set the number of periods  $T = 11$ ; available seats  $N = 39$ ; and the average total number of prospective travelers  $L = 1.2N$ .<sup>9</sup> Each simulated flight exhibits a possibly different selling path across periods; the number of seats to be removed after each period is obtained by deducting the mode of sold seats in each period after 50.000 runs.

## 2.1 Simulation outcomes

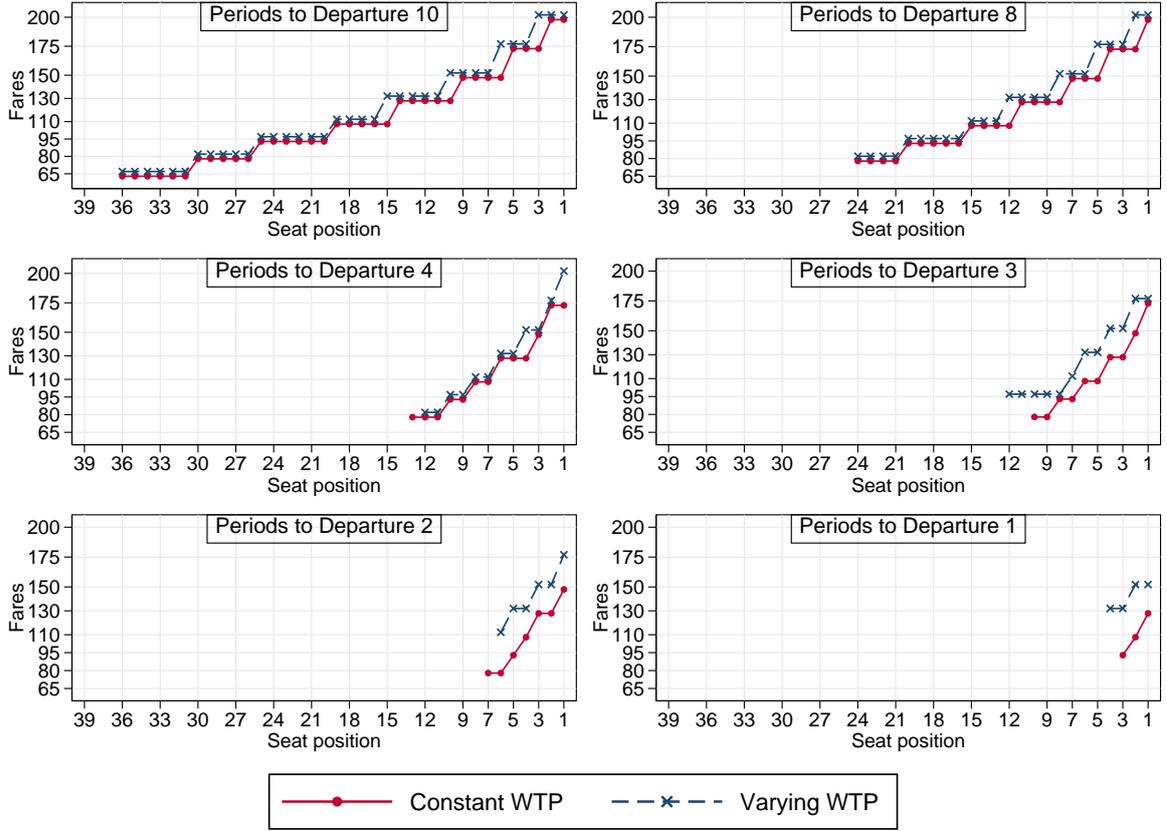
Figure 2 reports the simulated fares at the beginning of periods, with circles and crosses denoting the optimal fare for each seat under a constant and time-varying  $F(\theta)$ , respectively.<sup>10</sup> The Figure highlights several interesting characteristics.

First, the optimization problem is solved by a pricing rule resulting in a sequence of fare classes that, due to the discrete nature of  $F(\theta)$ , is non-strictly increasing ([Dana,](#)

<sup>9</sup>Thus, we set  $\varphi = L/(L + T) \simeq 0.81$ . The simulation solves (2) recursively, see Appendix A.2

<sup>10</sup>To uniquely identify a seat position in every period, positions are counted from right to left because seats, which are sold from left to right, disappear after their sale.

Figure 2: Simulated fare sequences



1999). This result is a direct consequence of the assumption that fares cannot be updated instantaneously within each period, forcing the carrier to set a fare for all capacity units. In all periods, the first seat on sale corresponds to the one positioned at the extreme left. Across periods, the size of each class varies, with lower classes disappearing when sold out. Second, sequences arise in equilibrium regardless of whether  $F(\theta)$  varies over time. Third, in early periods, the sequences with and without a time-varying  $F(\theta)$  tend to overlap, with a slightly larger number of seats allocated to the top classes in the time-varying case. Fourth, starting from period 3, i.e., when the proportion of high WTP travelers increases, more seats are optimally allocated to higher classes in the case of a time-varying  $F(\theta)$ .

Importantly, Figure 2 allows us to identify two forms of DP that are preprogrammed (i.e., they are based only on information available at the start) in the adjustment process. The first denotes seats moved to lower classes in the sequence. These changes over periods are denoted as “S-down”. For example, twelve seats are sold, between period 10, with thirty-six seats available, and period 8, with twenty-four available. Clearly, the optimal fare sequence was modified by moving seats from upper to lower classes; for example, the 200-class, the 150-class and the 95-class all lose one seat between period 10 and 8. This

explains why the 80-class still counts four seats in period 8, despite the simulated sale of twelve units. This form of DP is due to the perishable nature of the seat, which drives down its option value, i.e., the expected value of having an additional seat in a subsequent period (McAfee and te Velde, 2007; Sweeting, 2012). This downward movement of seats is independent of whether  $F(\theta)$  varies over time.

The second form of DP occurs when at least one seat is moved to an upper class in the sequence (henceforth, “S-up”). In the simulation, this instance happens only in the time-varying case and thus provides an insight into how inter-temporal price discrimination can be implemented by airlines. Between periods 4 and 3, the time varying sequence has the same number of available seats, but seats 12 and 11, that in period 4 belong to the 80-class, end up allocated to the 95-class. Thus, the simulation suggests that S-up is an optimal response to an anticipated increase in the proportion of high WTP customers.

To sum up, the simulation has highlighted two key aspects of the pricing algorithm. One, an increasing sequence of prices constitutes an optimal pricing rule when prices cannot be altered instantaneously. Two, the revision of the sequences is revealed by the shift, either upward or downward, of some units’ price. Both aspects will be further investigated and tested in the empirical part of the paper.

### 3 Data

The data collection employed a web crawler, as widely used in the literature (Gaggero and Piga, 2011; Clark and Vincent, 2012; Escobari, 2012; Obermeyer et al., 2013; Li et al., 2014; Alderighi et al., 2015; Bilotkach et al., 2015). It incorporated an experimental design explicitly aimed at retrieving a flight’s sequence of fares, assuming it is stored in the carrier’s reservation system. A successful identification of sequences similar to those in Figure 2 would in itself lend general support to the main assumption of the theoretical set-up that instantaneous price updating is not feasible.

For each flight and departure date, the crawler started by requesting the fare of one seat, and continued by sequentially increasing the number of seats by one unit. The sequence would stop either because the website’s maximum number of seats in a query, equal to 40, was reached or at a smaller number of seats. As in Alderighi et al. (2015), the latter case directly indicates the exact number of seats available on the flight on a particular query date, which we store in a variable called *Available Seats* to track how a

flight occupancy changes as the departure date nears. The former case corresponds to a situation where we know that at least 40 seats still remain to be sold on a given query date; i.e., *Available Seats* is censored at 40.

The crawler automatically connected to the website of easyJet, the second largest European low-cost carrier, and issued queries specifying the departure and arrival airports (the route), the date of departure and the number of seats to be booked. This information was saved together with the flight's departure and arrival time and the fare for the number of seats specified in the query. As it is standard in the literature, to double the data size, the query was for a return flight, with a return date 4 days after the first leg; fares are all denoted in the currency of the country of the originating airport (Bachis and Piga, 2011). In total, we surveyed 9,026 daily flights scheduled to depart during the period October 2016 - May 2017.<sup>11</sup>

With the exception of Escobari et al. (2019), the previous studies of airline pricing that used scraped data were limited to one daily observation per flight, and thus they were not suited to detect the existence and duration static pricing spells. Therefore, the crawler in our study was let to run uninterruptedly over the 14 days preceding the departure time, generating multiple daily queries. To increase the number of daily observations, we focused on 21 routes only and used two PCs to launch the same set of queries, but with a different order.<sup>12</sup> Finally, we checked that web scraping did not engender DP, as in Cavallo (2017).<sup>13</sup>

The saving of the exact time and date of each query allows to gauge the frequency with which each flight's fare sequence was surveyed. On average, the data collection process remained stable over the 14 days' booking period, with an average of 7.2 daily queries per flight, consistent with a median interval of about 3.8 hours between two consecutive observations for the same flight (Table 2).

In the remainder of the paper, the empirical analysis will be carried out only on the uncensored observations, those for which the number of available seats is less than 40.

---

<sup>11</sup>The carrier displays the fare for a query of  $n$  seats as the mean of the  $n$  fares inclusive of a booking fee. In Appendix A.1 we show how to build the fare sequences and its fare classes net of the booking fee, and verify that the size of the classes are consistent with an important information published on the website: the number of seats available at a given posted fare.

<sup>12</sup>The PCs were located in a British and an Italian university, respectively: interestingly, the retrieved data could be matched seamlessly, thus suggesting that the airline drew its posted fares from the same database, regardless of the query's geo-location.

<sup>13</sup>First, the cookie folder was cleaned every day; second, using computers located outside our university offices, we checked a sample of fares scraped by the computers in our university offices. No noticeable difference could be found.

Table 2: Frequency of web scraper’s queries

Days to depart.	Number of daily queries					Interval between queries (hrs)					Obs
	mean	p25	p50	p75	p90	mean	p25	p50	p75	p90	
0-2	6.5	4	6	8	11	4.0	2.0	3.8	4.9	6.2	46,134
3-4	7.0	5	6	8	10	3.9	2.0	3.9	4.9	6.1	86,393
5-6	7.2	5	7	9	11	3.8	2.0	3.8	4.9	6.1	93,075
7-8	7.4	5	7	9	11	3.7	2.0	3.8	4.8	6.0	96,596
9-10	7.3	5	7	9	11	3.7	1.9	3.8	4.9	5.9	98,706
11-14	7.3	5	7	9	11	3.5	1.7	3.5	4.7	5.8	120,854
Total	7.2	5	7	9	11	3.7	1.9	3.8	4.8	6.0	541,758

This allows a precise evaluation the seats’ position in the sequence and the determination of whether the sequence has changed its structure.

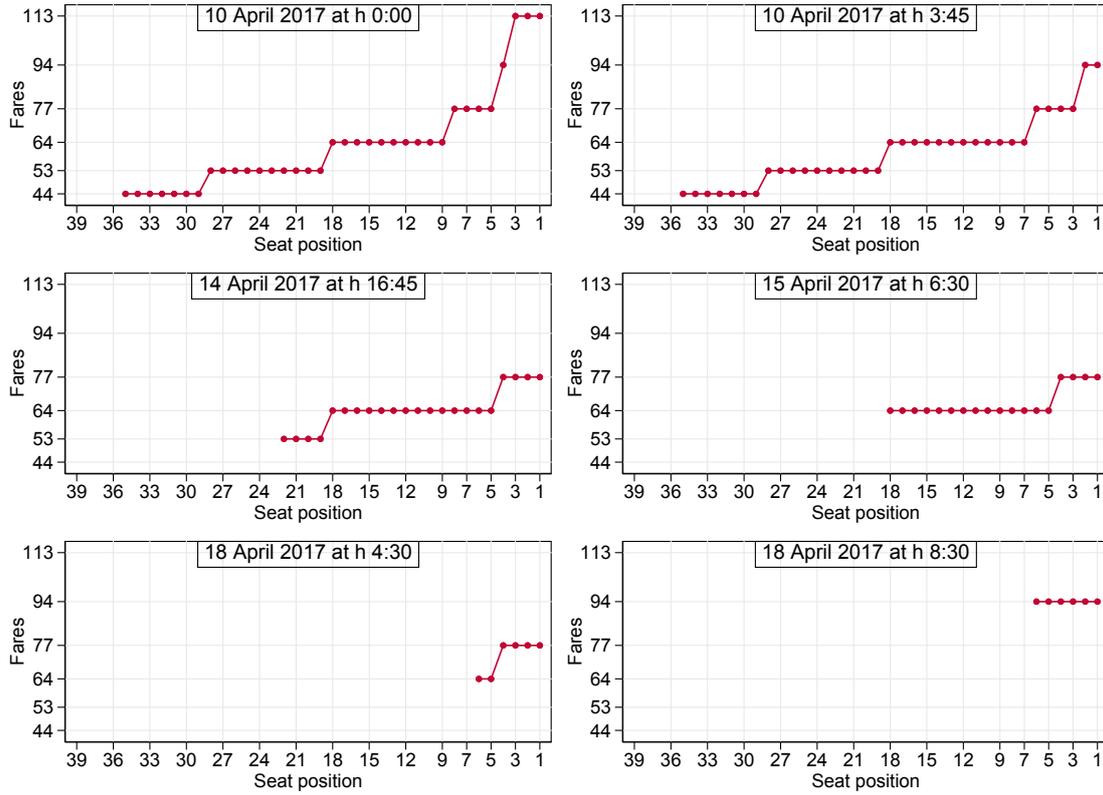
## 4 Linking data to the simulation

A key assumption of the theoretical model is that, due to organizational frictions, the airline may not constantly readjust its fares. A first straightforward test in support of such an assumption is to verify whether the carrier stores, in its reservation system, sequences similar to those generated by the simulation (Figure 2). With no exception, for all the 9,026 flights and at any point in time, the crawler always retrieved fares organized as non-strictly increasing sequences, representatively shown in Figure 3. These are qualitatively identical to the ones obtained from the model’s simulation, based on a discrete WTP distribution.

The equivalence between the theoretical and actual fare sequences provides the first piece of evidence in support of the theoretical model’s assumption that managerial and technological frictions hinder instantaneous offerings of fares. The adoption of fare sequences is also not limited to easyJet; the Appendix A.3 shows similar sequences constructed using data from Ryanair and Southwest. This suggests that this pricing approach is typical in the industry.

The adjustment process emerging from the data is also consistent with the theoretical results in Figure 2, i.e., when a flight’s sequence is changed, some seats’ fares are either moved down or up. For example, the panels in the top row of Figure 3 show a case of S-down. Comparing the left and right hand sides shows that the number of available seats (35) does not change during the interval. However, the €113-class disappears, with its three seats being placed in the €94-class and the €77-class. Moreover, seats 7 and 8,

Figure 3: Examples of easyJet’s fare sequences at various query times.



Note - Flight code U2 4541 Berlin Schoenfeld - Nice departing at 14.10 on 18 Apr 2017.

previously in the €77-class, end up in the €64-class. The structure of the fare sequence has thus changed; hence these panels identify an instance of DP, although the fare of the seat on sale, the one with position 35, has not changed.

The middle panels illustrate the possibility of no DP even in the combined event of a large number of seats being sold and a change in the selling fare. The number of available seats drops from 22 to 18, causing the fare of the first seat to increase from €53 to €64. This example illustrates how, within a static pricing spell, stochastic demand is managed by the sequence: after the four seats in the €53-class were sold, the system automatically posted as selling fare the one corresponding to the next non-empty class. No DP occurs during the almost 14 hours that separate the collection of the left and right middle panels, in the sense that the algorithm has not intervened to alter the structure of the fare sequence.

Finally, the bottom panels depict a case of “S-up”. The last six seats on the flight, which at 4.30am of the departure day were placed in the €64- and €77-classes, by 8.30am were all re-allocated to the higher, and previously closed, €94-class. Notably, the six

marked-up seats, which had stayed in the same fare class for at least four days (as the middle panel shows), were moved up just a few hours before departure, consistent with a standard application of intertemporal price discrimination targeting very last-minute, high-WTP travelers.

## 5 Duration of static pricing spells

The theoretical analysis so far has shown an algorithm with a pricing process that defines fare sequences whose revisions are set “once and for all” via an entirely preprogrammed adjustment process embedded in an IT system. Although some of its predictions found support in our data and play a central role in the remainder of the empirical analysis, no airline entirely relies on a fully automated system to manage its yield. Indeed, airlines hire several analysts to work within their Revenue Management division. Therefore, in Figure 1, revisions of fare sequences can be either preprogrammed or implemented by a RM analyst. In the adjustment process, the analyst takes up the tasks of both assessing external information (a demand shock due to, e.g., a pandemic, a terrorist attack or an important sport final that is relevant for the inhabitants of at least one of the route’s endpoints) and overseeing ‘exception reports’ produced by the IT system. In both cases, the human intervention would translate into a possible revision of the fare sequence. The presence of human analysts allows us to relax the theoretical assumption of static spells of fixed duration, as well as the stationarity of flight’s parameters, enabling the more realistic stance that the adjustment process can intervene at a positive, finite cost. Therefore, the decision of when and how to intervene to modify the fare sequences becomes endogenous. If the adjustment cost is very low, and sequence revision very frequent, we should observe that static spells’s duration in our data to be very close to the average interval between two consecutive scraper’s queries (3.8 hours in Table 2). This would cast serious doubts on the validity of both our theoretical and empirical approaches. However, the opposite conclusion would be supported if the evidence indicated that static pricing spells tend to last long. Therefore, this Section assesses the duration of static pricing spells from both a descriptive and econometric perspective. The next Section builds on this one by further investigating the drivers of the two forms of dynamic pricing, and the possible impact of human interventions.

We measure static pricing spells’ duration,  $S$ , by the time interval that separates the

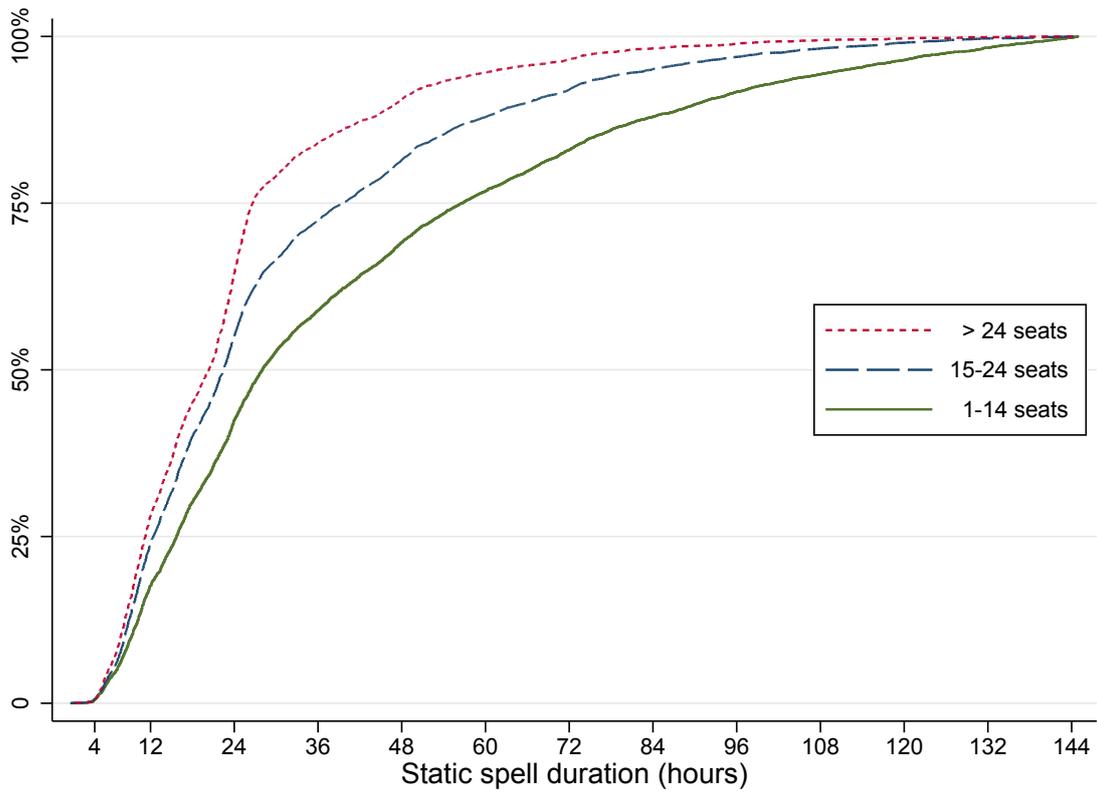
observation of two consecutive instances of DP (i.e., observed changes in the structure of a flight’s fare sequence). Although neither the start nor the end of a spell can be exactly recorded, both are recorded with a delay, because both the start and the end occur before the crawler detects them. However, because the crawler operated continuously with average intervals between two consecutive queries of 3.8 hours throughout the booking period (Table 2), we assume that the two delays cancel out on average; i.e., a late recorded start is offset by a late recorded end.

Table 3: Average duration of static pricing spells (hrs)

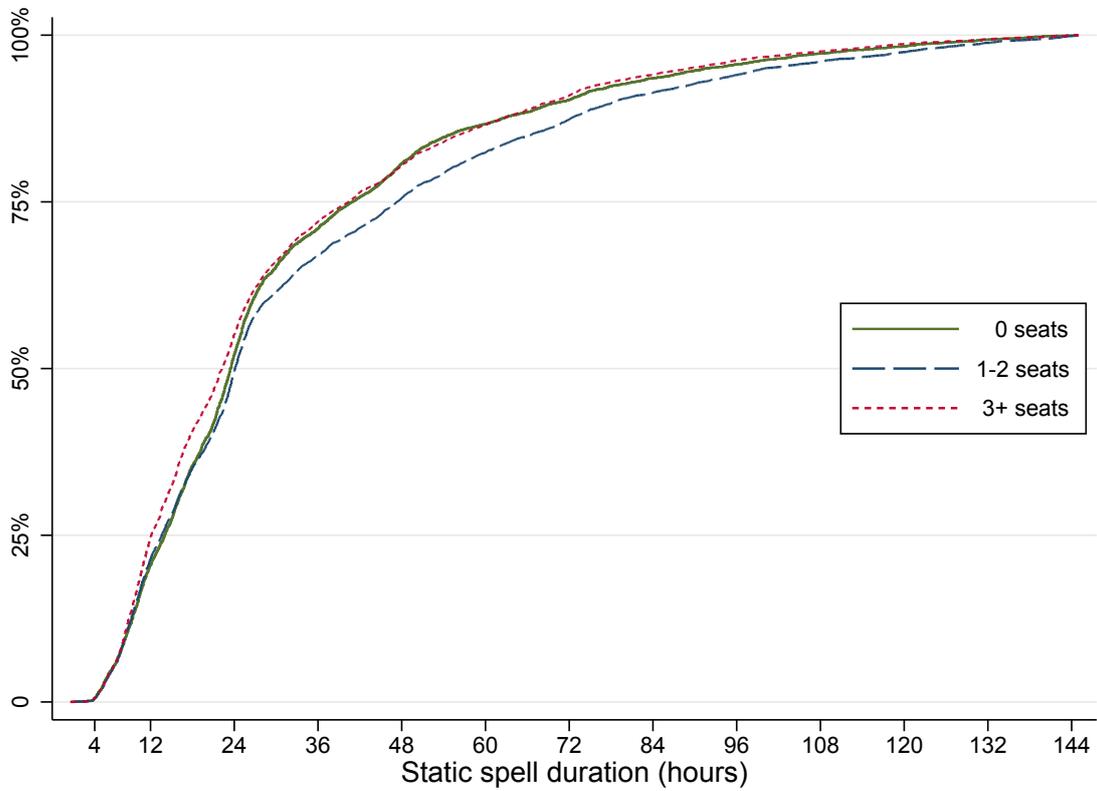
Routes	Hours to departure when spell ends								Overall
	0-23	24-71	72-119	120-167	168-215	216-263	264-311	312+	
AMS MXP	49.7	46.8	35.6	32.2	28.8	21.4	21.8	15.4	30.6
AMS NCE	41.5	39.5	40.6	31.5	29.5	22.8	18.9	15.4	30.7
BFS CDG	51.3	60.5	54.6	56.8	45.5	35.0	24.2	15.9	43.9
BOD BRU	48.0	49.1	32.1	44.2	35.0	25.2	22.3	16.2	33.9
BRU NCE	47.5	35.0	39.1	31.5	32.4	21.7	21.1	14.2	31.9
BUD CDG	70.5	61.6	48.4	40.2	41.5	34.8	22.4	16.3	41.1
BUD GVA	58.3	49.5	51.0	45.5	40.4	25.6	23.3	13.7	38.6
CDG KRK	58.6	51.5	50.2	37.1	32.5	29.8	20.9	14.0	36.0
CDG LGW	31.4	28.3	29.8	24.3	24.1	20.3	19.3	16.0	25.2
CDG LIS	80.0	68.9	51.3	49.6	40.6	28.5	22.6	16.7	41.6
CDG MAD	57.0	52.4	42.8	35.5	34.2	26.2	20.1	17.4	34.1
CDG PRG	52.2	65.6	64.4	53.9	51.9	33.7	23.2	15.8	47.9
FCO LYS	54.8	79.3	62.2	48.7	48.5	31.0	24.2	16.6	46.1
LGW LYS	57.0	42.1	46.0	38.4	39.5	27.3	21.6	15.4	36.4
LGW MRS	60.0	48.0	56.1	41.7	42.1	27.8	23.2	15.9	39.4
LGW VIE	58.5	53.5	52.3	43.9	42.5	29.0	21.2	15.5	39.5
LIN ORY	34.8	29.9	22.2	22.0	22.3	20.1	17.0	16.5	22.7
LYS MAD	51.8	64.5	61.7	46.2	44.7	32.2	22.1	15.3	41.5
NCE STN	23.8	33.4	37.2	35.2	37.1	30.8	19.0	13.4	31.0
NCE SXF	51.4	44.1	43.0	38.4	34.9	30.3	24.1	17.1	36.7
Total	50.8	47.1	42.5	36.1	35.1	26.1	21.7	15.7	34.8

The sample average of spells’ duration is 34.8 hours (Table 3). That is, a flight’s fare sequence remains unaltered for about a day and a half, but in some routes the duration of static pricing spells extends to over two days, depending on time remaining before the flight’s departure. Overall, the evidence confirms the existence of relevant static pricing spells throughout the entire booking period, that justify the use of fare sequences. We now turn the attention to which factors may affect their duration.

Figure 4: Kaplan-Meier cumulative distribution of spells' durations



(a) by median capacity during spell



(b) by selling rate over 24 hours

## 5.1 Drivers affecting a spell’s duration

To model the duration of static pricing spells we need to identify the drivers that are likely to trigger a fare sequence’s revision. The industry-based literature is useful in this regard. The quote from [Belobaba \(2009\)](#), reported in the Introduction, explains that computerized RM systems recommend a sequence’s update when the actual bookings differ significantly from the expected booking curve. The recent literature on airline pricing suggests two variables that can form adequate proxies for a flight’s actual booking curve: a flight’s available capacity and its selling rate ([Puller et al., 2009](#); [Escobari, 2012](#); [Alderighi et al., 2015](#); [Williams, 2017](#)). On the one hand, we hypothesize that when both variables are close to expected levels, the sequence’s revision is not triggered and the spell is extended; that is, no revenue management intervention occurs. On the other hand, two different reasons may explain when a spell is more likely to end. One, if a flight has significantly more available seats than expected or if it is not selling, the sequence’s revision would be in the form of a S-down. Two, a S-up would instead be used when a noticeably lower number of seats remain and/or when seats are being sold at a significantly higher than expected pace. In this second case, the spell of a well-performing flight may be generally longer but may be interrupted at some point to implement a revision upward of the fare sequence. Overall, after controlling for how much time is left before departure, we expect longer spells when available capacity is lower and the selling rate is higher.

Prima facie evidence is shown in [Figure 4](#), whose Panel (a) depicts the Kaplan-Meier cumulative distributions of spells’ duration in sub-samples that differ by the median number of available seats during a spell.<sup>14</sup> Duration is shorter as more seats remain available: the median duration is about 18 hours when 25 or more seats remain available; it increases to about 23 hours when there are between 24 and 15 seats left, and reaches 32 hours for spells occurring when less than 15 seats remain. The evidence is less clear-cut in Panel (b), which shows the number of seats sold from the start of the spell until 24 hours before the end of a spell: when no sale occurs, spells tend to be shorter than when at most two seats are sold.

The distributions in [Figure 4](#) are obtained from the 20,895 spells reported in [Table 4](#), which shows average spells duration broken down by available seats and number of seats

---

<sup>14</sup>Confidence intervals are not shown to enhance clarity, but they are generally very small, and point to significant differences among the distributions.

Table 4: Mean duration of SP spells (hours), by seats sold during the spell and available capacity at the last observation of the spell

Available Seats	Seats sold during the Static Spell					Total
	0	1	2-4	5-7	8+	
1-5	29.8	40.5	45.9	56.7	90.2	54.7
6-10	23.6	29.7	39.2	50.0	79.9	43.7
11-15	20.1	27.3	34.9	45.8	69.4	37.2
16-20	18.5	26.0	30.1	38.9	59.6	32.2
21-30	17.1	21.3	27.4	40.0	53.1	28.3
31+	16.2	21.0	25.2	30.9	40.4	21.0
Total	19.8	25.2	32.2	44.0	70.5	34.8
N col	6,823	2,202	5,602	3,189	3,079	20,895

sold during the spell. Holding the latter fixed, duration is shorter when at the end of the spell the number of seats remaining on sale is higher. However, duration increases as more seats are sold, holding available capacity fixed.<sup>15</sup>

## 5.2 Econometric analysis

The impact of available capacity and selling rate on spells' duration is formally tested using the following specification:

$$S_{idt} = \sum_t \beta_t D_t + \delta AvSeats_{idt} + \gamma Sold24_{idt} + \left( \sum_t \kappa_t D_t * AvSeats_{idt} \right) + \left( \sum_t \omega_t D_t * Sold24_{idt} \right) + \zeta_i + \varepsilon_{it}, \quad (3)$$

where  $S$  denotes the spell duration;  $d$  a spell within flight  $i$  (i.e., flightcode+date of departure);  $D_t$  a set of *Days before departure* dummy variables, with  $t$  being the median number of days before departure during the spell;  $AvSeats$  the number of available seats at the end of the spell;  $Sold24$  the number of seats sold during the spell up to 24 hours preceding the end of the spell;  $\zeta_i$  a flight fixed effect;  $\varepsilon_{it}$  a random disturbance.

The model is estimated first by excluding and then including, one at a time, the terms in square brackets, i.e., the interaction between the temporal dummies with  $Sold24$  and  $AvSeats$ . The latter are both treated as endogenous because  $\varepsilon_{it}$  may reflect the carrier's intervention induced by flight-specific shocks, such as an unexpected surge in demand or an unusual change in the pricing of competitors, which we do not observe. As instruments,

<sup>15</sup>The Table also supports the theoretical approach to allow for multiple units' sale in a period, as about 57% of spells report a sale of at least two seats.

following Alderighi et al. (2015) and Escobari (2012), we propose the mean of lagged values of *Sold24* and *AvSeats*. That is, for a flight departing at date  $x$  and holding the days to departure fixed, we take the mean values of both variables from the flights departing two to seven days before  $x$ . Flights' shocks  $\varepsilon_{it}$  are independent across dates of departure, so lagged values of *Sold24* and *AvSeats* are independent of  $S_{idt}$ , but they are correlated with  $Sold_{idt}$  and  $AvSeats_{idt}$  because flights on the same route share a similar booking curve summarising the airline's expectation on the optimal rate of sales over the booking periods.

For the calculation of the Sargan-Hansen test, in the first column of Table 5 each single instrument is interacted with the temporal dummies  $D_t$ : the statistic indicates that the instruments are valid exclusion restrictions. The estimates in the second column derive from a perfectly identified regression, where the two instruments were not interacted; the coefficients are very similar to the ones shown in the first column. Thus, columns (3) and (4) also report estimates obtained from perfectly identified regressions.

In Table 5, the coefficients of the temporal dummies  $D_t$  show that spells' duration decreases as the departure date nears; i.e., the adjustment process more actively manages the flights during the last few days, when presumably demand uncertainty is lower. Focusing on the first two columns where *AvSeats* and *Sold24* are not interacted with the temporal dummies, the coefficients indicate that an extra available seat shortens a spell's duration by about 3.33 to 3.61 hours. Moreover, selling an extra seat in the previous 24 hours extends the static pricing spell by about 6.30 to 6.67 hours. These findings suggest that the adjustment process intervenes less frequently in well-performing flights.

Similar conclusions hold in the regressions where either *AvSeats* or *Sold24* is interacted with the  $D_t$  dummies (columns 3 and 4). Spells' duration appears to be less sensitive to both *AvSeats* and *Sold24* in the last two days before departure, as their coefficients fall in value and lose significance. In the other periods, higher values of *AvSeats* (*Sold24*) shorten (extend) spells' duration in line with the values in the first two columns, with a slightly larger effect 7-8 days before departure. Overall, the full set of estimates in Table 5 support the hypotheses that the sequence's revision is more frequently triggered when the flight has more remaining capacity and when fewer seats are sold.

Table 5: Instrumental Variable regressions; dependent variable: Spell Duration.

	(1)	(2)	(3)	(4)
AvSeats	-3.335 (0.29)	-3.480 (0.31)		-3.609 (0.30)
(D 0-2)*AvSeats			0.387 (0.91)	
(D 3-4)*AvSeats			-2.351 (0.66)	
(D 5-6)*AvSeats			-4.138 (0.57)	
(D 7-8)*AvSeats			-5.670 (0.53)	
(D 9-10)*AvSeats			-5.729 (0.10)	
(D 11-12)*AvSeats			-3.533 (0.96)	
(D 13-14)*AvSeats			-0.296 (0.92)	
Sold24	6.305 (1.38)	6.669 (1.51)	6.339 (1.88)	
(D 0-2)*Sold24				-0.220 (2.96)
(D 3-4)*Sold24				4.227 (2.16)
(D 5-6)*Sold24				5.540 (1.64)
(D 7-8)*Sold24				12.245 (3.31)
(D 9-10)*Sold24				7.014 (2.14)
(D 11-12)*Sold24				9.378 (1.91)
(D 13-14)*Sold24				3.574 (1.43)
D 0-2	-65.731 (4.65)	-69.262 (5.03)	-52.357 (24.86)	-60.329 (10.00)
D 3-4	-37.808 (3.76)	-40.637 (4.16)	12.546 (28.89)	-43.451 (7.26)
D 5-6	-31.341 (4.18)	-33.920 (4.49)	45.374 (23.70)	-38.433 (4.17)
D7-8	-19.753 (3.46)	-21.864 (3.73)	86.694 (26.72)	-45.939 (9.77)
D 9-10	-12.479 (2.44)	-14.005 (2.66)	104.580 (23.56)	-21.536 (7.18)
D 11-12	-0.639 (1.10)	-1.338 (1.18)	76.224 (18.94)	-14.159 (5.13)
Constant	100.185 (7.30)	103.726 (8.03)	31.628 (25.76)	113.590 (7.80)
N	20,895	20,895	20,895	20,895
Sargan-Hansen	$\chi^2(12)=18.02$	-	-	-

Note: Standard errors in parentheses, clustered by route.

## 6 Forms of Dynamic Pricing

The previous Section does not distinguish which forms of DP cause the end of a static pricing spell, nor does it address the reasons why the algorithm chooses a specific form. Furthermore, the analysis so far has not accounted for any human role in the management of fare sequences, and how it relates to the operativity of the algorithm illustrated in Figure 1. This may be important for competition policy cases, as discussed in the Conclusions section.

Let  $UP_{idt} = 1$  denote an instance of S-up that at time  $t$  ends the spell  $d$  for flight  $i$ ;  $UP_{idt} = 0$  thus denotes a S-down case. That is, we consider the intervals studied in Table 5 and focus on the form of DP that ends the spell. Building on the previous Section, we use the following probit specification:

$$UP_{idt} = 1 \left[ \sum_t \chi_t H_t + \iota AvSeats + \tau Sold24_{idt} + \phi Night_{idt} + \sum_i \eta F_i + \left( \sum_t \varpi_t H_t * AvSeats_{idt} \right) + \left( \sum_t \varrho_t H_t * Sold24_{idt} \right) + \nu_{idt} \right], \quad (4)$$

where  $AvSeats$  and  $Sold24$  are as in eq. (3);  $H_t$  identifies the number of hours to departure, and  $F_i$  denotes an array of flight  $i$ 's characteristics, to control for fixed and time effects: route, departure time (morning, afternoon, evening), departure's day of week, and the week during the year when departure occurs.

Relative to eq. (3), a new variable,  $Night$ , is included to capture the possible role for the RM analysts. It is a dummy initialized when the DP case occurs overnight, between 10:00 p.m. and 6:00 a.m. of the next day, that is, when RM analysts are most likely off duty.<sup>16</sup> We do not have any prior regarding this variable; however, the literature suggests the following considerations.

First, informal discussion with RM analysts informed us that, overnight, carriers launch computer programmes to update their RM system based on the new information accrued during the previous day. These programmes generate recommendations in the form of 'exception reports', which human analysts then evaluate and, in case, implement during the working shift (Belobaba, 2009; Agrawal et al., 2019). Second, Escobari et al. (2019) show evidence of fares in the U.S. domestic market falling overnight, and increasing again

<sup>16</sup> $Night=0$  thus identifies 16 daily hours, representing two 8-hours work shifts.

during working hours. They interpret the increase as evidence of price discrimination against buyers who book during daytime and whose motivation for travel is therefore business-related. Our paper suggests an alternative mechanism, one where ‘exception reports’ identify those flights that RM analysts should consider as candidates for fare increases, for reasons not restricted to inter-temporal price discrimination. Therefore: S-up should be less likely observed overnight, when analysts are not working.

Table 6: Sample and predicted mean values of S-up

	Sample Means				Predictions†	
	Full	with Night	Night=0	Night=1	Full	with Night
<i>Hrs. to Dep.</i> 0-23	0.891	0.884	0.856	0.925	0.891	0.863
<i>Hrs. to Dep.</i> 24-47	0.716	0.690	0.566	0.854	0.716	0.671
<i>Hrs. to Dep.</i> 48-71	0.532	0.509	0.518	0.462	0.534	0.521
<i>Hrs. to Dep.</i> 72-95	0.443	0.422	0.497	0.125	0.448	0.428
<i>Hrs. to Dep.</i> 96-119	0.504	0.507	0.568	0.250	0.509	0.517
<i>Hrs. to Dep.</i> 120-143	0.507	0.529	0.625	0.140	0.511	0.532
<i>Hrs. to Dep.</i> 144-167	0.466	0.482	0.667	0.088	0.468	0.485
<i>Hrs. to Dep.</i> 168-191	0.605	0.616	0.721	0.189	0.608	0.619
<i>Hrs. to Dep.</i> 192-215	0.547	0.562	0.658	0.140	0.547	0.562
<i>Hrs. to Dep.</i> 216-239	0.449	0.455	0.626	0.082	0.451	0.455
<i>Hrs. to Dep.</i> 240+	0.329	0.342	0.592	0.044	0.330	0.340
Sample mean	0.498	0.500	0.623	0.216	0.499	0.500
N	20,895	15,076	10,520	4,556	20,895	15,076

Note: the complement to one of the value in each cell defines the proportion of S-down cases. † Predictions obtained from the estimation of eq. (4), using samples excluding and including *Night*.

*Night* can be measured precisely as either one or zero only if the query preceding the end of the spell also occurs at, respectively night or daytime; this results in some missing values.<sup>17</sup> Therefore, we delve deeper into some aspects of this variable to make sure that the missing values do not bias the analysis. In Table 6, the first two columns report the descriptive sample mean of a S-up, in the full sample and in the sample where *Night* is fully identified. No appreciable difference appears (e.g., sample averages equal to 0.498 and 0.50), indicating that the missing values do not select the sample arbitrarily. Similar values are found in both samples when we condition by hours to departure. Importantly, a much larger proportion of S-up occurs during working hours (column *Night* = 0 relative to *Night* = 1) when more than three days (72 hours) separate the intervention from the departure. This finding is consistent with a more intense human involvement in the decision of a fare increase earlier before departure. However, the distinction wanes completely

<sup>17</sup>Consider one DP observation taken at midnight, causing the end of a spell. If the previous observation was taken at 23:00 hours, then we know that the decision to revise the sequence was taken during nighttime. However, this could not be ascertained if the previous observation were retrieved at, say, 21:30.

during the last 48 hours before departure, when the airline may be committed to pursue an inter-temporal price discrimination strategy, a task which appears to be largely delegated to the algorithm as it occurs with similar probabilities both day and night.

To sum up, the foregoing discussion suggests the following hypotheses. First, we expect S-up to reflect inter-temporal price discrimination if, all else equal, its adoption is more often observed during the last few days before departure (Hp.1). Second, higher values of *Sold24* (*AvSeats*) are expected to enhance (reduce) the probability to observe  $UP_{idt} = 1$  (Hp.2); third, this probability is lower during nighttime (Hp.3).

### 6.1 Dissecting the dynamic pricing decision

These hypotheses are formally tested by estimating eq. (4) with a probit model where *AvSeats* and *Sold24* are treated as endogenous and instrumented using the same exclusion restrictions of eq. (3).

To gauge the probit model’s goodness of fit, the last two columns of Table 6 report the predicted mean values of  $UP_{idt}$ ; they are very similar to the samples’ means reported in the first two columns, indicating that the probit models accurately predict which form of DP will be selected by the adjustment process during the fortnight preceding a flight’s departure. The same estimates also support Hp.1, as they confirm the larger proportion of S-up over the last two days, relative to earlier periods.

Given the good predictive properties of the probit models, we further investigate the other hypotheses using Table 7, which reports the marginal effects of *Night*, *Sold24* and *Avseats* on the probability to observe  $UP = 1$ , using the samples with and without *Night* to estimate eq. (4). Each column reports an average effect for each variable, which is then disaggregated by hours before departure. The first column predicts an overall 35.3% lower probability of observing S-up overnight; further support for Hp.3 arises from similar values being found in any period prior to departure, except within the last 24 hours when the effect falls to  $-24.3\%$ . This result further strengthens the previous conclusion that the airline routinely commits to an inter-temporal price discrimination strategy during the last few days before departure. Overall, the evidence suggesting that S-up is generally applied during daytime is compelling and point to an organizational process where computerized systems can be used to augment more sensitive decision-making by humans (Agrawal et al., 2019).

Table 7: Marginal Effects on probability of  $UP$

Variable	<i>Night</i>	<i>Sold24</i>		<i>AvSeats</i>	
		w/Night <sup>‡</sup>	Full <sup>‡</sup>	w/Night <sup>‡</sup>	Full <sup>‡</sup>
Full Avg. Effect	-0.353 (0.023)	0.062 (0.006)	0.055 (0.007)	-0.003 (0.004)	-0.004 (0.003)
<i>Hrs. to Dep.</i> 0-23	-0.243 (0.031)	0.019 (0.004)	0.016 (0.005)	0.006 (0.002)	0.005 (0.002)
<i>Hrs. to Dep.</i> 24-47	-0.400 (0.038)	0.044 (0.009)	0.043 (0.01)	0.014 (0.004)	0.015 (0.005)
<i>Hrs. to Dep.</i> 48-71	-0.399 (0.026)	0.055 (0.008)	0.054 (0.01)	0.005 (0.006)	0.002 (0.005)
<i>Hrs. to Dep.</i> 72-95	-0.342 (0.018)	0.054 (0.009)	0.049 (0.008)	-0.007 (0.004)	-0.008 (0.004)
<i>Hrs. to Dep.</i> 96-119	-0.384 (0.026)	0.056 (0.008)	0.053 (0.009)	-0.003 (0.004)	-0.004 (0.004)
<i>Hrs. to Dep.</i> 120-143	-0.354 (0.026)	0.064 (0.006)	0.059 (0.008)	-0.008 (0.005)	-0.009 (0.004)
<i>Hrs. to Dep.</i> 144-167	-0.374 (0.028)	0.062 (0.008)	0.051 (0.008)	-0.005 (0.004)	-0.008 (0.003)
<i>Hrs. to Dep.</i> 168-191	-0.369 (0.025)	0.064 (0.008)	0.061 (0.009)	-0.007 (0.005)	-0.008 (0.004)
<i>Hrs. to Dep.</i> 192-215	-0.366 (0.027)	0.082 (0.006)	0.075 (0.008)	-0.004 (0.004)	-0.004 (0.004)
<i>Hrs. to Dep.</i> 216-239	-0.345 (0.023)	0.069 (0.007)	0.063 (0.008)	-0.009 (0.003)	-0.008 (0.003)
<i>Hrs. to Dep.</i> 240+	-0.338 (0.023)	0.074 (0.007)	0.061 (0.008)	-0.002 (0.003)	-0.004 (0.003)
N	15,076	20,895	15,076	20,895	15,076

Note: Standard errors in parentheses, clustered by route. <sup>‡</sup> Specification of eq. (4) with and without *Night*.

As far as the impact of *Sold24* is concerned, the evidence generally supports Hp.2, as a unit increment in *Sold24* increases the probability to observe a S-up by about 6.2% on average. Interestingly, the impact is stronger seven days (i.e., 168 hours) or more before departure but it falls well below the sample average during the last 48 hours. The effect of *AvSeats* is also in the right direction, but quite small in magnitude (less than  $-1.0\%$ ). These results further corroborate the hypothesis that in the last few days S-up is predominantly used for inter-temporal price discrimination purposes, as the algorithm appears not to be sensitive to the occupancy rate and the selling rate. Instead, the robust effect of *Sold24* in earlier periods reflects an unexpected positive shock in demand and thus links earlier manifestations of S-up to reasons other than inter-temporal price discrimination, i.e., the need to manage the yield of well-performing flights.

## 7 Conclusions

In the absence of managerial frictions and technological bottlenecks, the airline's pricing problem would consist in deciding just one fare for all seats, as in [Williams \(2017\)](#) and [Lazarev \(2013\)](#). Such a fare could be changed instantaneously.

The present study investigated the characteristics of the solution to an optimal pricing problem for multiple units of capacity, when instantaneous pricing is not feasible. In line with [Dana \(1999\)](#), the solution involves the definition of a price sequence, specifying a (possibly different) price for each unit of capacity. The data collected from easyJet's website show sequences to be the way airlines store their fares on distribution channels.

A managerial implication of our study is that sequences serve as an insurance mechanism, by protecting the carrier against the risks that using a unique fare for all capacity would entail. Imagine if the carrier set the same low fare for all the seats and the demand for the flights turned out to be extremely high; the impossibility to adjust the fare instantaneously would cause the under-selling of most or all of flight capacity. If, as a remedy, only a few seats were made available at a single low fare, and they sold out, then no more seats could be sold, leading the system to reject possible subsequent higher bids' requests. A non-strictly increasing sequence efficiently addresses and solves both aspects of this problem. This insight is relevant for other industries where multiple units of capacity are sold in advance (e.g., hotels, ferries, cruise, car rentals, railways).

Various findings in this paper have implications for competition policy. In merger

cases, our analysis implies that one possible effect is the adoption by the target firm of the acquiring firm's revenue management methods, including the parts of the algorithm responsible for the creation of the fare sequences, inter-temporal price discrimination and dynamic pricing. Post-merger fares will be more generally driven by the way the acquiring firm defines and updates its sequences. [Dobson and Piga \(2013\)](#), for instance, find that the inter-temporal profiles of fares became steeper after Ryanair took over the routes previously served by the acquired company, Buzz, with early fares becoming cheaper and late fares more expensive.

Furthermore, the Air Berlin-Lufthansa case has highlighted the possibility that algorithms may be used to abuse a dominant position ([www.thelocal.de](http://www.thelocal.de), 2017). Following the Air Berlin collapse, Lufthansa denied changing its pricing methods, arguing that its fully-automated booking system was simply responding to a spike in demand and displaying higher average prices as a result. On the one hand, in our analysis such increases may be due to more seats being sold, i.e., to movements along a fixed and predetermined sequence, which would support Lufthansa's argument. On the other hand, Lufthansa's revenue management analysts could have been involved in the decision to alter the fare sequences for flights in affected routes. In the light of our finding, a selective intervention indicative of a possible abuse of dominant position, would be identified by a drastic shift upward of the sequences only on the affected routes but not on non-affected ones.<sup>18</sup> Overall, our analysis illustrates how, even in an industry renowned for its extensive reliance on revenue management techniques, decision-making by human agents continues to play a crucial role. It is left to future research to evaluate to what extent improvements in the Artificial Intelligence technology will diminish the human role.

---

<sup>18</sup>In May 2018 the German Cartel Office decided not to investigate Lufthansa for market abuse over rising ticket prices following the collapse of Air Berlin, see <https://uk.reuters.com/article/us-lufthansa-fares/no-antitrust-probe-for-lufthansa-over-fares-after-air-berlin-collapse-idUKKCN1IU10G>

## References

- Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction. *Journal of Economic Perspectives*, 33(2):31–50.
- Alderighi, M., Gaggero, A. A., and Piga, C. A. (2020). Hidden prices with fixed inventory: evidence from the airline industry. <https://ssrn.com/abstract=3085585>.
- Alderighi, M., Nicolini, M., and Piga, C. A. (2015). Combined Effects of Load Factors and Booking Time on Fares: Insights from the Yield Management of a Low-Cost Airline. *Review of Economics and Statistics*, 97:900–915.
- Assad, S., Clark, R., Ershov, D., and Xu, L. (2020). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. mimeo.
- Bachis, E. and Piga, C. A. (2011). Low-Cost Airlines and online Price Dispersion. *International Journal of Industrial Organization*, 26(6):655–667.
- Belobaba, P. P. (2009). Fundamentals of pricing and revenue management. In Belobaba, P., Odoni, A., and Barnhart, C., editors, *The Global Airline Industry*, chapter 4, pages 73–112. Wiley & Sons, Chichester, UK.
- Bilotkach, V., Gaggero, A. A., and Piga, C. A. (2015). Airline pricing under different market conditions: Evidence from European Low-Cost Carriers. *Tourism Management*, 47:152–163.
- Bloom, N. and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, 122(4):1351–1408.
- Calvano, E., Calzolari, G., Denicoló, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing and collusion. *American Economic Review*, 110(10):3267–3297.
- Cavallo, A. (2017). Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers. *American Economic Review*, 107(1):283–303.
- Cho, S., Lee, G., Rust, J., and Yu, M. (2018). Optimal dynamic hotel pricing. mimeo.
- Clark, R. and Vincent, N. (2012). Capacity-contingent pricing and competition in the airline industry. *Journal of Air Transport Management*, 24:7–11.

- Dana, J. D. (1999). Equilibrium price dispersion under demand uncertainty: the roles of costly capacity and market structure. *Rand Journal of Economics*, 30(4):632–660.
- DellaVigna, S. and Gentzkow, M. (2019). Uniform pricing in U.S. retail chains. *The Quarterly Journal of Economics*, 134(4):2011–2084.
- Deneckere, R. and Peck, J. (2012). Dynamic Competition with Random Demand and Costless Search: a Theory of Price Posting. *Econometrica*, 80(3):1185–1247.
- Dobson, P. W. and Piga, C. A. (2013). The Impact of Mergers on Fares Structure: Evidence from European Low-Cost Airlines. *Economic Inquiry*, 51(2):1196–1217.
- Ellison, S. F., Snyder, C., and Zhang, H. (2018). Costs of managerial attention and activity as a source of sticky prices: Structural estimates from an online market. Working Paper 24680, National Bureau of Economic Research.
- Escobari, D. (2012). Dynamic pricing, advance sales, and aggregate demand learning in airlines. *Journal of Industrial Economics*, 60(4):697–724.
- Escobari, D., Rupp, N. G., and Meskey, J. (2019). An Analysis of Dynamic Price Discrimination in Airlines. *Southern Economic Journal*, 85(3):639–662.
- Gaggero, A. A. and Piga, C. A. (2011). Airline Market Power and Intertemporal Price Dispersion. *Journal of Industrial Economics*, 59(4):552–577.
- Gallego, G. and van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020.
- Gerardi, K. and Shapiro, A. (2009). Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry. *Journal of Political Economy*, 117:1–37.
- Gershkov, A. and Moldovanu, B. (2009). Dynamic Revenue Maximization with Heterogenous Objects: A Mechanism Design Approach. *American Economic Journal: Microeconomics*, 1(2):168–198.
- Harrington, J. E. (2019). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law and Economics*, forthcoming.

- Kleinberg, J., Himabindu, L., Jure, L., Jens, L., and Mullainathan, S. (2018). Human Decisions and Machine Predictions. *Quarterly Journal of Economics*, 133(1):237–293.
- Koenigsberg, O., Muller, E., and Vilcassim, N. (2008). easyJet pricing strategy: should low-fare airlines offer last-minute deals? *Quantitative Marketing and Economics*, 6:279–297.
- Lazarev, J. (2013). The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in U.S. monopoly markets. *mimeo*. 2013.
- Li, J., Granados, N., and Netessine, S. (2014). Are Consumers Strategic? Structural Estimation from the Air-Travel Industry. *Management Science*, 60:2114–2137.
- McAfee, R. P. and te Velde, V. (2007). Dynamic pricing in the airline industry. In Hendershott, T. J., editor, *Handbook of Economics and Information Systems*, volume Vol 1. Elsevier Science, New York, NY, USA.
- Miklós-Thal, J. and Tucker, C. (2019). Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers? *Management Science*, 65(4):1552–1561.
- Misra, K., Schwartz, E. M., and Abernethy, J. (2019). Dynamic Online Pricing with Incomplete Information Using Multiarmed Bandit Experiments. *Marketing Science*, 38(2):226–252.
- Möller, M. and Watanabe, M. (2010). Advance Purchase Discounts and Clearance Sales. *Economic Journal*, 120:1125–1148.
- Nijs, V. R., Misra, K., and Hansen, K. (2014). Outsourcing Retail Pricing to a Category Captain: The Role of Information Firewalls. *Marketing Science*, 33(1):66–81.
- Obermeyer, A., Evangelinos, C., and Püschel, R. (2013). Price dispersion and competition in European airline markets. *Journal of Air Transport Management*, 26:31–34.
- Pang, Z., Berman, O., and Hu, M. (2015). Up then down: Bid-price trends in revenue management. *Production and Operations Management*, 24(7):1135–1147.

- Puller, S. L., Sengupta, A., and Wiggins, S. (2009). Testing theories of Scarcity pricing and Price Dispersion in the Airline Industry. N. 15555, December, NBER Working Paper Series.
- Swan, W. (2002). Airline demand distributions: passenger revenue management and spill. *Transportation Research Part E*, 58:253–263.
- Sweeting, A. (2012). Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets. *Journal of Political Economy*, 120(6):1133–1172.
- Travel Weekly (2018). Airlines inching closer to dynamic pricing. [www.travelweekly.com/Travel-News/Airline-News/Airlines-inching-closer-to-dynamic-pricing](http://www.travelweekly.com/Travel-News/Airline-News/Airlines-inching-closer-to-dynamic-pricing).
- Williams, K. R. (2017). Dynamic Airline Pricing and Seat Availability. Discussion paper no. 3003, Cowles Foundation.
- [www.thelocal.de](http://www.thelocal.de) (2017). Watchdog accuses Lufthansa of price hikes in wake of Air Berlin bankruptcy. [www.thelocal.de/20171228/is-lufthansa-raising-prices](http://www.thelocal.de/20171228/is-lufthansa-raising-prices).
- Zbaracki, M., Ritson, M., Levy, D., Dutta, S., and Bergen, M. (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *Review of Economics and Statistics*, 86(2):514–533.

## A Appendix - For Online Publication

### A.1 Building the distributions from easyJet’s posted fares.

This Section contains further details on the procedure we applied to derive the fare distributions from the posted fares.

Through data visual inspection, we learnt that the carrier’s posted fares follow this rule:

$$PF(n) = \frac{C + \sum_{j=1}^n p_j}{n}, \quad (\text{A.1})$$

where  $n$  denotes the number of seats in the query,  $PF(n)$  the corresponding posted fare,  $p_j$  the fare of each seat, starting from the first one available for sale and  $C$  is a fixed charge which we interpret as a fixed commission per booking. The presence of  $C$  implies that the distribution of posted fares over seats is generally U-shaped, with the decreasing part due to the commission being spread over more seats and the increasing part due to the increasing values of fare classes, as in Figure 3.

To find  $C$ , we rely on the fact that in most cases the first and the second seat are likely to belong to the same fare class. Therefore  $C$  (and the value of the first fare class) can be obtained by solving the following system of two linear equations in two unknowns, using the identity  $p_1 = p_2 = p$ :

$$\begin{aligned} PF(1) &= C + p \\ PF(2) &= (C + 2p)/2 \end{aligned}$$

After finding  $C$ , using (A.1) we can derive the price level of a fare class,  $P_j$ :

$$P_j = j * PF(j) - (j - 1) * PF(j - 1) \text{ with } j \in [2, 40], \quad (\text{A.2})$$

with  $P_1 = PF(1) - C$ .<sup>19</sup>

Two aspects are noteworthy. First, the procedure to derive the fare classes’ levels does not impose any restriction on the monotonicity of the distribution. Second, and most importantly, the distributions we derive correspond exactly to the distributions advertised on the carrier’s website. As discussed in the Data Collection section, for each query the crawler retrieved the information that appears on the booking page regarding the “number

---

<sup>19</sup>For simplicity, cents and pennies are rounded to unity.

of seats available at that fare”. We can then gauge the extent to which the size of each fare class, obtained from (A.2), conforms with the information provided by the carrier. It turns out that the above procedure generates buckets’ sizes that perfectly correspond to the sizes implied by the information posted by the carrier on the number of seats available at a given fare. We take this aspect as a strong indication that we succeeded in reverse-engineering the carrier’s pricing approach.

## A.2 Recursive Solution

Rewrite eq. (2) as:

$$V(t, M) = \max_p \{q(p) [p + V(t, M - 1) - V(t - 1, M)]\} + V(t - 1, M) \quad (\text{A.3})$$

with boundary conditions  $V(t, 0) = 0$  and  $V(0, M) = 0$ , for any  $t \in \{0, \dots, T\}$  and  $M \in \{0, \dots, N\}$ . To find a solution for the problem described in (A.3), we consider the following steps.

**Step 1.** Find the solution for  $\max_{p \in \Theta} q(p) (p + X)$ . Because  $\Theta$  is compact, there exists a solution for the problem.

**Step 2.** Set  $t = 1$  and  $M = 1$ .

**Step 3.** Compute  $X = V(t, M - 1) - V(t - 1, M)$  and use Step 1 to get  $p(t, M)$ . Replace it in (A.3) to obtain  $V(t, M)$ .

**Step 4.** Set  $m = m + 1$ . Repeat Step 3 until  $m = N$ .

**Step 5.** Set  $t = t + 1$  and  $m=1$ . If  $t < T$ , then go back to Step 3.

### A.3 Fare sequences from other airlines

Figure A1: Ryanair's sequences



Figure A2: Southwest's sequences

