International Journal of Industrial Organization 71 (2020) 102639

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Dynamic oligopoly pricing: Evidence from the airline industry

Caspar Siegert[a], Robert Ulbricht[b,*]

[a] Bank of England, London, United Kingdom
[b] Boston College, Newton, MA, USA

ARTICLE INFO

Article history:
Received 21 October 2019
Revised 19 May 2020
Accepted 26 May 2020
Available online 10 June 2020

JEL classification:
D43
D92
L11
L93

Keywords:
Airline industry
Dynamic oligopoly pricing
Intertemporal price dispersion
Price discrimination

ABSTRACT

We explore how pricing dynamics in the European airline industry vary with the competitive environment and with customer heterogeneity. We document three main findings. First, the rate at which prices increase towards the scheduled departure date is significantly reduced in more competitive markets. Second, the sensitivity of the intertemporal slope to competition increases in the heterogeneity of the customer base. Third, ex-ante predictable advance purchase discounts account for 83 percent of within-flight dispersion in prices and for 17 percent of cross-market variation in pricing dynamics.

1. Introduction

The tendency for airline ticket prices to rise as the scheduled departure date approaches is a well-known regularity of airline markets. Yet, while there is an active literature researching pricing dynamics in monopoly markets, our understanding of dynamic oligopoly markets is still limited. How does competition affect the scope for intertemporal price differentiation? Which markets are most likely to see steep intertemporal price gradients? Are price increases predictable conditional on the market environment, or do they appear random from the perspective of an outside observer?

In this paper, we use novel data on the time path of prices from the European airline industry to study these questions. We empirically explore how pricing dynamics vary with the competitive environment, document a pivotal role of customer heterogeneity for determining the intertemporal gradient, and investigate the importance of ex-ante predictable advance purchase discounts relative to (possibly stochastic) residual volatility for realized price dynamics.

Main findings

To explore how moving from monopoly markets to oligopoly markets affects pricing dynamics, we begin our analysis by estimating the intertemporal slope of prices and its sensitivity to competition. Overall, we find that prices in our sample increase substantially over time, but at a rate that is highly sensitive to competition. While monopoly prices increase by an average of 1.31 percent with every day that a customer waits to book, this slope is reduced to 1.19 percent in duopolies and continues to decrease monotonically to a slope of 0.90 percent in markets with 5+ competing airlines. A nonparametric

* Corresponding author.
E-mail addresses: caspar.siegert@bankofengland.co.uk (C. Siegert), ulbricht@bc.edu (R. Ulbricht).

https://doi.org/10.1016/j.ijindorg.2020.102639
0167-7187/

treatment of pricing dynamics further reveals that these differences are mainly driven by the last 5 weeks before departure. Over this period, prices on monopoly routes increase by roughly twice as much as they do on routes served by 5+ competitors (159 vs. 80 percent).

In light of this discrepancy, a natural question then is, "Does competition universally flatten the intertemporal price gradient in all markets, or does it do so only in certain markets?" We approach this question by looking at one potentially important enabling factor—customer heterogeneity. To do so, we construct a novel index of customer heterogeneity that is based on the realization that high levels of customer heterogeneity are likely to induce *monopoly* airlines to engage in intertemporal price differentiation.[1] Building on this realization, our index combines eight market indicators, each designed to capture variations in the intensity of tourist to business travelers, by identifying the particular linear combination that is most likely to steepen the intertemporal price gradient on monopoly markets.

Our results provide strong support for the notion that competition affects pricing dynamics *differentially* by customer heterogeneity. In markets with a highly heterogeneous customer base, competition flattens the intertemporal slope from a daily rate of 1.42 percent in monopoly markets to 0.90 percent in markets with five competitors. By contrast, in markets with little customer heterogeneity, we find virtually no effect of competition on the intertemporal slope.

Having documented these patterns, we then ask how much of the observed within-flight dispersion in prices is due to the documented advance purchase discounts, and how much is due to residual price volatility around the competition-specific trends? We find that for the average flight the documented advance purchase discounts account for 83 percent of the overall intertemporal dispersion in ticket prices. That is, while there is some residual volatility around the systematic price gradient, a large share of the intertemporal price dispersion that we observe is explained by ex-ante predictable advance purchase discounts. Lastly, looking at the *differences* in pricing dynamics across markets, we find that about 17 percent of the observed cross-market variation is explained by competition and customer heterogeneity.

Interpretation

From a theoretical perspective, our understanding of pricing dynamics in *oligopoly* markets is currently limited, mainly due to the challenge of incorporating strategic interactions into stochastic control problems. While a few studies explore dynamic oligopoly models, these treatments are rare, typically focus on stylized settings with two selling periods as in Anton et al. (2014) and Dana and Williams (2019), and their predictions are sensitive to modeling choices.[2]

The primary aim of this paper is to provide empirical guidance for future theoretical developments. Nevertheless, we offer some *prima facie* interpretation of our results in Section 6. In our view, the documented impact of customer heterogeneity on pricing dynamics hints towards intertemporal price discrimination as a likely source of advance purchase discounts. Intuitively, monopoly airlines discriminate against late booking customers with inelastic demand, but are restrained in their ability to do so in more competitive environments. In line with this intuition, we expect that the intertemporal price gradient is particularly sensitive to competition when there is a high potential to discriminate against late booking customers in the first place, which is precisely what our findings regarding customer heterogeneity have shown.

On the other hand, our findings also indicate that even in highly competitive markets, prices tend to systematically increase over time, suggesting another force at play. One possibility is that airlines face aggregate uncertainty regarding their demand, which in combination with capacity constraints may support advance purchase sales to low-valuation customers even in perfectly competitive markets (Dana, 1998).[3] A corollary to such a stochastic demand interpretation is that even though prices may *on average* increase over time, the *realized* price path will depend on the realized demand and cannot be perfectly predicted by ex-ante market characteristics. Our results regarding the predictability of within-flight price dispersion and cross-market differences indicate that this is indeed the case.

Relation to empirical literature

From an empirical perspective, the analysis of pricing dynamics has proven difficult mainly due to a lack of public data. In the airline industry, public price data is available only at a route-quarter level that pools prices across different itineraries and travel dates, preempting the study of pricing dynamics for a given flight. For a long time, the literature has therefore focused on the impact of competition on broad measures of price dispersion (e.g, Borenstein and Rose, 1994; Hayes and Ross, 1998; Gerardi and Shapiro, 2009; Dai et al., 2014).

In this paper, we address this issue using posted price data collected from a leading online booking website. In particular, we construct a panel including about 1.4 million prices for airline tickets on the intra-European market where for each route-date pair, we record a time series of posted prices ranging from 10 weeks to 1 day prior to departure. Using this time series dimension permits us to shift the focus on pricing dynamics and their determinants.

To relate our findings to the earlier price dispersion literature, we also use our data to disentangle the impact of competition on different dimensions of price dispersion. In line with the documented pricing dynamics, we find that competition has an unambiguously negative impact on within-flight price dispersion. However, if we compute price dispersion by pool-

[1] While theory does not provide clear guidance regarding the scope to which *oligopolists* can exploit heterogeneity in the customer base, it is optimal under a wide range of assumptions for monopolist airlines to offer advance purchase discounts when late booking customers have a higher willingness to pay (e.g., Gale and Holmes, 1993; McAfee and Te Velde, 2007; Williams).

[2] There exists a much larger literature developing dynamic pricing theories for monopoly markets (see, e.g., Talluri and Van Ryzin (2006) for a textbook treatment). Applications to the airline industry include Gale and Holmes, 1993, McAfee and Te Velde (2007), Lazarev (2013), and Williams (2017). At the opposite extreme of the competitive spectrum, Dana (1998, 1999) develops dynamic pricing theories for perfectly competitive markets.

[3] See also Gale and Holmes, 1992, 1993 and Dana (2001) for similar arguments.

ing across different travel dates or across different flights within a given route, that relation becomes diluted (in the former case) or even overturned (in the latter case). Hence, while we find competition to have an unambiguous negative impact on intertemporal price dispersion, the relation between competition and cross-flight dispersion in our data is less clear, which may explain seemingly contradictory findings in the earlier literature.[4] Related to our attempt to disentangle various dimensions of price dispersion are Puller et al. (2015), and especially Gaggero and Piga (2011) who have also documented a negative relationship between competition and within-flight price dispersion.

In using hand-collected data from the internet, we join a recent generation of papers with similar data strategies. Lazarev (2013) and Williams (2017) use scraped data to estimate dynamic monopoly models, but do not have data for other competitive environments. By contrast, Gaggero and Piga (2011), Escobari (2012) and Escobari et al. (2019) have price data for oligopoly markets, but are interested in the determinants of price levels and dispersion rather than pricing dynamics. Specifically, Escobari (2012) explores how prices adjust to demand shocks, and Escobari et al. (2019) study price-discrimination between bookings in business hours versus bookings in the evening. To the best of our knowledge, this is the first paper which investigates empirically how pricing dynamics vary with the competitive environment.

Layout

The paper is structured as follows. Section 2 describes the data. Section 3 documents how pricing dynamics vary with the competitive environment. Section 4 introduces our measure of customer heterogeneity, and investigates how it interacts with competition in determining pricing dynamics. Section 5 decomposes price dispersion along various dimensions. Section 6 offers some interpretation of the findings.

2. Data

Our primary dataset is a panel of airline ticket prices, collected from a leading online booking website, on 92 intra-European routes and 41 distinct travel dates, where for each route-date pair we record a time series of prices ranging from 10 weeks to 1 day prior to departure. Fig. 1 illustrates the cross-section of routes; Appendix Appendix A provides a full listing of routes and describes the selection process.[5] The data covers virtually all direct flights offered by traditional airlines on these routes, as well as most low cost carriers. We exclude indirect flights, because on intra-European routes these are arguably no close substitutes to direct flights.

Prices are recorded for flights taking off between October 31, 2010 and March 26, 2011, which defines the 2010/2011 European winter flight schedule.[6] We record prices for all return flights leaving on Friday and returning on Sunday within that period, as well as for all flights leaving on Monday and returning on Thursday. This gives us two distinct travel dates per week, resulting in 41 distinct travel dates per route (uniquely defined by the date of the outbound flight). We refer to these “route × travel-date” combination as “markets”.

For each market, we monitor available flights and ticket prices once a week, starting 10 weeks prior to departure. In the last week prior to departure, we monitor prices daily to account for an increased frequency of price changes. For each combination of outgoing and incoming flights, we record the lowest fare available at each of the 17 potential purchase dates, treating code-sharing flights as distinct flights unless the involved airlines are affiliated through equity (see below).[7]

Here and throughout the paper, we reserve the term “flight” to refer to a specific physical flight (identified by its flight identification number) *at a specific travel date*. By contrast, we use the term “itinerary” to refer to a specific routing (identified by its flight identification number) *on any of the 41 potential travel dates*. For example, in our terminology the Paris–London routing involving the outbound British Airways flight BA 333 and the inbound flight BA 334 would be a specific roundtrip *itinerary*, while the same routing for any particular travel date would constitute a combination of (outgoing and incoming) *flights*.

Overall we have data on 3762 out of 3772 distinct markets (41 travel dates times 92 routes).[8] Each market averages 377 prices that are recorded over up to 17 different dates prior to departure for an average of 41.9 flight combinations (i.e., roundtrip combinations) per market. In total, our data set consists of 1.42 million individual prices (92 routes times 41 travel dates times 377 recorded prices per market). Routes are on average 560 miles long and connect metropolitan areas with an average of 3.9 million inhabitants. The share of domestic routes in our sample is roughly 13 percent (12 out of 92 routes).

Prices are collected from an online booking website, which accounts for a major share of bookings on the European market. The recorded prices in our sample range from 27 to 2581 Euros, with a weekly average of 364 Euros and a standard

[4] E.g., Gerardi and Shapiro (2009) document a negative relation between competition and price dispersion, whereas Borenstein and Rose (1994) document a positive relation, and Dai et al. (2014) find a non-monotonic relation.

[5] For each route, the direction of the outbound flight is randomly selected, so that for each city pair one of the two cities is the origin of the outbound flight, whereas the other one is the origin of the return flight.

[6] Flight schedules and routings within Europe are planned on a semiannual basis. Routings rarely change outside these schedules.

[7] The reasoning behind this choice is that in so-called “block space” codeshare agreements, each codesharing partner is typically granted an *ex ante* fixed amount of seats with considerable freedom to set prices independently. In line with that, prices in our data differ substantially across different codesharers: The median standard deviation among tickets sold at the same day for the same physical flight across different codesharers is 60.13 Euros. See Appendix Appendix B, and in particular Footnote [28], for details and further discussion, and for a robustness exercise where we only consider the cheapest available fare for a given physical flight across *all* codesharing offers.

[8] Of the ten markets without any data, seven are missing on the route Brussels–Leeds, where we did not find any flights offered on seven travel dates; the other three missing markets are on the routes Bordeaux–Madrid, Moscow–Budapest and Stockholm–Berlin.

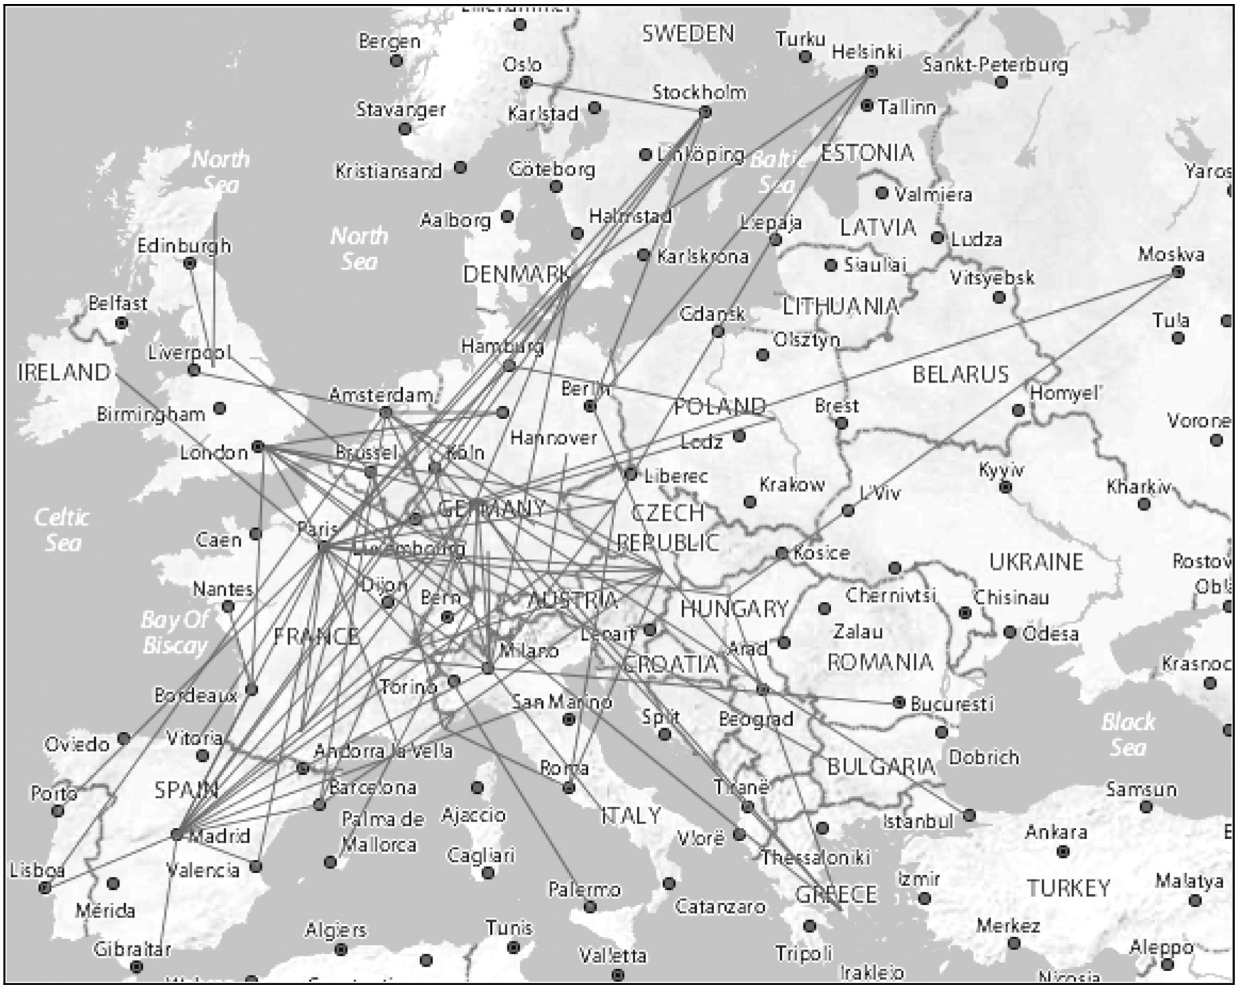

Fig. 1. Map of routes.

deviation of 466. Fig. 2 shows how average prices evolve as the scheduled departure date approaches. Prices increase from an average of 327 Euros ten weeks prior to departure to more than 500 Euros within the last week before departure.

To investigate the impact of competition on the observed pricing dynamics, we use the number of airlines that compete in a given market as our baseline measure of competition.[9] For these purposes, we treat airlines that are affiliated to each other *through equity* as single competitors (e.g., Lufthansa and Germanwings). Specifically, we treat an airline as member of a single affiliate group if that group owns more than 25% of the airline's equity (see Appendix A for details).[10] Table 1 summarizes the resulting distribution of competitors. Alternative competition measures are explored in Appendix B.

3. Impact of competition on price dynamics

We begin our analysis by estimating the intertemporal slope of prices and its sensitivity to competition. Exploiting the within-flight time series structure of our data, our baseline empirical model estimates the intertemporal slope of log prices using only the intertemporal variations of prices *within* flights. All variation in prices that is route-, time- or itinerary-specific is absorbed by fixed effects specified for the travel date, the purchase date, and both the outbound and return itinerary (using flight numbers as identifiers). Section 3.2 generalizes our baseline setup to allow for nonlinear price dynamics, and Appendix B studies the robustness of our results to alternate competition measures.

[9] In 7.9 percent of our sample, the number of airlines offering services on the outbound leg differs from the number of airlines offering services on the return leg. This may arise, for instance, if an airline does not offer services on every day of the week. In these cases, we use the mean number of competitors (rounded up).

[10] Note that our notion of affiliation through equity is distinct from the broader grouping of airlines into "airline alliances" such as "One World", "Sky Team", and "Star Alliance". Supporting the treatment of affiliation groups as single competitors, the median within-flight price dispersion between affiliated airlines is 0 Euros, whereas the median within-flight dispersion between non-affiliated codesharers (typically belonging to the same alliance) is 60.13 Euros.

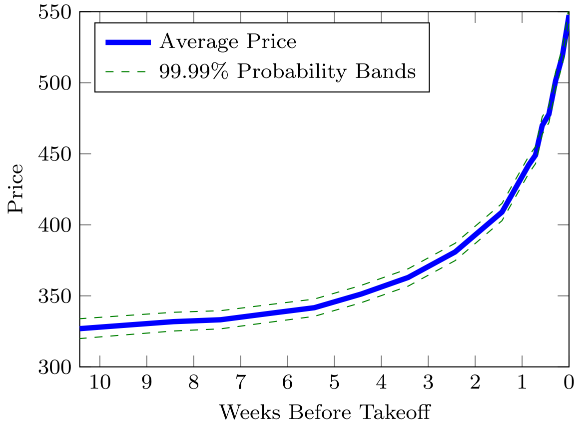

Fig. 2. Average prices (in Euros) as a function of time remaining until departure.

Table 1
Competition in the sample.

	Prices		Markets	
Competing airlines	Frequency	Percent	Frequency	Percent
1	229 196	16.17	904	24.03
2	648 242	45.73	1 696	45.08
3	275 831	19.46	657	17.46
4	185 051	13.05	382	10.15
5+	79 315	5.59	123	3.27

3.1. Baseline specification

Let $Price_{ijtd}$ denote the price for a round trip that involves the outbound itinerary i and the return itinerary j (both identified by their flight numbers), for which the outbound flight departs at date t, and which is purchased at date d. Further, let $Comp_{ijt}$ denote a vector of dummy variables that includes all five competition categories, and let $Daysleft_{td}$ denote the difference between t and d in days. As a baseline, we estimate the following equation:

$$100 \times \ln(Price_{ijtd}) = (\alpha + \beta Daysleft_{td}) \times Comp_{ijt} + \lambda_i + \mu_j + \nu_t + \xi_d + \varepsilon_{ijtd}, \tag{1}$$

where we treat λ_i, μ_j, ν_t, and ξ_d as fixed effects.[11] Here, α is a vector of competition-specific constants and β is the coefficient-vector on $Comp_{ijt} \times Daysleft_{td}$. Note that λ_i and μ_j both nest a complete set of route specific fixed effects since flight numbers uniquely pin down the city of departure and arrival. Together the specified set of fixed effects absorbs all itinerary-related effects such as departure time or length of flight; all route characteristics such as alternative means of transportation and city fixed effects; and all time-related effects of either the travel or purchase date.

The impact of competition on the observed pricing dynamics is captured by our estimate for β. Table 2 reports the estimated coefficients. Our estimates for the corresponding standard errors are adjusted for clustering at the market level. All reported coefficients are statistically significant at the 0.1 percent level.

For all competition categories prices increase as the scheduled travel date approaches. However, the intertemporal slope at which prices increase is substantially smaller for routes with more competitors. While in monopoly markets prices increase by an average of 1.31 percent with every day that a customer waits to book, this slope is reduced to 1.19 percent in duopoly markets and 0.90 percent in markets with 5+ competitors. These *differences* in slopes are statistical significant for all pairwise combinations, except for the duopoly vs $Comp_{ijt} = 3$ slopes for which we cannot reject equality (see Table 3 for details.)

Comment on endogeneity

A potential concern may be that competition is possibly endogenous. To the extend that our route fixed effects control for the effect of unobserved heterogeneity on average prices, we only need to worry about heterogeneity that affects prices *differentially* across purchase dates.

In particular, one might worry that routes with more scope for price discrimination are more profitable and attract more competitors. If such a correlation exists, it seems likely that it would result in a *positive* link between competition and the

[11] Because our sampling is weekly for all but the last week before departure, a daily specification of ξ_d would partially absorb the impact of ($Daysleft_{td} < 7$) on prices; ξ_d is therefore coded at a weekly frequency.

Table 2
Baseline estimation of intertemporal slopes.

	Coefficients	Clustered Std. Errors
$(Comp_{ijt} = 1) \times Daysleft_{td}$	−1.31	0.09
$(Comp_{ijt} = 2) \times Daysleft_{td}$	−1.19	0.09
$(Comp_{ijt} = 3) \times Daysleft_{td}$	−1.15	0.09
$(Comp_{ijt} = 4) \times Daysleft_{td}$	−1.07	0.09
$(Comp_{ijt} = 5+) \times Daysleft_{td}$	−0.90	0.10
Observations	1 417 628	
R-squared (adj.)	0.58	

Notes: The dependent variable is $100 \times \ln(Price_{ijtd})$. The estimation controls for levels of $Comp_{ijt}$ and fixed effects λ_i, μ_j, ν_t and ξ_d. Standard errors are clustered at the market level. All reported coefficients are significant at the 0.1 percent level.

Table 3
Significance of pairwise difference in slopes.

	$(Comp_{ijt} = m) \times Daysleft_{td}$			
	$m = 2$	$m = 3$	$m = 4$	$m = 5+$
$(Comp_{ijt} = 1) \times Daysleft_{td}$	15.9***	21.1***	38.0***	75.1***
$(Comp_{ijt} = 2) \times Daysleft_{td}$		1.6	12.3***	45.6***
$(Comp_{ijt} = 3) \times Daysleft_{td}$			5.1**	31.8***
$(Comp_{ijt} = 4) \times Daysleft_{td}$				12.8***

Notes: The table shows F-statistics from pairwise Wald tests for equality of $(Comp_{ijt} = m) \times Daysleft_{td}$ with $(Comp_{ijt} = n) \times Daysleft_{td}$. Degrees of freedom in the denominator are 3760. ***, **, * indicate significance at the 1%, 5%, and 10% level.

intertemporal slope. We therefore suspect that—if at all—our results *underestimate* the true impact of competition on the intertemporal slope.

To get a sense about how large such bias might be, we have computed the correlation between competition and 8 different proxies for customer heterogeneity (further described below). All of these correlations are essentially zero (see Fig. 4). While we cannot fully rule out the existence of confounding factors, we therefore consider it likely that they are of second order compared to the direct impact of competition on pricing dynamics that we have documented.

3.2. Nonlinear advance purchase discounts

Our analysis so far has documented the impact of competition on the *average* intertemporal slope. In this section, we now generalize our baseline specification to disaggregate the average impact over time. To this end, we estimate a variant of (1) where we specify $Daysleft_{td}$ as a vector of dummy variables including all values in the support of $Daysleft_{td}$. Note that this specification imposes virtually no parametric restrictions on the evolution of prices over time. Our empirical model can thus effectively be written as:

$$100 \times \ln(Price_{ijtd}) = \phi(Daysleft_{td}, Comp_{ijt}) + \lambda_i + \mu_j + \nu_t + \xi_d + \varepsilon_{ijtd}, \tag{2}$$

where $\phi : \{0, \ldots, 73\} \times \{1, \ldots, 5+\} \to \mathbb{R}$ is an arbitrary function determined by the data.

Fig. 3 shows the estimated impact of competition on the intertemporal slope. In the left panel, the predicted price is normalized relative to the log-price predicted for the day of departure.[12] The y-axis hence reflects the estimated advance purchase discount relative to the price charged on the day of departure. It can be seen that, although nonlinear, the slopes are again monotonically decreasing in the number of competitors. That is, the relative discount for booking a flight in advance is less pronounced on routes that are served by a larger number of competitors, reinstating the conclusion drawn from our baseline estimation.

Taking a closer look at the identified pricing dynamics, it can further be seen that prices are increasing at similar slopes until about five weeks before takeoff. Only in the last five weeks, prices in less competitive routes have a significantly steeper slope than prices in more competitive routes. To illustrate this further, the right panel re-plots the predicted price normalized relative to 38 days before departure. Until about five weeks before takeoff, log prices increase along a common trajectory across all competitive environments, showing an increase of about 0.44 percent for each day a customer waits to book. This translates into an overall discount of approximately 14 percent for purchasing tickets ten weeks before departure compared to five weeks before departure.

[12] Ticket prices for same day departures are crawled between 3.15am and 3.45am. The earliest recorded flights depart at 4.00am.

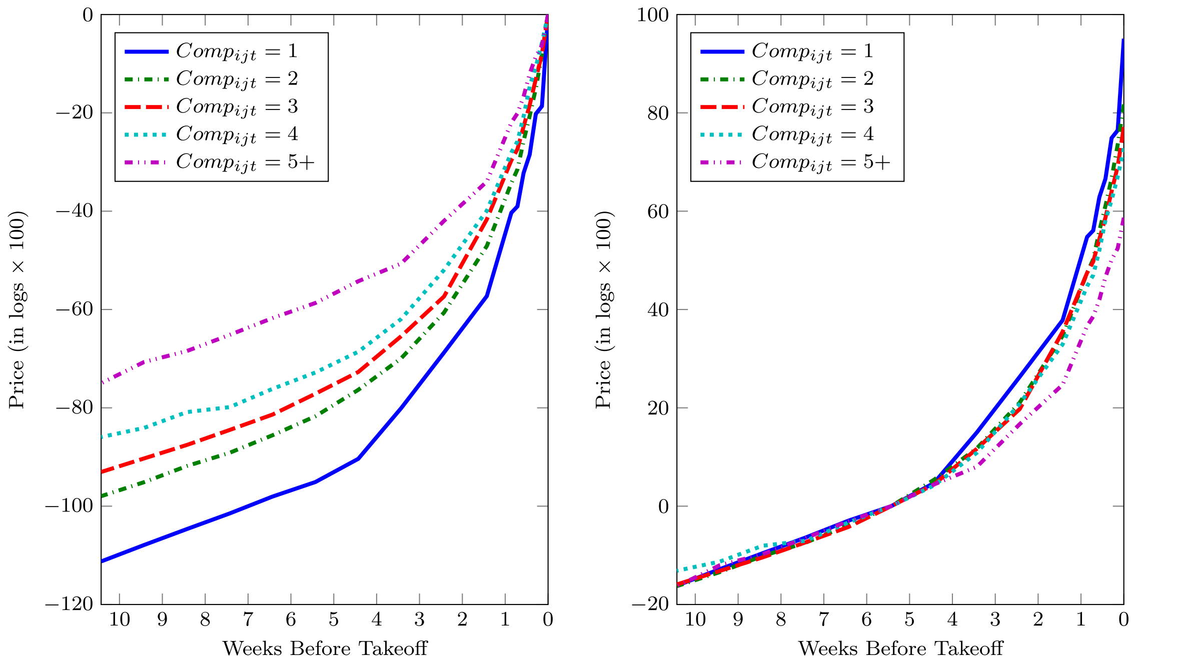

Fig. 3. Nonlinear estimation of pricing dynamics. Notes: The figure shows the estimated impact of $Daysleft_{td}$ and $Comp_{ijt}$ on $100 \times \ln(Price_{ijtd})$. The estimation controls for levels of $Comp_{ijt}$ and fixed effects λ_i, μ_j, ν_t and ξ_d. Prices are normalized relative to the predicted log-price on the day of departure (left panel) and 38 days before departure (right panel).

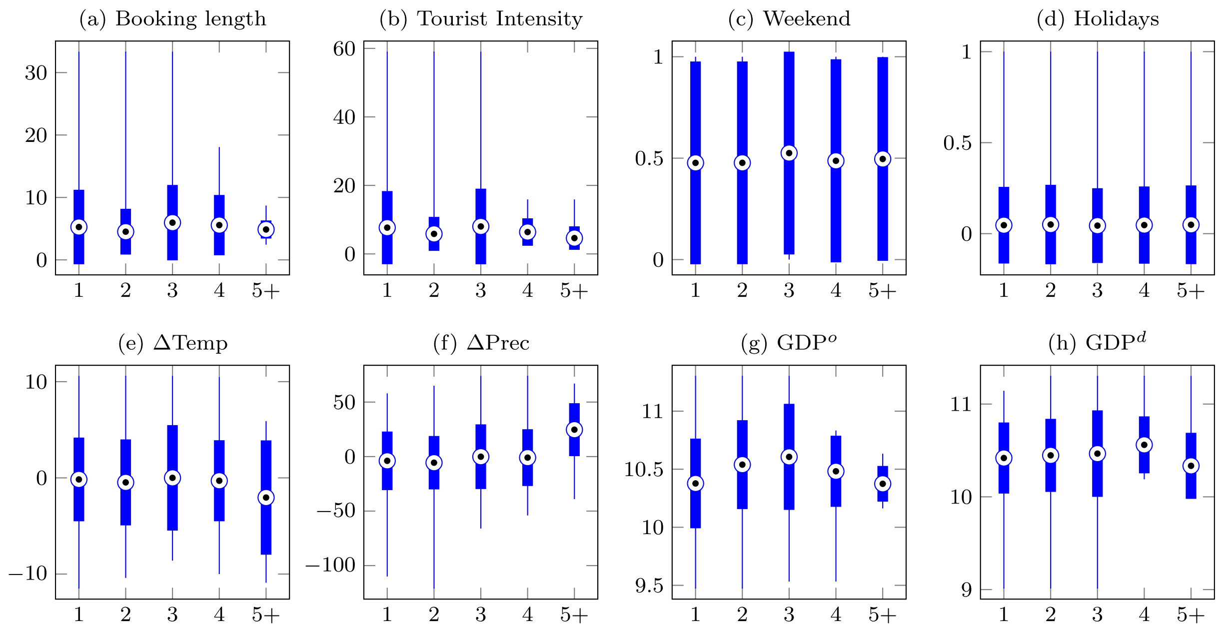

Fig. 4. Summary of tourist indicators. Notes: Means and $\pm$ 1 standard deviation bands are indicated by the circled dots and boxes. Supports are indicated by the extending lines.

Starting about five weeks before departure, prices increase significantly faster, with prices on the least competitive environments increasing the fastest. On monopoly routes, customers pay a premium of 159 percent for purchasing their ticket on the day of departure rather than five weeks in advance.[13] This premium is reduced to 126 percent in duopoly markets,

[13] I.e., $100 \cdot (e^{y_0/100} - 1) = 159$ percent, where $y_0 = \phi(0, Comp_{ijt} = 1) - \phi(38, Comp_{ijt} = 1)$ is the intercept with the (right) y-axis in the right panel of Fig. 3.

Table 4
Significance of pairwise difference in same day departure premiums.

	$\Delta\phi(\cdot, Comp_{ijt} = m)$			
	$m = 2$	$m = 3$	$m = 4$	$m = 5+$
$\Delta\phi(\cdot, Comp_{ijt} = 1)$	40.5***	62.1***	69.4***	142.6***
$\Delta\phi(\cdot, Comp_{ijt} = 2)$		6.2**	14.7***	70.7***
$\Delta\phi(\cdot, Comp_{ijt} = 3)$			3.1*	41.3***
$\Delta\phi(\cdot, Comp_{ijt} = 4)$				19.5***

Notes: The table shows F-statistics from pairwise Wald tests for equality of $\phi(0, m) - \phi(38, m)$ with $\phi(0, n) - \phi(38, n)$. Degrees of freedom in the denominator are 3760. ***, **, * indicate significance at the 1%, 5%, and 10% level.

and is further reduced to 116 percent in markets with three competitors, 107 percent in markets with four competitors, and 80 percent in markets with 5+ competitors. These premiums are all statistically different from one another (see Table 4 for details.)

In sum, the nonlinear estimation reinforces the clear pattern of dynamic oligopoly pricing revealed by our baseline specification: While airlines offer substantial advance purchase discounts across all market structures, the magnitude of these discounts is highly sensitive to competition.

4. Impact of customer heterogeneity on price dynamics

A number of channels may link rising airline prices to heterogeneous demand elasticities. For instance, a common view is that demand elasticities are correlated with the ability or willingness to book in advance, allowing airlines to discriminate against late-booking customers.

This section presents results for the effects of customer heterogeneity on pricing dynamics. To measure variations of customer heterogeneity across different markets, we construct a composite index of a variety of indicators on the composition of leisure and business travelers across markets. Heuristically, our measure of customer heterogeneity corresponds to the predicted impact of these indicators on the intertemporal slope *in monopoly markets.* The main identifying assumption is that monopolists are able to exploit higher degrees of customer heterogeneity by selling tickets cheap to price sensitive customers that book early, while charging higher prices for less sensitive customers that book late. Using this measure, we then explore how customer heterogeneity affects the impact of competition on the intertemporal slope.

4.1. Measuring customer heterogeneity

In the airline industry the co-existence between tourists (and other leisure travelers) and business travelers is arguably the largest source of heterogeneity in the customer base. Our approach to evaluate customer heterogeneity aims to capture such variations in the intensity of leisure and business travelers using a variety of market indicators.

4.1.1. List of market indicators

The first two indicators are based on the importance of tourism at the destination.[14]

(a) Booking Length: The indicator measures the average number of nights per visit booked in hotels and similar residencies at the destination, proxying for how attractive the destination is for tourists.
(b) Tourist Intensity: The indicator measures the number of nights booked in hotels and similar residencies relative to the population at the destination (in thousands), measuring the intensity of the tourist industry at the destination.

Confirming our intuition, these measures are largest in pro-typical tourist locations such as Mallorca (Booking Length = 33.4 nights, Tourist Intensity = 59.0) and Innsbruck (24.3, 47.2), and smallest in locations that are unlikely tourist locations such as Leeds (2.1, 1.4), Manchester (2.6, 1.8), or Duesseldorf (2.2, 1.9).

The next set of indicators aims to measure the relative importance of leisure and business travelers based on travel dates.

(c) Weekend: The dummy distinguishes weekend trips (departing on Friday and returning on Sunday), which presumably attract fewer business travelers than trips that depart on Monday and return on Thursday.
(d) Holiday: The dummy indicates roundtrips over the Christmas and New Years holidays.

[14] The two indicators are based on 2011 Eurostat data at the NUTS-2 level. For large cities the NUTS-2 level typically coincides with the city level (e.g., Berlin, Lisbon or Prague), while smaller cities or less densely populated areas are typically clustered into urban areas or regions (e.g., Manchester into Greater Manchester or Aberdeen into North Eastern Scotland).

The next two indicators are based on an idea by Brueckner et al., 1992 who argue that a good proxy for the share of leisure travelers on a route are temperature differentials between the destination and origin.[15]

(e) **Δ**Temp: The indicator measures the temperature differential between destination and origin during the month of travel.
(f) **Δ**Prec: The indicator measure precipitation differentials between destination and origin during the month of travel.

Finally, we use regional GDP data as another proxy for business activity.[16]

(g) GDP$^{\boldsymbol{o}}$: GDP at the origin of the itinerary.
(h) GDP$^{\boldsymbol{d}}$: GDP at the destination of the itinerary.

Fig. 4 gives an overview over all eight indicators, showing their means (marked by the circled dots), standard deviation bands (marked by the boxes), and supports (marked by the lines) conditional on the number of competitors.

4.1.2. Construction of the composite indicator for customer heterogeneity

Let Z_{ijt} define a subset of the eight market indicators. In our main specification, Z_{ijt} includes all eight indicators, but we also explore the cases where Z_{ijt} includes only a single indicator at a time.

Our measure of customer heterogeneity aims to identify variations in the customer base, as indicated by Z_{ijt}, that induce a *monopolist* airline to more strongly differentiate prices between early and late booking travelers. To this aim, we estimate the following first-stage projection *in the subsample of monopoly markets*:

$$\ln\left(Price_{ijtd}\right) = \alpha_0 + \alpha_z Z_{ijt} + \alpha_d Daysleft_{td} + \beta Z_{ijt} \times Daysleft_{td} + \lambda_i + \mu_j + \nu_t + \xi_d + \epsilon_{ijtd} \quad \text{if } Comp_{ijt} = 1.$$

The resulting measure of customer heterogeneity is defined as

$$Het_{ijt} = -\beta Z_{ijt},$$

standardized to have zero mean and unit variance. Intuitively, Het_{ijt} is a linear combination of Z_{ijt}, constructed so that higher realizations of Het_{ijt} *steepen* the intertemporal slope in monopoly markets. (In the case where Z_{ijt} includes a single indicator, the standardized measure Het_{ijt} is simply the original indicator Z_{ijt} normalized by $-\text{sign}\{\beta\}$).

The motivation behind our approach of measuring customer heterogeneity is that one would expect that—absent competition—the *scope* to raise prices over time is positively correlated with the customer heterogeneity on a given route. Using $(-\beta)$ to weigh the various market indicators, we are thus likely to identify variations in the composition of leisure and business travelers that are associated with a more heterogeneous customer base.

Our baseline index, which is based on all eight indicators, is given by:

$$\begin{aligned} Het_{ijt} \quad = \quad & .5594 \times \text{Booking Length}_{ij} - .1841 \times \text{Tourist Intensity}_{ij} - .8011 \times \text{Weekend}_t - .1397 \times \text{Holidays}_t \\ & + .4650 \times \Delta\text{Temp}_{ijt} - .0440 \times \Delta\text{Prec}_{ijt} + .0212 \times \text{GDP}^o_{ij} + .0627 \times \text{GDP}^d_{ij}. \end{aligned} \tag{3}$$

Here all coefficients are denominated in units of standard deviations of the corresponding indicator. The estimated index coefficients are indicative of the presumption that markets are most heterogeneous if they are frequented by both tourists *and* business travelers: On the one hand, if temperature differentials are positive and hotels at the destination are booked for many nights, customer heterogeneity is predicted to be high, suggesting that tourism is key for heterogeneity. On the other hand, if the destination is likely to be a pure vacation resort (indicated by a high tourist intensity) or if the booking is for weekends or holidays, customer heterogeneity is predicted to be low, presumably due to a lack of business travelers on such itineraries.

4.2. Baseline specification

We are now ready to investigate how customer heterogeneity impacts the sensitivity of the intertemporal slope to competition. As a baseline, we estimate the following model:

$$100 \times \ln(Price_{ijtd}) = (\alpha + \beta Daysleft_{td}) \times X_{ijt} + \lambda_i + \mu_j + \nu_t + \xi_d + \epsilon_{ijtd}, \tag{4}$$

with

$$X_{ijt} = (1,\ N_{ijt},\ Het_{ij},\ N_{ijt} \times Het_{ijt}),$$

where N_{ijt} is the number of competitors, Het_{ijt} is the baseline customer heterogeneity measure (based on all eight indicators, and standardized to have zero mean and unit variance), and λ_i, μ_j, ν_t and ξ_d are treated as fixed effects.

Table 5 reports the estimated "slope" coefficients (β) and their standard errors (clustered at the market level). First, prices are again increasing over time with a slope that—on average—flattens in more competitive markets (recall that Het_{ijt} is normalized to have zero mean). This confirms the results in Section 3.

[15] See also Stavins (2001) and Goolsbee and Syverson (2008), who have used similar temperature-based proxies to measure for tourist intensities.
[16] We use GDP data provided by the European commission at the NUTS-2 level. See also Footnote 14.

Table 5
Impact of customer heterogeneity on pricing dynamics.

	Coefficients	Clustered Std. Errors
(a) Estimated Slope Coefficients		
$Daysleft_{td}$	−1.19	0.10
$Daysleft_{td} \times N_{ijt}$	0.10	0.01
$Daysleft_{td} \times Het_{ij}$	−0.33	0.03
$Daysleft_{td} \times N_{ijt} \times Het_{ij}$	0.09	0.01
(b) Sensitivity and Overall Slope		
Low heterogeneity markets		
sensitivity to competition	0.01	0.02
implied monopoly slope	−0.85	0.10
implied competitive slope ($N_{ijt} = 5$)	−0.81	0.11
High heterogeneity markets		
sensitivity to competition	0.19	0.01
implied monopoly slope	−1.33	0.10
implied competitive slope ($N_{ijt} = 5$)	−0.56	0.10
Observations	1 215 441	
R-squared (adj.)	0.57	

Notes: The dependent variable is $100 \times \ln(Price_{ijtd})$. The estimation controls for levels of $Daysleft_{td}$, N_{ijt}, Het_{ij} and $N_{ijt} \times Het_{ij}$, and fixed effects λ_i, μ_j, ν_t and ξ_d. Het_{ij} is standardized to have zero mean and unit variance. Standard errors are clustered at the market level. Low and high heterogeneity markets in Panel (b) are defined as markets with $Het_{ij} = \pm 1$ standard deviation relative to the mean.

Next, it can be seen that the intertemporal slope steepens in customer heterogeneity as long as the number of competitors is small ($N_{ijt} \leq 3$). Recalling that Het_{ijt} was constructed to steepen the slope in monopoly markets, this is expected at a qualitative level. Quantitatively, we find that in monopoly markets, a one standard deviation increase in customer heterogeneity steepens the daily slope by 0.24 percentage points, which is more than twice as much as the impact of an additional competitor.

We now discuss the main question: How does customer heterogeneity affect the *sensitivity* of the slope to competition. The main result is that the estimated slope coefficient on $N_{ijt} \times Het_{ij}$ is statistically significant and positive. That is, competition flattens the slope *more* in markets where customer heterogeneity is high. Panel (b) of Table 5 explores the economic significance of the result by looking at markets where customer heterogeneity is one standard deviation above and below its mean. Recalling that Het_{ijt} is normalized to have unit variance, the impact of competition on the slope in these markets is given by,

$$\left.\frac{\partial^2 \ln(Price)}{\partial Daysleft_{td}\, \partial N_{ijt}}\right|_{Het_{ij}=\pm 1} = \beta_N \pm \beta_{N\times Het}.$$

Panel (b) uncovers a stark disparity: While increasing competition flattens the intertemporal slope in high heterogeneity markets by 0.19 percentage points per day, competition has virtually no effect on the slope in low heterogeneity markets (the estimated impact of 0.01 is statistically insignificant).

To further illustrate the magnitudes of these effects, we explicitly compute the slopes in monopoly markets and markets with 5+ competitors. In low heterogeneity markets, the implied monopoly slope of 0.85 percent per day barely differs from the competitive slope of 0.81 percent per day. By contrast, in high heterogeneity markets, the monopoly slope of 1.33 percent per day is more than twice as large as the competitive slope of 0.56 percent.

4.3. Robustness specifications

We conclude this section by demonstrating robustness of the findings to a variety of alternate specifications

4.3.1. Alternative customer heterogeneity measures

In our baseline, we use all eight market indicators to construct Het_{ijt}. Table 6 reports results for the cases where Het_{ijt} is based on a single indicator at a time.

Overall, the baseline findings are confirmed by all but one of these alternatives. Specifically, with the exception of column (8), we again find that the impact of competition on the slope is significantly dampened in low heterogeneity markets; i.e., the sensitivity of the slope in N_{ijt}, reported in Panel (b), is smaller in low heterogeneity markets than in high heterogeneity markets.[17] The one exception is the case where Het_{ijt} is based on GDP at the destination, in which case higher customer

[17] In four of the specifications (booking length, tourist intensity, weekend, and ΔPrec), the sensitivity in low heterogeneity markets is virtually zero (and statistically insignificant), mirroring the baseline case. For Holidays, ΔTemp, and GDP^o, the effects are qualitatively the same as in the baseline, but

Table 6
Alternative customer heterogeneity measures.

Customer Heterogeneity Measure:	(1) Book. Len.	(2) Tour. Int.	(3) Weekend	(4) Holidays	(5) ΔTemp	(6) ΔPrec	(7) GDP^o	(8) GDP^d
	(a) Estimated Slope Coefficients							
$Daysleft_{td}$	−1.41	−1.37	−1.00	−1.33	−1.35	−1.33	−1.43	−1.37
	(0.10)	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.10)
$Daysleft_{td} \times N_{ijt}$	0.07	0.07	0.08	0.08	0.07	0.07	0.09	0.07
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$Daysleft_{td} \times Het_{ijt}$	−0.13	−0.17	−0.26	−0.09	−0.02	−0.06	−0.11	0.07
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)
$Daysleft_{td} \times N_{ijt} \times Het_{ijt}$	0.06	0.08	0.06	0.02	0.02	0.04	0.03	−0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
	(b) Implied Sensitivity to Competition and Overall Slope							
Low heterogeneity markets								
sensitivity	0.01	−0.01	0.02	0.06	0.05	0.03	0.06	0.10
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.020)	(0.02)
monopolies	−1.27	−1.21	−0.73	−1.18	−1.28	−1.25	−1.26	−1.35
	(0.10)	(0.10)	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)	(0.10)
competitive ($N_{ijt} = 5$)	−1.23	−1.25	−0.67	−0.94	−1.07	−1.14	−1.02	−0.97
	(0.10)	(0.10)	(0.10)	(0.09)	(0.09)	(0.10)	(0.10)	(0.11)
High heterogeneity markets								
sensitivity	0.13	0.15	0.15	0.10	0.10	0.11	0.13	0.05
	(0.01)	(0.01)	(0.013)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
monopolies	−1.42	−1.38	−1.11	−1.31	−1.28	−1.30	−1.41	−1.25
	(0.10)	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.10)
competitive ($N_{ijt} = 5$)	−0.90	−0.76)	−0.53	−0.92	−0.90	−0.88	−0.90	−1.05
	(0.10)	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)	(0.10)	(0.10)
Observations	1 255 118	1 286 680	1 417 628	1 417 628	1 417 628	1 417 628	1 377 951	1 286 680
R-squared (adj.)	0.57	0.57	0.58	0.58	0.58	0.58	0.58	0.57

Notes: The dependent variable is $100 \times \ln(Price_{ijtd})$. The estimation controls for levels of $Daysleft_{td}$, N_{ijt}, Het_{ijt} and $N_{ijt} \times Het_{ijt}$, and fixed effects λ_i, μ_j, ν_t and ξ_d. Het_{ijt} is standarized to have zero mean and unit variance. Clustered standard errors (at the market level) are reported in parentheses. Low and high heterogeneity markets in Panel (b) are defined as markets with $Het_{ijt} = \pm 1$ standard deviation relative to the mean.

heterogeneity reduces the impact of competition. Given that GDP contributes comparably little to our baseline index (cf., Eq. (3)), this may reflect that GDP is not a good proxy for customer heterogeneity by itself.

4.3.2. Last five weeks before departure

In Section 3.2 we document that most of the increase in ticket prices (and virtually all the impact of competition on the slope) is confined to the last five weeks before departure. It is therefore natural to ask inasmuch our findings regarding customer heterogeneity apply to those last five weeks.

Table 7 reports for the case where we re-estimate 4 for the five week window preceding departure. The results mirror the baseline case qualitatively, but are quantitatively amplified. While the sensitivity to competition in low heterogeneity markets continues to be statistically insignificant, the sensitivity in high heterogeneity markets doubles to 0.38 percentage points per day. This reflects that the overall slope is approximately twice as steep in the last five weeks compared to the average slope over the last ten weeks: In low heterogeneity markets, the monopoly slope amounts to 1.82 percent per day, whereas the competitive slope ($N_{ijt} = 5+$) amounts to 1.95. In high heterogeneity markets, the monopoly slope amounts to 2.86 percent per day, whereas the competitive slope amounts to 1.36 percent.

4.3.3. Competition-specific impact of heterogeneity

We now generalize specification (4) to allow for a competition-specific impact of customer heterogeneity. Specifically, we estimate the following model:

$$\ln(Price_{ijtd}) = (\alpha + \beta Daysleft_{td}) \times X_{ijt} + \lambda_i + \mu_j + \nu_t + \xi_d + \epsilon_{ijtd}, \tag{5}$$

with

$$X_{ijt} = (Comp_{ijt}, Comp_{ijt} \times Het_{ijt}),$$

quantitatively smaller. This is not surprising given that these are likely to be relatively more noisy measures of customer heterogeneity and therefore subject to greater attenuation bias.

Table 7
Last five weeks before departure.

	Coefficients	Clustered Std. Errors
(a) Estimated Slope Coefficients		
$Daysleft_{td}$	−2.52	0.09
$Daysleft_{td} \times N_{ijt}$	0.17	0.02
$Daysleft_{td} \times Het_{ij}$	−0.73	0.06
$Daysleft_{td} \times N_{ijt} \times Het_{ij}$	0.20	0.02
(b) Sensitivity and Overall Slope		
Low heterogeneity markets		
sensitivity to competition	−0.03	0.03
implied monopoly slope	−1.82	0.09
implied competitive slope ($N_{ijt} = 5$)	−1.94	0.11
High heterogeneity markets		
sensitivity to competition	0.38	0.03
implied monopoly slope	−2.86	0.09
implied competitive slope ($N_{ijt} = 5$)	−1.36	0.10
Observations	814 479	
R-squared (adj.)	0.58	

Notes: The dependent variable is $100 \times \ln(Price_{ijtd})$. The estimation controls for levels of $Daysleft_{td}$, N_{ijt}, Het_{ij} and $N_{ijt} \times Het_{ij}$, and fixed effects λ_i, μ_j, ν_t and ξ_d. Het_{ij} is standardized to have zero mean and unit variance. Standard errors are clustered at the market level. Low and high heterogeneity markets in Panel (b) are defined as markets with $Het_{ij} = \pm 1$ standard deviation relative to the mean.

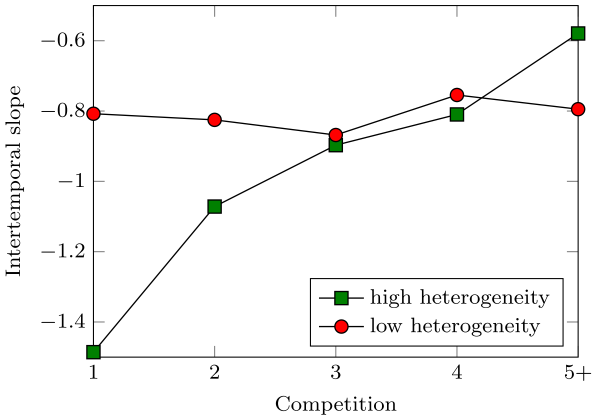

Fig. 5. Predicted intertemporal slope as a function of competition and customer heterogeneity.Notes: The intertemporal slopes are denominated in daily price changes in percent. High and low heterogeneity lines correspond to $Het_{ijt} = \pm 1$ standard deviation.

where $Comp_{ijt}$ is now treated as a categorical variable as in Section 3. To facilitate the interpretation, Het_{ijt} is now standardized to have zero mean *within* each competition category (and unit variance across the full sample as before); i.e.,

$$Het_{ijt}^{\text{standardized}} = \frac{Het_{ijt} - \mathbb{E}[Het_{ijt}|Comp_{ijt}]}{\text{s.d.}[Het_{ijt}]}.$$

Table 8 reports the estimated slope coefficients (β) and standard errors (clustered at the market level). Confirming our baseline findings, the impact of competition on the intertemporal slope is again amplified in more heterogeneous markets. Fig. 5 visualizes this by plotting the predicted intertemporal slope as a function of the number of competitors in high and low heterogeneity markets ($Het_{ijt} = \pm 1$ standard deviation). In line with our baseline findings, competition significantly flattens the intertemporal slope in high heterogeneity markets, while it has little effect in low heterogeneity markets.

To provide a direct comparison with the baseline specification, Panel (b) of Table 8 reports the average sensitivity to competition, corresponding to the average slope-to-competition gradient in Fig. 5:

$$\frac{1}{4}\sum_{N=2}^{5+}\left\{\left[\beta_{Comp=N} - \beta_{Comp=N-1}\right] \pm \left[\beta_{(Comp=N)\times Het} - \beta_{(Comp=N-1)\times Het}\right]\right\}.$$

Not surprisingly, given the approximately linear relations in Fig. 5, these average sensitivities are very similar to the ones predicted by the baseline model.

Table 8
Competition-specific impact of heterogeneity.

	Coefficients	Clustered Std. Errors
(a) Estimated Slope Coefficients		
$(Comp_{ijt} = 1) \times Dayslef t_{td}$	−1.15	0.10
$(Comp_{ijt} = 2) \times Dayslef t_{td}$	−0.95	0.10
$(Comp_{ijt} = 3) \times Dayslef t_{td}$	−0.88	0.10
$(Comp_{ijt} = 4) \times Dayslef t_{td}$	−0.78	0.10
$(Comp_{ijt} = 5+) \times Dayslef t_{td}$	−0.69	0.11
$(Comp_{ijt} = 1) \times Het_{ijt} \times Dayslef t_{td}$	−0.34	0.02
$(Comp_{ijt} = 2) \times Het_{ijt} \times Dayslef t_{td}$	−0.12	0.02
$(Comp_{ijt} = 3) \times Het_{ijt} \times Dayslef t_{td}$	−0.01	0.02
$(Comp_{ijt} = 4) \times Het_{ijt} \times Dayslef t_{td}$	−0.03	0.03
$(Comp_{ijt} = 5+) \times Het_{ijt} \times Dayslef t_{td}$	0.11	0.05
(b) Sensitivity and Overall Slope		
Low heterogeneity markets		
average sensitivity to competition	0.00	0.02
implied monopoly slope	−.81	0.10
implied competitive slope $(Comp_{ijt} = 5+)$	−.80	0.12
High heterogeneity markets		
average sensitivity to competition	0.23	0.02
implied monopoly slope	−1.49	0.10
implied competitive slope $(Comp_{ijt} = 5+)$	−0.58	0.11
Observations	1 215 441	
R-squared (adj.)	0.57	

Notes: The dependent variable is $100 \times \ln(Price_{ijtd})$. The estimation controls for levels of $Comp_{ijt}$, $Comp_{ijt} \times Het_{ijt}$, and fixed effects λ_i, μ_j, ν_t and ξ_d. Het_{ijt} is standardized to have zero mean *within* each competition category, and unit variance in the full sample. Standard errors are clustered at the market level. Low and high heterogeneity markets in Panel (b) are defined as markets with $Het_{ijt} = \pm 1$ standard deviation relative to the mean in the corresponding competition category.

5. Decomposing price dispersion

In this section, we explore how much the documented pricing dynamics contribute to overall price dispersion for airline tickets, and how much of the *variation* in pricing dynamics across markets is accounted for by the documented impact of competition and customer heterogeneity.

5.1. Within-flight vs. between-flight dispersion

The most natural approach to measuring price dispersion for our purpose is to focus on the intertemporal dispersion of prices *within* a given physical price. This within-flight dispersion reflects systematic advance purchase discounts, as documented in the previous two sections, and further reflects unpredicted residual fluctuations in ticket prices across different purchase dates. Before exploring the relative contribution of these two factors to the total within-flight dispersion, we briefly contrast our measure of within-flight dispersion with more broadly defined dispersion measures adopted by the previous literature.

In particular, due to data limitations, previous literature often focuses on dispersion measures for airfares that do not disentangle intertemporal dispersion within flights from the dispersion in prices across different flights (pooling both, different itineraries and different travel dates). To relate to this literature, we exploit the double-panel structure of our data to disentangle the intertemporal price dispersion for a given flight from various types of cross-flight dispersion. Following the literature, we use the Gini coefficient to measure the dispersion within a given set of prices. To make our measures comparable to the previous literature, we compute all Gini coefficients after up-sampling observations with more than seven days left to departure in order to compensate for the increased sampling frequency in the last week.

Summary of dispersion measures

Panel (a) of Table 9 summarizes the average Gini coefficients in our data set. Column 1 reports the average price dispersion across all flights (and dates of ticket offers) offered by a given competitor on a given route. As in the previous literature, this pools together different itineraries (within the same route) and different flights across different travel dates. We find an average dispersion of 64 percent of the average offer.[18] Column 4 contrast this with the intertemporal dispersion within a

[18] The Gini coefficient corresponds to half the expected price difference in terms of the average price. A Gini coefficient of 0.32 therefore represents an expected absolute difference between two randomly selected prices of 64 percent of the average price.

Table 9
Advance purchase discounts and price dispersion.

	Data				Projected Data	
Dispersion Category:	Routes (1)	Markets (2)	Itineraries (3)	Flights (4)	Flights (5)	Flights (6)
	(a) Price dispersion within affiliate group × dispersion category					
	0.32 (0.13)	0.28 (0.14)	0.17 (0.09)	0.12 (0.09)	0.11 (0.04)	0.10 (0.03)
	(b) Competition Effects on G^{lodd}_{ijt}					
N_{ijt}	0.11 (0.01)	0.30 (0.03)	−0.22 (0.03)	−0.45 (0.07)	−0.54 (0.04)	−0.52 (0.04)
Observations	1 215 148	1 215 148	1 215 148	1 207 253	1 215 148	1 215 148
R-squared (adj.)	0.62	0.32	0.17	0.06	0.07	0.06

Notes: Panel (a) reports average Gini coefficients and standard deviations (in parenthesis) within routes (Column 1), markets (Column 2), itineraries (Column 3) and flights (Column 4). In all cases, price dispersion is within-airlines. Columns 5 and 6 of Panel (a) project within-flight Gini coefficients based on the pricing dynamics identified by the baseline regression (1) and the nonlinear regression (2). Panel (b) estimates the effect of competition on the Gini log-odds ratio, $G^{lodd}_{ijt} = \ln[G_{ijt}/(1 - G_{ijt}]$, in the data and in the projected sample. The estimations control for fixed effects ν_t and ξ_d and the controls listed in Footnote [19]. Clustered standard errors (at the market level) are reported in parentheses. All reported coefficients are significant at the 0.1 percent level.

given flight (defined by a particular itinerary at a particular travel date). We see that the within-flight dispersion contributes about one third to the overall price dispersion within routes.

Columns 2 and 3 further dissect the difference between these two dispersion measures by looking at the two intermediate cases where we pool flights across only itineraries (within markets) and across only travel dates (within itineraries). Evidently, prices vary more in flight characteristics (e.g., early morning vs. late evening flights) than they do across different travel dates (e.g., a specific Monday-Friday roundtrip vs. the exact same flight combination one week later).

Impact of competition on price dispersion

The pricing dynamics documented in this paper suggest a negative impact of competition on within-flight price dispersion. By contrast, using the broader within-route dispersion measure, previous literature reached ambiguous conclusions regarding the relation between competition and price dispersion, ranging from positive (Borenstein and Rose, 1994), over no clear relation (Hayes and Ross, 1998), to negative (Gerardi and Shapiro, 2009).

In an attempt to reconcile our findings with the previous literature, we now explore the impact of competition on the four different measures of price dispersion. We do no longer include itinerary specific fixed effects, since these would absorb any dispersion on the itinerary and route level. Instead we include a large number of control variables X_{ijt} that were previously nested into λ_i and μ_j.[19] Moreover, since the Gini coefficient is bounded between zero and one, we follow the previous literature and transform it into an unbounded statistic, using instead the Gini log-odds ratio $G^{lodd}_{ijt} = \ln[G_{ijt}/(1 - G_{ijt})]$.

We estimate the following empirical model:

$$G^{lodd}_{ijt} = \alpha + \beta \times N_{ijt} + \gamma \times X_{ijt} + \nu_t + \xi_d + \epsilon_{ijt}, \tag{6}$$

where N_{ijt} is the number of competitors, X_{ijt} contains our set of controls, and ν_t and ξ_d are vectors of fixed effects for the travel date and the date of the price offer. The estimated coefficients are reported in Panel (b) of Table 9. All reported coefficients are significant at the 0.1 percent level (using standard errors that are clustered at the market level).

In line with the pricing dynamics identified in Section 3, competition has a negative impact on the within-flight price dispersion (Column 4).[20] However, once we consider broader measures of price dispersion the impact is diluted (Column 3) or even overturned (Columns 1 and 2). A potential explanation is that in more competitive environments, airlines offer a larger number of connections to compete in departure dates and times, and that there is a meaningful dispersion of prices across these connections.[21] This suggests that the seemingly contradicting findings in the earlier literature may be driven by confounding different dimensions of dispersion.

5.2. Contribution of advance purchase discounts to within-flight dispersion

We now return to the question of how much of the within-flight dispersion is due to the systematic advance purchase discounts documented in Section 3 as opposed to unpredicted residual fluctuations in prices across different purchase dates.

[20] See also Gaggero and Piga (2011).

[21] The positive relation between competition and cross-flight dispersion is also consistent with arguments by Borenstein (1985) and Holmes (1989), who argue that the dispersion *across* flights may increase in competition when consumers' cross-price elasticities between different airlines is lower than the elasticity of industry demand.

Table 10
Contribution of competition and customer heterogeneity to variation in advance purchase discounts.

Part of X_{rt}	(1)	(2)	(3)
$Comp_{rt}$	Yes	Yes	Yes
$Comp_{rt} \otimes Het_{rt}$		Yes	Yes
$Comp_{rt} \otimes (Route_r,\ Flightdate_t)$			Yes
R-squared	0.05	0.17	0.63

Notes: The dependent variable is the intertemporal slope in market rt (S_{rt}). $Comp_{rt}$, $Route_r$ and $Flightdate_t$ are vectors containing a full set of competition, route, and flight date dummies, and Het_{rt} is a vector stacking all eight customer heterogeneity indicators. Number of observations is 3 311 markets.

To answer this question, we use our estimates of the baseline model (1) and of the nonlinear model (2) to construct projected samples of counterfactual price data in which the intertemporal variation in prices is exclusively due to systematic advance purchase discounts. Columns 5 and 6 of Panel (a) in Table 9 report the within-flight dispersion for these two samples.

For the average flight the projected within-flight dispersion closely resembles the one in the data.[22] Based on the nonlinear estimation, 83 percent of the within-flight dispersion of the average flight (0.10 out of 0.12) can be attributed to systematic advance purchase discounts as seen in Fig. 3. Systematic advance purchase discounts therefore seem to be the major source of within-flight price dispersion, as opposed to residual volatility around the trend. While we do not attempt to explain these residual components, an obvious source of residual price dispersion are dynamic adjustments in prices in response to stochastic fluctuations in demand.

Closely related, one may ask how much of the documented impact of competition on within-flight price dispersion is due to its impact on the intertemporal slope relative to its potential impact on the residual volatility around the trend. One way to approach this question is re-estimate (6) in the projected samples and compare the impact coefficients with the one found earlier. The results are reported in Pabel (b) of Table 9. Comparing the competition impact in the projected samples (Columns 5 and 6) with its impact in the data (Column 4) reveals that the impact of competition on the intertemporal slope in fact overpredicts its impact on price dispersion. This suggests that the impact of competition on price dispersion is mainly driven by its impact on the intertemporal slope rather than competition-driven heteroskedasticity in the residual dispersion.[23]

5.3. Contribution of competition and customer heterogeneity to advance purchase discounts

Lastly, we explore how much of the *variation* in pricing dynamics across markets is accounted by the impact of competition and customer heterogeneity on the intertemporal slope. We proceed in two stages. In stage 1, we split our data into 3762 market-specific sub-samples, defined by the combination of a route r and a travel date t. In each of these market samples, we run the following first stage projection:

$$\ln(Price_{ijtd}) = \alpha_{rt} + S_{rt} Daysleft_{td} + \epsilon_{ijtd},$$

where α_{rt} and S_{rt} are the estimated market-specific coefficients. Collecting S_{rt} from all of these projections, this gives us a sample of market-specific intertemporal slopes.

In stage 2, we then regress the intertemporal slopes S_{rt} on a number of explanatory variables X_{rt} to assess their share in explaining the observed variation in S_{rt}. Specifically, we estimate

$$S_{rt} = \alpha + \beta \times X_{rt} + \epsilon_{rt}, \tag{7}$$

where X_{rt} varies across specifications. For each specification, we use the R-squared statistic to identify the fraction of the cross-market variations in S_{rt} explained by X_{rt}. Table 10 reports the results.

In Column 1, we set $X_{rt} = Comp_{rt}$ to assess the cross-market variation in the intertemporal slope that can be attributed to variation in the competitive environment alone (analogous to the patterns identified in Section 3). We find that competition accounts for about 5 percent of the cross-market variation.

In Column 2, we expand the set of explanatories to include all eight heterogeneity indicators introduced in Section 4. While individually each of these indicators only captures some variation in customer heterogeneity, it seems reasonable

[22] The difference between the two projections is likely due to the increased sampling frequency in the last week before departure. Because prices are generally increasing at a steeper slope during that period, the sample-average slope exceeds the time-average in our sample. Projecting the sample-average slope throughout the ten week horizon does therefore overestimate the contribution of systematic advance purchase discounts relative to unsystematic volatility. This bias vanishes, once we allow slopes to vary with the time left until departure, as we do in the projection based on our nonlinear estimation.

[23] Specifically, if the overall impact of competition on price dispersion were due to its impact on the intertemporal slope and its impact on residual dispersion in proportion to their relative contribution to within-flight dispersion, then a back of the envelope computation suggests that the impact of competition on residual dispersion is about one fifth of its impact on the intertemporal slope (i.e., choose x so that $-0.45 = -0.83 \cdot 0.54 + 0.17 \cdot x$, yielding $x = -0.11 \approx \frac{-0.52}{5}$).

that their combination spans a large part of the actual variation in customer heterogeneity. To allow for the patterns documented in Section 4, we interact each of the eight indicators with all five competition categories. Using this specification, X_{rt} accounts for 17 percent of the variation in S_{rt}, suggesting that a large share of the cross-market *variation* in slopes is explained by other factors.

To get an idea about what these other factors might be, we lastly set X_{rt} to include a full set of route and travel date fixed effects (again interacted with $Comp_{rt}$). The explanatory power of this fixed effect specification is 63 percent, providing an upper bound on the explanatory power of *any* route or travel data specific characteristics. The remaining 37 percent are only explainable using factors that jointly vary across travel dates *and* routes. While it is beyond the scope of this paper to identify the source of these residual variation in the intertemporal slope, we find it hard to reconcile them fully with any deterministic pricing scheme. We therefore conjecture that they represent, at least in part, price adjustments to unpredictable events such as demand shocks.

6. Discussion

From a theoretical perspective, it is far from obvious how competition affects the scope and profitability of raising ticket prices as the scheduled departure date approaches, and different theories deliver different predictions. If the main force behind rising prices is price discrimination against late booking customers with relatively inelastic demand, then a flattening of the slope in more competitive environments seems likely (e.g., Dudey, 1992; Martínez-de Albéniz and Talluri, 2011; Dana and Williams, 2019). On the other hand, if prices reflect scarcity of seats, then prices may increase even in perfectly competitive environments (Dana, 1998).

While the primary goal of this paper is to provide empirical guidance for developing new theories, we like to conclude this paper with a *prima facie* interpretation of our findings. In our view, our findings point to a two-faceted rational behind the observed pricing dynamics. On the one hand, the pivotal role of customer heterogeneity is indicative of intertemporal price discrimination: When customers are heterogeneous, a monopoly airline discriminates against late-booking customers with inelastic demand such as business travelers and other travelers that are desperate to fly. In more competitive environments, this ability to price-discriminate against late booking customers is restrained, flattening the intertemporal slope. However, this restraint is relevant mainly in high-heterogeneity markets where there is scope for price-discrimination in the first place, explaining the diminished impact of competition in low-heterogeneity markets.

On the other hand, our findings also indicate that even in markets with 5+ competitors prices tend to systematically increase over time at an average daily slope of 0.90 percent, suggesting another force at play. One possibility is that airlines face aggregate uncertainty regarding their demand. As first demonstrated by Dana (1998), uncertainty about demand may support advance purchase sales as airlines may lower the cost of holding excess inventories, even in perfectly competitive markets.[24] A corollary of such a stochastic demand interpretation is that even though prices may *on average* increase over time, the *realized* price path will depend on the realized demand and cannot be perfectly predicted by ex-ante market characteristics. Our results regarding the predictability of within-flight price dispersion and cross-market differences indicate that this is indeed the case.[25]

Acknowledgements

We are extremely grateful to Anton Vasilev for invaluable help and technical support throughout the process of collecting our dataset. We also would like to thank the editor, Frank Verboven, two anonymous referees, Ken Boyer, Meghan Busse, Niko Matouschek, Mar Reguant, Klaus M. Schmidt, Mauricio Varela, and various seminar and conference participants for useful comments and discussions. Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 and GRK 801 is gratefully acknowledged. Finally, we like to thank Xueqian Chen and Markus Deák for research assistance. The views in this paper are those of the authors, and not necessarily those of the Bank of England or its policy committees.

Appendix A. Data construction

A1. Routes

Our cross-section of routes is sampled from the existing connections between a set of 60 European cities with international airports. All routes are defined on the city-level. In case there exist multiple airports within one city, we include routes to all airport combinations (e.g., routes between London and Paris cover all offered combinations between {LCY, LGW, LHR, LTN} and {CDG, ORY}).

[24] See also Gale and Holmes, 1992, 1993 and Dana (2001) for similar arguments.

[25] See also Alderighi et al. (2015), Escobari (2012), and Puller et al. (2015) for direct evidence that realized seat scarcity matters for ticket prices.

Table A.1
List of routes.

Origin	Destination	Origin	Destination	Origin	Destination
Aberdeen	Manchester	London	Bordeaux	Paris	Dublin
Amsterdam	Barcelona	London	Frankfurt	Paris	Hamburg
Amsterdam	Zurich	London	Hannover	Paris	London
Athens	Budapest	London	Prague	Paris	Madrid
Athens	London	London	Sofia	Paris	Marseille
Barcelona	Lyon	London	Zurich	Paris	Prague
Belgrade	Vienna	Liverpool	Amsterdam	Paris	Stockholm
Berlin	Helsinki	Lyon	Madrid	Paris	Turin
Berlin	Vienna	Madrid	Barcelona	Paris	Valencia
Bilbao	Paris	Madrid	Copenhagen	Paris	Warsaw
Bologna	Madrid	Madrid	Lisbon	Palermo	Turin
Bordeaux	Madrid	Madrid	Milan	Prague	Helsinki
Bordeaux	Nantes	Madrid	Stockholm	Prague	Milan
Brussels	Leeds	Madrid	Valencia	Prague	Rome
Brussels	London	Madrid	Zurich	Rome	Nice
Bucharest	Milan	Malaga	Madrid	Rome	Vienna
Budapest	Munich	Milan	Copenhagen	Stockholm	Berlin
Copenhagen	Geneva	Milan	Duesseldorf	Stockholm	Duesseldorf
Copenhagen	Helsinki	Milan	Frankfurt	Stockholm	Oslo
Duesseldorf	Athens	Milan	Lyon	Stuttgart	Milan
Edinburgh	Manchester	Milan	Paris	Strasbourg	Paris
Frankfurt	Innsbruck	Moscow	Budapest	Toulouse	Brussels
Frankfurt	Istanbul	Munich	Athens	Toulouse	Paris
Frankfurt	Madrid	Munich	Madrid	Vienna	Amsterdam
Frankfurt	Moscow	Munich	Paris	Vienna	Barcelona
Frankfurt	Paris	Munich	Vienna	Vienna	Frankfurt
Frankfurt	Toulouse	Naples	Milan	Vienna	Lyon
Hamburg	Warsaw	Nice	Brussels	Vienna	Paris
Hannover	Amsterdam	Nuremberg	Amsterdam	Zurich	Frankfurt
Leipzig	Munich	Oporto	Paris	Zurich	Mallorca
Lisbon	Amsterdam	Paris	Copenhagen		

Notes: Prices are recorded for 41 distinct travel dates for each route. In 10 instances, we did not find any of our roundtrip combinations offered; 7 of them missing on the route Brussels–Leeds; the remaining 3 markets are missing on the routes Bordeaux–Madrid, Moscow–Budapest and Stockholm–Berlin.

The 60 cities were chosen to ensure regional variety as well as variety in the size and importance of the residing airports. To this end, the set includes the cities with the four largest airports in each of the EU5 countries (measured by 2009 total passenger traffic):

- France: Paris, Nice, Lyon, Marseille
- Germany: Frankfurt, Munich, Duesseldorf, Berlin
- Italy: Rom, Milan, Venice, Catania
- Spain: Madrid, Barcelona, Palma de Mallorca, Malaga
- UK: London, Manchester, Edinburgh, Birmingham

The remaining 40 cities are selected from both the EU5 and the rest of Europe (including Russia and Turkey): Aberdeen, Amsterdam, Athens, Belgrade, Bilbao, Bologna, Bordeaux, Brussels, Bucharest, Budapest, Copenhagen, Dublin, Geneva, Hamburg, Hannover, Helsinki, Innsbruck, Istanbul, Leeds, Leipzig, Lisbon, Liverpool, Moscow, Nantes, Naples, Nuernberg, Oporto, Oslo, Palermo, Prague, Sofia, Stockholm, Strasbourg, Stuttgart, Toulouse, Turin, Valencia, Vienna, Warsaw, and Zurich.

From the simplex of routes spanned by those 60 cities, we then sampled 100 random routes, disregarding all routes for which not at least one direct daily connection was offered at the beginning of our sampling period (October 31, 2010). To this routes, we added, if not yet contained, the ten routes connecting the cities with the largest airports in each of the EU5 countries (Paris, Frankfurt, Milan, Madrid, and London).

A limitation of our data source is that it does not contain prices set by Ryanair, a major competitor in the intra-European market. To prevent our data from being affected by an unobserved competitor, we therefore excluded all routes that were served by Ryanair within a 40 miles (65 km) radius of the corresponding city centers. From the above city list, this applies to all (existing) Ryanair route combinations between the following cities: Barcelona, Birmingham, Berlin, Bologna, Bordeaux, Brussels, Budapest, Catania, Dublin, Edinburgh, Leeds, Leipzig, Lisbon, Liverpool, London, Madrid, Malaga, Manchester, Marseille, Milan, Nantes, Nice, Nuernberg, Oporto, Palma de Mallorca, Palermo, Rome, Strasbourg, Turin, Valencia, Venice, and Warsaw. However, the majority of the possible combinations among those cities are *not* served by Ryanair, so that only a small number of drawn routes were affected.

These steps give the 92 connections between cities underlying our final sample. For each of them, we randomly assigned one of the two cities as departing city for the outbound flight and the other one as the departing city for the return flight. Table A.1 lists the resulting cross-section of routes.

A2. Affiliate groups

Our baseline measure for competition treats airlines that are affiliated through cross-holdings as single competitors. An airline is matched to an affiliate group if a member of that group owns more than 25% of the airline's equity. This is in line with the observed pricing practices in our sample, which show little variation *within* affiliate groups.[26] Out of the airlines observed in our sample, we have identified the following affiliate groups based on this criterion.

- Aegean, Aegean Airlines, Olympic
- Air France, KLM
- Air One, Alitalia, Meridiana fly, Wind Jet
- British Airways, Iberia, Vueling Airlines
- Air Dolomiti, Austrian Airlines, bmi, Brussels Airlines, Condor, Germanwings, Lufthansa, SunExpress, Swiss International Air Lines
- Blue 1, Cimber Sterling, Norwegian Air Shuttle, SAS, Spanair
- airberlin, Niki
- LAN Airlines, TAM Airlines, TAM Brazilian Airlines
- Singapore Airlines, Virgin Atlantic
- Air Seychelles, Etihad Airways
- Aeroflot-Russian Airlines, Malev Hungarian Airlines

Appendix B. Alternative competition measures

This appendix contains some robustness analysis with respect to our baseline approach to measure competition.

First, we consider two variations of our baseline approach to identify the number of competitors. Specifically, in our baseline measure, we treat codesharing airlines as competitors. Accordingly, if the same physical connection is marketed under different flight numbers that correspond to different (non-affiliated) airlines, this increases our measure of competition. The reasoning behind this choice is that in so-called "block space" codeshare agreements, each of the codesharing partner still controls a distinct, *ex ante* fixed amount of seats. In practice, by the pricing agreements between the carrier operating a service and the codesharing partner, the codesharer is usually granted considerable freedom to set prices independently.[27] In line with that, prices in our data differ indeed substantially across different codesharers.[28]

To evaluate whether implicit pricing agreements between codesharing airlines systematically affect our findings, we consider N_{ijt}^{csh} as a first alternative, which defines competition as the number of airlines that *operate* their own services on a particular market.[29] Similarly, we also consider N_{ijt}^{allie}, which counts the number of competing airline alliances, for which a similar concern might be raised.[30] Finally, we also use the Herfindahl index, which is a common alternative used in the literature.[31] For each of these measures, we consider both a log and a level variant.[32]

Table B.12 reports the results to the following estimation:

$$\ln(Price_{ijtd}) = (\alpha + \beta Daysleft_{td}) \times Comp_{ijt} + \lambda_i + \mu_j + \nu_t + \xi_d + \epsilon_{ijtd}, \tag{B.1}$$

where $Comp_{ijt}$ is a stand-in for the eight competition measures that we consider, and where λ_i, μ_j, ν_t, and ξ_d are treated as fixed-effects.

[26] The median within-*ijtd* standard-deviation of prices among members of the same affiliate group is 0 Euros, whereas the corresponding median dispersion for the same physical flight offered on the same date among codesharing partners *across* affiliate groups is 60.13 Euros (see also Footnote [28]).

[27] See, e.g., the report by the European Commission, "Competition impact of airline code-share agreements: Final report" (2007), available on the EC Website (last checked: October 2012).

[28] In our data set, 21.1 percent of the same physical roundtrip combinations *ij* for a given travel date *t* and a given date of the ticket offer *d* are on average offered by more than one codesharing operator (not counting codesharing within our affiliate group definitions). Among those observations, prices for a given flight sold at a given day differ across codesharing partners in 83.3 percent of the cases with a median within-*ijtd* standard deviation across codesharers of 60.13 Euros, suggesting that there is indeed significant leeway for independent pricing among codesharing partners.

[29] To be consistent with this approach, we also pool all physically identical roundtrips into a single observation, where at each date the pooled roundtrip is assigned the lowest price offered by any of the codesharing partners.

[30] Specifically, N_{ijt}^{allie} treats all airlines within "Star Alliance", "Sky Team" and "One World" as single competitors. Airlines not belonging to any of these alliances are counted as independent competitors according to our baseline competition measure.

[31] Since a Herfindahl index of one indexes the highest concentration of market power, we use $-H_{ijt}$ to make it qualitatively comparable with the other measures.

[32] To facilitate the comparison across the different measures which share different supports (the Herfindahl index is defined on the real line, whereas N_{ijt}, N_{ijt}^{csh}, and N_{ijt}^{allie} are natural numbers with different ranges), we adopt linear specifications for each of the competition measures in our estimation.

Table B.12
Alternative competition measures.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Baseline Slope							
$Daysleft_{td}$	−1.37 (0.09)	−1.32 (0.09)	−1.39 (0.09)	−1.29 (0.09)	−1.36 (0.09)	−1.29 (0.09)	−0.89 (0.09)	−1.38 (0.09)
	(b) Competition Effects							
N_{ijt}	0.08 (0.01)							
$\ln N_{ijt}$		0.19 (0.02)						
N^{csh}_{ijt}			0.11 (0.02)					
$\ln N^{csh}_{ijt}$				0.20 (0.03)				
N^{allie}_{ijt}					0.09 (0.01)			
$\ln N^{allie}_{ijt}$						0.18 (0.03)		
$-H_{ijt}$							0.50 (0.05)	
$-\ln H_{ijt}$								0.31 (0.03)
	(c) Implied Overall Slope							
Least competitive routes	−1.29	−1.32	−1.28	−1.29	−1.27	−1.29	−1.39	−1.38
Most competitive routes	−0.87	−0.98	−0.94	−1.02	−0.95	−1.02	−1.00	−0.91
Observations	1 417 635	1 417 635	1 417 635	1 417 635	1 417 635	1 417 635	1 417 635	1 417 635
R-squared (adj.)	0.58	0.58	0.58	0.58	0.58	0.58	0.58	0.58

Notes: The dependent variable is $\ln(Price_{ijtd})$. Unreported but included in the estimations are levels of the competition measures and fixed effects λ_i, μ_j, ν_t and ξ_d. Clustered standard errors (at the market level) are reported in parentheses. All reported coefficients are significant at the 0.1 percent level and are multiplied by 100. Overall slopes are linear combinations of the baseline slope and competition effects evaluated at the minimum and maximum value of the corresponding competition measure. E.g., in (4), N^{csh} ranges from 1 to 4, implying an overall slope for the most competitive route of $\ln 4 \times 0.1974 - 1.2902$, whereas in (7) the Herfindahl index H_{ijt} is ranging from 1 to 0.2223, implying an overall slope for the least competitive route of $-1 \times 0.5034 - 0.8904$.

The first thing to note, as can be seen in Panel (b), is that all competition measures yield the same qualitative conclusions as our baseline setup. In order to assess the quantitative implications across the variety of competition measures, we compute the linear combination of the baseline slope (Panel a) with the competition effects for the minimum and maximum realization of competition for each of the measures. The implied range of intertemporal slopes for each of our measures is reported in Panel (c). Here it can be seen that also quantitatively, the different measures yield very similar conclusions: While the different measures find a daily price increase ranging from 1.27 percent to 1.39 percent for the least competitive routes in our sample, the daily price increase estimated for the most competitive routes ranges from 0.87 percent to 1.02 percent.

References

Martínez-de Albéniz, V., Talluri, K., 2011. Dynamic price competition with fixed capacities. Manag. Sci. 57 (6), 1078–1093.
Alderighi, M., Nicolini, M., Piga, C.A., 2015. Combined effects of load factors and booking time on fares: insights from the yield management of a low-cost airline. Rev. Econ. Stat. 97 (4), 900–915.
Anton, J.J., Biglaiser, G., Vettas, N., 2014. Dynamic price competition with capacity constraints and a strategic buyer. Int. Econ. Rev. (Philadelphia) 55 (3), 943–958.
Borenstein, S., 1985. Price discrimination in free-entry markets. Rand J. Econ. 16 (3), 380–397.
Borenstein, S., Rose, N.L., 1994. Competition and price dispersion in the US airline industry. J. Polit. Econ. 102 (4), 653–683.
Brueckner, J.K., Dyer, N.J., Spiller, P.T., 1992. Fare determination in airline hub-and-spoke networks. Rand J. Econ. 23 (3), 309–333.
Dai, M., Liu, Q., Serfes, K., 2014. Is the effect of competition on price dispersion nonmonotonic? Evidence from the us airline industry. Rev. Econ. Stat. 96 (1), 161–170.
Dana, J.D.J., 1998. Advance-purchase discounts and price discrimination in competitive markets. J. Polit. Econ. 106 (2), 395–422.
Dana, J.D.J., 1999. Using yield management to shift demand when the peak time is unknown. Rand J. Econ. 30 (3), 456–474.
Dana, J.D.J., 2001. Monopoly price dispersion under demand uncertainty. Int. Econ. Rev. (Philadelphia) 42 (3), 649–670.
Dana, J. D. J., Williams, K. R., 2019. Intertemporal price discrimination in sequential quantity-price games. Unpublished Manuscript.
Dudey, M., 1992. Dynamic Edgeworth-Bertrand competition. Q. J. Econ. 107 (4), 1461–1477.
Escobari, D., 2012. Dynamic pricing, advance sales and aggregate demand learning in airlines. J. Ind. Econ. 60 (4), 697–724.
Escobari, D., Rupp, N., Meskey, J., 2019. An analysis of dynamic price discrimination in airlines. South Econ. J. 85 (3), 639–662.
Gaggero, A.A., Piga, C.A., 2011. Airline market power and intertemporal price dispersion. J. Ind. Econ. 59 (4), 552–577.
Gale, I.L., Holmes, T.J., 1992. The efficiency of advance-purchase discounts in the presence of aggregate demand uncertainty. Int. J. Ind Organ. 10 (3), 413–437.
Gale, I.L., Holmes, T.J., 1993. Advance-purchase discounts and monopoly allocation of capacity. Am. Econ. Rev. 83 (1), 135–146.
Gerardi, K.S., Shapiro, A.H., 2009. Does competition reduce price dispersion? New evidence from the airline industry. J. Polit. Econ. 117 (1), 1–37.

Goolsbee, A., Syverson, C., 2008. How do incumbents respond to the threat of entry? Evidence from the major airlines. Q. J. Econ. 123 (4), 1611–1633.
Hayes, K.J., Ross, L.B., 1998. Is airline price dispersion the result of careful planning or competitive forces? Rev. Ind. Organ. 13 (5), 523–541.
Holmes, T.J., 1989. The effects of third-degree price discrimination in oligopoly. Am. Econ. Rev. 79 (1), 244–250.
Lazarev, J., 2013. The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in us monopoly markets. Unpublished Manuscript.
McAfee, R.P., Te Velde, V., 2007. Dynamic pricing in the airline industry. In: Handbook on Economics and Information Systems.
Puller, S. L., Sengupta, A., Wiggins, S. N., 2015. Does scarcity drive intra-route price dispersion in airlines?Unpublished Manuscript.
Stavins, J., 2001. Price discrimination in the airline market: the effect of market concentration. Rev. Econ. Stat. 83 (1), 200–202.
Talluri, K.T., Van Ryzin, G.J., 2006. The Theory and Practice of Revenue Management, 68. Springer Science & Business Media.
Williams, K. R., 2017. Dynamic airline pricing and seat availability. Unpublished Manuscript.