Avoiding Competition-Enhancing Price Discrimination: Evidence from the U.S. Airline Industry

Matthew S. Lewis*

November 30, 2018

Abstract

Certain forms of price discrimination in oligopoly markets can lead to more aggressive competition and lower profits, yet few empirical studies examine either the extent to which this occurs or the ability of firms to avoid this lower-profit equilibrium. I explore the potential for competition-enhancing discrimination in the airline industry and formulate a new empirical approach that utilizes fare quotes to measure price discrimination while controlling for cost heterogeneity. The empirical findings reveal that carriers within the U.S. domestic market systematically avoid this form of discrimination, likely allowing carriers to maintain higher profits.

*John E. Walker Department of Economics, Clemson University, 228 Sirrine Hall, Box 341309, Clemson, SC 29634-1309, mslewis@clemson.edu

I Introduction

Price discrimination is often viewed as a valuable tool that firms with market power can use to extract additional surplus and increase profits. This intuition is based on the fact that price discrimination by a monopolist will always generate profits that are at least as large as under uniform pricing. In oligopoly markets, the same intuition does not necessarily apply. Price discrimination by oligopolists can result in higher equilibrium profits for all firms. However, when it is possible to target discounts to customers who have a higher willingness to pay for a rival firm's product, equilibrium profits can be lower for all firms under price discrimination than under uniform pricing. In a variety of settings (Thisse and Vives, 1988; Shaffer and Zhang, 1995, 2002; Corts, 1998; Chen et al., 2001), the ability to price discriminate can result in a prisoner's dilemma in which each firm has a dominant strategy to price discriminate despite the fact that profits would be higher for all if discrimination were not possible.

Given the potential for price discrimination to intensify oligopoly competition, it seems natural to wonder how frequently these situations arise in real-world markets. Surprisingly, despite an extensive theoretical literature, little empirical evidence exists on the prevalence or extent of pro-competitive price discrimination or the ability of firms to potentially prevent aggressive competition by avoiding the use of price discrimination in these situations.[1] This study aims to continue filling this gap in the literature by examining the role of pro-competitive price discrimination in an important and commonly-studied market: airline travel.

Economists often use airline pricing as a canonical example of price discrimination across consumer groups, and an extensive literature has documented and examined the use of various forms of price discrimination in this market (e.g., Stavins, 2001; Puller and Taylor, 2012; Puller et al., 2009; Dana, 1998; Borenstein and Rose, 1994; Gerardi and Shapiro,

[1]Examples of some of the few empirical papers that do consider price discrimination targeting rivals' customers include Asplund et al. (2008) who study subscription discounts offered by local newspapers in Sweden and Miller and Osborne (2014) who study location-based price discrimination in the Southeastern U.S. cement industry. Villas-Boas (2009) and Grennan (2013) estimate structural models of discriminatory bargaining in wholesale markets for coffee (Villas-Boas) and coronary stents (Grennan) and present counterfactual results suggesting that policies establishing uniform prices would soften competition. Nevo and Wolfram (2002) also present evidence that the use of coupons in the breakfast cereal market intensifies competition.

2009; Chandra and Lederman, 2018). Airlines are known to frequently charge more for tickets on the same flights to passengers that purchase, for example, one-way tickets rather than round-trip tickets, itineraries that don't include a Saturday night stay-over, or tickets that are eligible for refunds or exchanges. The goal in offering these different fares is to encourage passengers with a less elastic demand for travel (e.g., business travelers) to pay higher fares while allowing more elastic passengers (e.g., leisure travelers) to select lower fares. Such strategies, however, do not segment consumers based on the likelihood that they prefer one particular airline over another and, therefore, cannot be used to target discounts to customers who prefer a rival carrier.

In light of this, my study begins with a review of the theoretical conditions under which price discrimination can enhance competition and a discussion of what such strategies might look like in the airline market. The empirical analysis then focuses on a particular form price discrimination which meets three important criteria: 1) it has the potential to enhance competition; 2) it is clearly feasible for airlines to implement; and 3) it is possible to identify and measure empirically. Specifically, I examine the potential for airlines to price discriminate based on the airport from which a passenger's itinerary originates.

Due to frequent flyer and and other incentive programs, many passengers prefer to concentrate their trips on one airline. However, it is easier to maintain loyalty to an airline if that airline flies to more destinations from the customer's home airport. Previous work has established that customers traveling on a given route are more likely to choose the airline that has the largest presence at the originating airport (e.g., Borenstein, 1991). As a result, if an airline is flying on a route between airports A and B and has a significantly larger airport presence at A, they have an opportunity to price discriminate by charging passengers originating out of airport A (i.e., flying a round trip from A to B to A) more than those originating out of airport B (traveling from B to A to B).

Using data on airline price quotes for specific flights I implement a new empirical design that isolates fare differences resulting from this type of directional price discrimination. Previous studies of airline price discrimination compare the prices of tickets purchased at different times or for travel on different flights which may also reflect cost differences generated by capacity constraints. My approach eliminates potentially confounding cost differences by comparing different prices quoted at the *exact same time* for seats on the *exact same flight*. More specifically, I compare the price charged for a specific round-trip itinerary from airport A to airport B to the price that a round-trip passenger traveling from B to A would pay to fly on the exact same flights.[2]

The results reveal that carriers serving domestic routes within the United States do not directionally price discriminate, choosing instead to systematically avoid offering discounted fares to customers who prefer to fly on rival carriers. Over 99% of itineraries exhibit identical prices for passengers originating at different endpoints. This behavior is particularly striking given that the same airlines engage so regularly in other forms of (profit-enhancing) price discrimination on these same U.S. routes. Considering that a relatively small number of airlines interact frequently across a large number of routes, the complete lack of price discrimination targeting rivals' customers could plausibly result from a tacit coordination to restrict the strategy set and prevent aggressive price competition. This possibility is further supported by the fact that these same airlines do engage in this exact form of directional price discrimination over 95% of the time on international routes (from U.S. airports) where they tend to face a much more diverse set of rival carriers and where directional price discrimination is less likely to result in more intense competition.

The empirical facts described above are new to the literature and represent the first empirical measures of airline price discrimination identified using an approach that entirely eliminates unobserved cost variation. In addition, of the numerous empirical studies of price discrimination in the airline industry, this is, to my knowledge, the first study to investigate airlines' use of potentially competition-enhancing forms of price discrimination. A recent working paper by Luttmann (forthcoming) also considers the idea of directional price discrimination, but focuses on the possibility of discriminating based on differences in endpoint city demographics rather than targeting rival airlines' loyal customers. Moreover, Luttmann's analysis relies on quarterly route-level average fares derived from the U.S. De-

[2]While it would be impossible for the same two flights to actually serve as a round trip in both directions since the outbound flight must occur before the inbound flight, I detail in Section III how this comparison can still be made under the assumption that the prices of specific outbound and inbound legs are independent of their pairings.

partment of Transportation's Origin and Destination Survey (DB1B) which are more highly aggregated than my flight level price quote data. With these aggregate data it is not possible to identify whether different price levels are being charged or whether customers in some cities are simply more likely to purchase higher priced tickets (with fewer restrictions, for example) than consumers in other cities. Interestingly, Luttmann concludes from his findings that airlines do engage in directional price discrimination, whereas I am able to confirm with more detailed data that this does not appear to be the case.

II Oligopoly Price Discrimination and Airline Pricing

Theoretical studies including Borenstein (1985) and Holmes (1989) establish that price discrimination in non-monopoly markets is sustainable and can generate higher profits. Unlike in the monopoly case, however, conditions can arise under which discrimination leads to lower prices for consumers and lower profits for firms. As illustrated by Corts (1998), if competing firms agree on which types of consumers are more elastic, equilibrium prices under price discrimination will be higher for less-elastic consumers and lower for more-elastic consumers—mirroring the impact of price discrimination under monopoly. In contrast, when firms disagree on which groups are more elastic and when each firm tends to sell more to the group they view as less elastic, then price discrimination can result in "all-out competition" where prices are lower for all consumer groups than under uniform pricing. In this case, the firms may collectively prefer *not* to have the ability to price discriminate, as uniform pricing by all firms produces higher profits. In the language of Corts (1998), the former case represents a market exhibiting *best-response symmetry* while the latter case exhibits *best-response asymmetry*. Though Corts (1998) demonstrates this result for differentiated-product oligopolists engaging in third-degree price discrimination, similar outcomes are shown to arise, for example, when spatially-differentiated firms engage in individualized location-specific pricing (Thisse and Vives, 1988; Shaffer and Zhang, 2002) or use coupons targeted at consumers in certain locations (Shaffer and Zhang, 1995).

Airlines do not have the the ability to use individualized pricing, and they often lack the information necessary to directly classify an individual as being of a particular customer type. Instead, they tend to rely on a small set of observable factors that indicate whether a consumer is more likely to be of a particular type: the day and time of the flight, how long they stay at their destination, when the flight is booked, whether the fares include cancelation restrictions, etc. These factors are used to separate leisure travelers who exhibit more elastic demand from business travelers who tend to have more inelastic demand. As a result, discriminatory pricing based on these factors will exhibit best-response symmetry in that all airlines prefer to target leisure travelers with lower fares and business travelers with higher fares, and such strategies can be expected to result in higher airline profits rather than more aggressive price competition.

Price discrimination schemes exhibiting best-response asymmetry require some information about the consumer that is correlated with the consumer's relative valuations of products from different sellers. The firm may directly observe some characteristic of the consumer or it may infer the consumer's preferences based on who they have purchased from in the past (Fudenberg and Tirole, 2000; Chen, 1997). In the airline market, however, this information is quite limited, particularly when prices are quoted before customers reveal any identifying personal information. Inference regarding a consumer's preferences over carriers must be drawn from the characteristics of the flight they are looking for. Though most flight characteristics are better suited for separating business and leisure travelers, airlines could utilize the origin airport of the passengers' itinerary to take advantage of geographic heterogeneity in preferences for different airlines.

Research on the so-called "hub premium" has provided clear evidence that consumers prefer to fly on airlines that offer more flights originating out of their city's airport. Frequent flyer programs and other corporate loyalty programs encourage customers to concentrate their flying with one airline, and airlines that fly to more destinations allow customers to both accumulate loyalty credit on more of their trips and utilize loyalty rewards to fly to more destinations. As a result, when an airline flies a round-trip between one airport from which it serves many routes (like a hub airport) and another from which it serves only a few routes, customers on the route originating from the first airport are likely to have a stronger preference for this airline than those originating from the second airport. Direct empirical evidence is provided by Borenstein (1991), who shows that airlines with a substantially higher overall airport passenger share at airport A than at airport B tend to have a significantly higher market share on the round-trip route from A to B to A than they they do on the route B to A to B. In other words, passengers on the route originating from airport A prefer to fly on the airline with the highest airport share at A while passengers on the same route originating from B prefer to fly on the airline with the largest airport B share. In Appendix A I replicate the main analysis of Borenstein (1991) using more recent data and obtain remarkably similar results, confirming that airport presence still strongly influences directional route market shares.

Based on this geographic heterogeneity, airlines may have an incentive to charge higher prices to passengers originating out of airports where they are more highly valued. On routes where the airline has a significantly larger presence at one endpoint airport than the other, this could be accomplished through the use of directional price discrimination where passengers on the same route are charged different prices based on their origin airport. Moreover, if the carriers on the route differ in their relative presence at the two endpoint airports, then directional price discrimination could very likely exhibit best-response asymmetry causing airline 1 to have a higher price on A to B to A than on B to A to B while airline 2 has a lower price. Therefore, unlike other forms of price discrimination commonly studied in airline markets, directional price discrimination has the potential to enhance competition and result in lower equilibrium profits.

A simple theoretical example in the spirit of Corts (1998) makes this point clear. Consider a standard differentiated-product price-setting environment in which two firms each sell tickets in both directional non-stop markets between airports A and B: the round trip from A to B to A and the round trip from B to A to B.[3] Suppose that firm 1 has a larger airport presence at A than B and firm 2 has a larger presence at B than A, so that firm 1 is valued more highly by travelers originating from A than from B while firm 2 is valued more highly by those originating from B than from A. To reflect this, assume that firm 1's best response, $b^1(p^2)$, to any price by firm 2 is always higher in the A to B to A market than the

[3]The results presented by Corts (1998) are quite general and require only a few assumptions standard for pricing games: that profit functions are convex, that the game exhibits strategic complementarity, and that demand is more responsive to own prices than prices of rivals (assuring a unique equilibrium pricing vector).

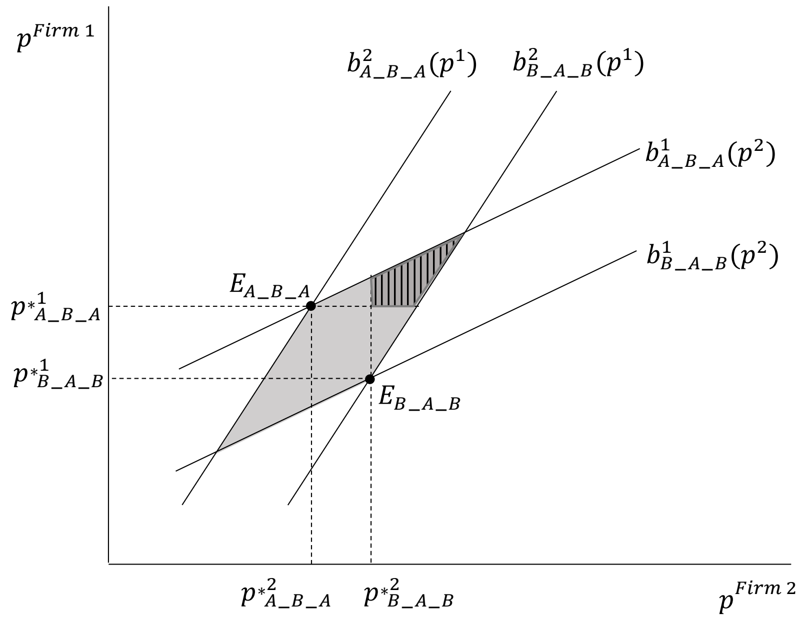

Figure 1: Directional Price Discrimination with Best-Response Asymmetry

B to A to B market, and the reverse is true for firm 2:

$$\begin{aligned} b^1_{A\_B\_A}(p^2) &> b^1_{B\_A\_B}(p^2) \quad \forall\, p^2 \\ b^2_{B\_A\_B}(p^1) &> b^2_{A\_B\_A}(p^1) \quad \forall\, p^1. \end{aligned} \tag{1}$$

Competition under these conditions is referred to as exhibiting best-response asymmetry, where the A to B to A market represents the *strong* market for firm 1 and B to A to B represents the *strong* market for firm 2.

When the firms engage in directional price discrimination the equilibrium prices in each directional market are determined by the intersections of these best response curves as depicted in Figure 1. Moreover, when directional price discrimination is not possible and firms set uniform prices across the directional markets, Corts (1998) shows that all uniform price equilibria lie within the shaded region bounded by each firm's single-market best response functions. Where the equilibrium is located within this region depends how on each firm's sales are split across markets. Firms that receive a larger share of sales from a particular directional market will have a profit-maximizing uniform price that is closer to their single-market best response for that directional market. As a result, if each firm obtains a majority of their sales from their strong market, uniform price equilibria will be located in the northeast corner of the shaded region in Figure 1 and will tend to compare favorably with prices obtained under directional discrimination. In fact, when uniform-price equilibria lie within the sub-region shaded with vertical lines, the use of directional price discrimination intensifies competition leading both firms to earn lower prices and profits in both markets than under uniform pricing. Such circumstances appear likely to arise in the U.S. domestic airline market. As the results of Borenstein (1991) and my own analysis in Appendix A confirm, differences in directional route market shares on U.S. domestic routes correspond strongly with differences in endpoint airport share, suggesting that airlines sell many more tickets in their stronger directional market than in their weaker directional market. Consequently, under conditions of best-response asymmetry a prisoners dilemma is likely to arise in which the use of directional price discrimination represents a dominant strategy for airlines yet leads to lower profits in equilibrium.

In Appendix B I provide a more concrete example of a specific setting in which best-response asymmetry and competition-enhancing directional discrimination will arise. Using a version of Hotelling's spatial model to illustrate competition in each directional market, I show that firms charge lower prices and earn lower profits when directional price discrimination is possible. The result is also shown to hold even when capacity constraints are binding, which may become relevant as airlines get close to selling all available seats on their scheduled flights.

It is important to note that directional price discrimination can also be useful under conditions of best-response symmetry. Consider a route in which all travelers from endpoint city A have a substantially lower willingness to pay to travel from A to B to A than do those in city B traveling from B to A to B. Travelers originating from A may still prefer a different carrier than those originating from B, but if these carrier preferences are outweighed by the differences in the general willingness to pay for travel, then all carriers on the route will view travelers originating from A as more elastic and charge lower fares on A to B to A than on B to A to B. In this case best-response symmetry holds, and directional discrimination will be more profitable for all firms than uniform pricing. Consequently, if directional discrimination is observed on a route, determining whether competing airlines charge higher fares to passengers originating from the same airport or different airports can serve as an indicator of whether the market exhibits best-response symmetry or best-response asymmetry and, hence, whether discrimination increases profits or intensifies competition. In contrast to the U.S. domestic market, my results for international routes (in and out of the U.S.) suggest that best-response symmetry may be more common due to large differences between the endpoint countries in demand for travel on the route.

III Empirical Strategy

The goal of my empirical analysis is to reveal the extent to which airlines engage in competition-enhancing price discrimination that targets discounts to customers who prefer to travel on rival carriers. I focus in particular on examining directional price discrimination for several important reasons. First, there is strong empirical evidence revealing substantial differences in passenger preferences between origin airports on domestic routes. Second, airlines have the ability to charge different prices to passengers departing from different origin airports. Third, comparing prices for round trip itineraries on the same route utilizing the same flights eliminates unobserved cost differences and results in a more accurate measure of the degree of price discrimination.

In this setting price discrimination on the non-stop route between airports A and B represents a difference in fares between passengers traveling on a round trip from A to B to A and passengers traveling from B to A to B that cannot be explained by differences in the airline's underlying costs. The opportunity costs of selling a seat on a specific flight can be influenced by many factors that are difficult to observe or control for. At any point in time, however, that opportunity cost should be the same regardless of what passenger purchases the ticket. In fact this same notion of opportunity cost is directly reflected in the *bid prices* used within airline revenue management systems (Airline Tarriff Publishing Company, 2017). Bid prices are determined for each flight segment within the route network to reflect the shadow value (in terms of expected future revenue) of selling an additional seat, and potential fare levels are only opened for booking if the fare exceeds the sum of the current bid prices for each leg in the itinerary. As a result, any observed differences in fares offered concurrently to different groups for a seat on the same flight can be interpreted as resulting from price discrimination.

Any flight between airports A and B will undoubtedly carry some passengers who originated their round trip at A and others that originated at B. On the other hand, a particular pair of flights from A to B and B to A can only represent a round trip itinerary from A to B to A if the A to B flight occurs before the B to A flight, and vise versa for the B to A to B itinerary. No passengers with different origin airports can ever have identical itineraries. This appears to undermine the empirical strategy of comparing round-trip fares offered on the exact same flights. Fortunately, with one additional assumption (that is strongly supported in the data) it is possible to recreate this comparison.

The necessary assumption (as stated in Assumption 1) is that airlines determine prices for the outbound and inbound legs of a round trip independently of which leg they are paired with.

Assumption 1 *At any particular time, there exists a unique set of latent prices for each potential outbound flight and for each potential inbound flight on a given route such that the observed price of any round trip itinerary on the route will be the sum of the latent prices of the associated outbound and inbound flights.*

To clarify the implications of the assumption, denote the latent price for outbound flight X as $\tilde{P}_X$, the latent price for inbound flight Y as $\tilde{P}_Y$, and the observed price for a round trip itinerary consisting of outbound flight X and inbound flight Y as P_{X_Y}. If the assumption holds, then for any two outbound flights X and X' and inbound flights Y and Y' it follows that:

$$P_{X\_Y} - P_{X'\_Y} = (\tilde{P}_X + \tilde{P}_Y) - (\tilde{P}_{X'} + \tilde{P}_Y) = (\tilde{P}_X + \tilde{P}_{Y'}) - (\tilde{P}_{X'} + \tilde{P}_{Y'}) = P_{X\_Y'} - P_{X'\_Y'}. \quad (2)$$

This assumption is not likely to hold for all possible itineraries. The price for a particular itinerary (on particular airline) generally depends on whether the departure dates and times of the outbound and inbound flights meet certain fare rules and restrictions. For example, some fares may only be available for itineraries involving flights that depart at certain times of the day or fly on certain days of the week. There can also be itinerary-specific fare rules such as minimum-stay or Saturday-night-stay requirements, and the (latent) price associated with a particular outbound flight may change depending on whether it is paired with an inbound flight that satisfies a particular fare rule. Since, the identifying assumption is more likely to hold when comparing itineraries that all satisfy the same rules and restrictions, my investigation of directional price discrimination focuses specifically on itineraries that are similar in these respects. Fortunately, by comparing the observed itinerary prices of different pairs of outbound and inbound flights, it will actually be possible to demonstrate that the condition in Equation 2 holds in the data, providing further support for the main identifying assumption.

Given this assumption, it becomes possible to compare round-trip fares quoted to a passenger originating at A and a passenger originating at B for travel on the exact same flights. More specifically, consider round trip itineraries involving flights on two out of four possible dates: Date 1, Date 2, Date 3 and Date 4, and denote flights based on origin, destination, and date such that an itinerary including a flight from A to B on Date 1 and from B to A on Date 2 would be represented by AB1_BA2. For a passenger originating at airport A, let the price of this itinerary be represented as $P^{A}_{AB1\_BA2}$. Consider three possible itineraries originating from A: AB1_BA2, AB3_BA4, and AB1_BA4. Figure 2 depicts each of these itineraries using dotted lines connecting their two component flights. The third itinerary shares one flight each with the first two. Therefore, if each flight (or leg) is priced independently of what itinerary it belongs to, then $P^{A}_{AB1\_BA2}+P^{A}_{AB3\_BA4}-P^{A}_{AB1\_BA4}$ should be equal to $P^{A}_{BA2\_AB3}$, representing the price it would cost a passenger originating at A to take the BA2 and AB3 flights. Of course, this pair of flights does not represent a viable round trip itinerary for a passenger originating at airport A because the BA flight occurs before the AB flight. However, if each flight is priced independent of its itinerary, then it does represent the incremental price charged for this pair of flights to a passenger originating out of airport A and it can be directly compared to the price that a passenger originating from B would pay to fly this itinerary: $P^{B}_{BA2\_AB3}$.

Following this logic, the difference between the price charged to a passenger originating at airport A and the passenger originating at airport B for a round trip ticket on the

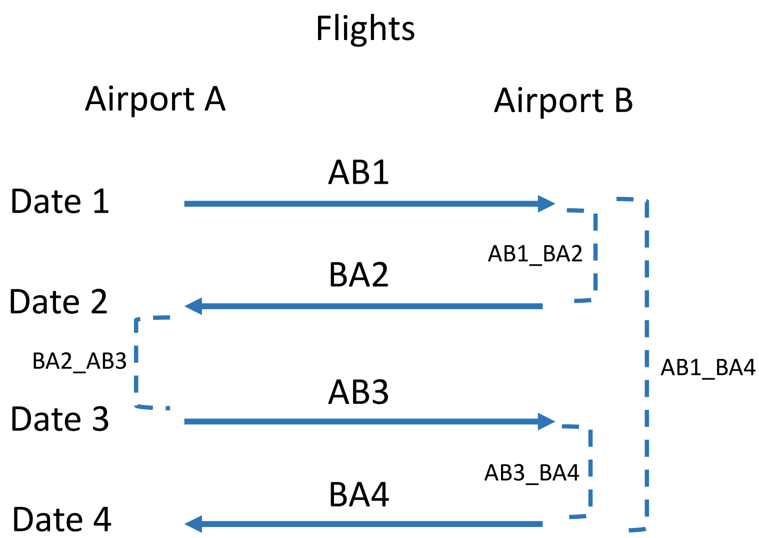

Figure 2: Flights and Itineraries Used to Identify Directional Price Discrimination

route between A and B can be measured as:

$$P^A - P^B = P^A_{AB1\_BA2} + P^A_{AB3\_BA4} - P^A_{AB1\_BA4} - P^B_{BA2\_AB3}. \tag{3}$$

By design, this represents a difference in the prices quote d to two different consumer types for travel on the exact same flights (specifically, flights BA2 and AB3), so it cannot be explained by differences in cost of service. Consequently, I adopt this difference as my measure of the degree of directional price discrimination on the route between A and B.

Notice that the directional price difference measured in Equation 3 is based on round trip itineraries involving the specific flights AB1, BA2, AB3, and BA4. If an airline flies more than one flight per day on the route, then more than one measure of the price difference can be constructed. As an example, if the airline offers 2 flights per day in each direction, then there are $2^4 = 16$ possible combinations of flights that could be used, and therefore 16 different observations of directional price differences for that airline on the AB route.

IV Data

The primary data used for analysis are price quotes for specific flight itineraries obtained from an airfare aggregator website. The aggregator itself collects itinerary and fare information in real time from a Global Distribution System (GDS) that disseminates the current fares provided by each airline to the Airline Tariff Publishing Company (ATPCO). My analysis will focus on non-stop round-trip coach-class fares[4] from the top 500 most heavily-traveled U.S. domestic routes and, for comparison, the top 300 non-stop international routes involving one U.S. airport.[5] For each quoted itinerary, the data include: the fare, ticketing carrier, operating carrier, origin and destination airports, and the flight times and dates of the outbound and inbound flights. Itineraries involving flights by more than one ticketing carrier are excluded.

Fares obtained by the airfare aggregator through the GDS come directly from the airlines and reflect the prices one would find offered for sale on each airline's website. Often, the aggregator will also present alternative quotes for the exact same flight (on the same airline) that it has collected from booking engines such as Expedia, Orbitz, Travelocity and a number of smaller travel sites. These fares are usually the same as that from the carrier website, but not always. Fares from these other sources may include additional discounts or markups added by the booking agent that do not reflect the strategic pricing decisions of the airline. As a result, my analysis will only utilize fares obtained by the aggregator directly from the airline.

Domestic Routes

Pricing information is observed for all major domestic airlines with the exception of Southwest Airlines, whose prices are consistently absent from airfare aggregator websites. In

[4]"Basic economy" fares from airlines such as American, United, and Delta are not included in the sample. Only fares classified by these carriers as "economy" are considered. Fares from airlines such as Spirit and Frontier are also included even though the amenities associated with these fares typically differ from the economy fares of legacy carriers. However, my analysis focuses on within-carrier differences in fares, so cross-carrier differences in amenities are inconsequential.

[5]The top 500 non-stop routes were selected based on total enplanements recorded in the U.S. Department of Transportation's (DOT) Origin and Destination Survey (DB1B) during the year 2016. Similarly, the top 300 international routes were selected based on total passengers reported in then U.S. DOT's T-100f during 2016.

all cases, the ticketing carrier is considered the carrier of interest. On nearly all domestic flights, the operating carrier is either the same as the ticketing carrier or is a regional airline (e.g., Piedmont Airlines, Republic Airlines, Mesa Airlines) who has been contracted to operate the flight on behalf of the ticketing airline. On the occasion where a major carrier is serving as an operating carrier for different ticketing airline (usually American selling tickets for a flight operated by Alaska Airlines, or vise versa), these price quote observations are dropped.

Computing the measure of directional price discrimination described in the previous section requires price quotes observed at a particular point in time for four different round-trip itineraries on each potential route, as depicted in Figure 2. Four different rounds of data collection are performed for domestic routes. Each round is defined by the date on which the price quotes are observed and the 4 travel dates that are used to construct quoted itineraries. The first round consists of quotes collected in early-December of 2017 for itineraries involving travel on four consecutive days: Monday Feb. 5th through Thursday Feb. 8th of 2018. The second round consists of quotes collected in mid-January of 2018 for travel on four consecutive Tuesdays: March 6th, 13th, 20th, and 27th of 2018. The third round contains quotes collected in early February 2018 for travel on four consecutive days: Tuesday March 6th through Friday March 9th. The fourth round includes quotes obtained between March 24th to 28th for travel on consecutive Wednesdays: April 25th and May 2nd, 9th, and 16th. Though quotes within a round may be collected across different days, all price quotes for itineraries on particular route are collected within minutes of each other, insuring that all within-route price comparisons are appropriately based on fares that were offered concurrently. The travel dates are chosen so that itineraries constructed using each round of data will all be relatively comparable. Itineraries in Rounds 2 and 4 will all include a weekend stayover while those in Rounds 1 and 3 will all include at least one night spent at the destination but do not include a weekend stayover.[6]

A route is defined by its unique paring of endpoint airports A and B, and is not di-

[6] Given that airlines have been known to charge different prices depending on whether travelers stay in the destination overnight or over a weekend, this approach attempts to avoid price differences between itineraries within each price comparison group that might arise for reasons other than the direction of the round trip. In other words, travel dates are selected to maximize the likelihood that Assumption 1 is satisfied.

rectional. To simplify the exposition, the term *route* and the subscript r are used to denote a particular pair of endpoint airports from a particular data collection round. Correspondingly, P^A_{cri} represents the price of the round-trip itinerary i originating out of airport A on route r provided by carrier c, and P^B_{cri} represents the price of the round-trip itinerary i originating out of airport B on route r provided by carrier c.[7] The prices P^A_{cri} and P^B_{cri} are constructed as in Equation 3 and used to generate the directional price difference for itinerary i served by carrier c on route r:

$$\textbf{DirectionalPriceDiff}_{cri} = P^A_{cri} - P^B_{cri}.$$

It will also helpful be to measure the price difference in percentage terms, which is similarly constructed as follows:

$$\textbf{DirectionalPriceDiff}\%_{cri} = \frac{P^A_{cri} - P^B_{cri}}{P^B_{cri}}.$$

In addition to obtaining price quote data for the sampled routes, I have also collected 2016 aggregate ticket sales and passenger enplanements for each carrier on each route from the U.S. DOT's Origin and Destination Survey (DB1B). Passenger enplanement counts are used to construct airport market shares while non-stop round-trip ticket sales are used to construct route-level market shares. The variables are defined for carrier c, route r, and airport a as follows:

$\textbf{EndpointAirportMarketShare}_{cra}$**:** The sum of all passenger enplanements on flights by carrier c departing from airport a not including those on route r divided by the sum of all passenger enplanements (by any carrier) departing from airport a not including those on route r.

$\textbf{RivalsEndpointAirportMarketShare}_{cra}$**:** The weighted average of $\textbf{EndpointAirportMarketShare}_{kra}$ across all carriers $k \in K$ that compete with carrier c on route r, weighted according to each competitor's relative route share: $\textbf{RouteShare}_{kra}$.

[7] In this context, itinerary i actually refers to a combination of four specific round-trip itineraries (with specific flight times) as in Figure 2 that are used to generate one observed P^A and P^B pair.

$\mathbf{RouteShare}_{cra}$**:** The sum of all non-stop round-trip tickets sold by carrier c on route r originating from airport a divided by the sum of all non-stop round-trip tickets sold (by any carrier) on route r originating from airport a.

As described in the previous section, my empirical strategy is based on studying directional differences in round trip prices charged to passengers on the same exact flights but originating out of different endpoint airports. Correspondingly, for each of the variables above, the values at the two endpoints airports of the route r are differenced to generate the following:

$\mathbf{EndpointShareDiff}_{cr}$**:** The difference in the values of $\mathbf{EndpointAirportMarketShare}_{cra}$ for endpoint airports A and B of route r.

$\mathbf{RivalsEndpointShareDiff}_{cr}$**:** The difference in the values of $\mathbf{RivalsEndpointAirportMarketShare}_{cra}$ for endpoint airports A and B of route r.

$\mathbf{RouteShareDiff}_{cr}$**:** The difference in the values of $\mathbf{RouteShare}_{cra}$ for endpoint airports A and B of route r.

International Routes

Price quotes for international routes were obtained for four different sets of travel dates. The first set consists of quotes collected during the second week of December for travel on four consecutive Wednesdays: Feb. 21st and 28th and Mar. 7th and 14th of 2018. The second set was collected during the third week of January for travel on four consecutive Tuesdays: Mar. 6th, 13th, 20th, and 27th of 2018. The third set was collected during the first week of April for travel on consecutive Thursdays: May 17th, 24th, 31st, and June 7th. These first three rounds are similar in that they all consider travel booked roughly 2 months in advance on flights that always include a weekend stayover. In contrast, the fourth and final round of data collection was designed to examine flights with shorter stays and very little advance purchase. These quotes were collected between March 28th and 31st for travel on four consecutive days: Monday April 9th through Thursday April 12th. Across all rounds, fare information is unavailable for a handful of foreign carriers such as

China Eastern and Interjet, but these airlines collectively represent less than 1% of total observations on the sampled routes.

As in the domestic data, the ticketing carrier is considered the carrier of interest. However, code sharing arrangements are much more common in the international market where major carriers frequently sell tickets for flights operated by another major airline, usually when both airlines are members of the same airline alliance (e.g., Star Alliance, Oneworld, Sky Team). As a result, several ticketing carriers are often observed selling tickets on the same flight, sometimes at similar prices to each other and sometimes at different prices. In order to avoid the complexities of code sharing as much as possible, I choose to focus on the "primary" carrier for each airline alliance on each route, where primary is defined as the airline within each alliance that operates the largest number of flights on the route. All quoted fares where the primary alliance carrier serves as the ticketing carrier are then included in the sample regardless of whether they also serve as operating carrier for the flights. Airlines that are not members of an international alliance serve as both ticketing and operating carriers on their flights, so all quotes by these airlines are also included.

Directional price difference measures are constructed exactly as they are for domestic routes. To make interpretation of the results more straightforward the difference measure is calculated as the price of a round trip originating out of the U.S. endpoint airport minus the price of the equivalent round trip originating out of the foreign endpoint airport. Observed fares are inclusive of all taxes and fees which may vary substantially across airports and countries, but directional price differences remain unaffected because taxes and fees are always applied separately to each flight segment and are not determined by the origin of the round trip itinerary.

Unfortunately, the supplemental data on airport market shares and demographics gathered for domestic airports are not available for foreign airports or cities. As a result, I focus mainly on measures of U.S endpoint airport share for international routes:

$\mathbf{U.S.EndpointAirportMarketShare}_{cr}$**:** The sum of all passenger enplanements on flights by carrier c departing from the U.S. endpoint airport not including those on route r divided by the sum of all passenger enplanements (by any carrier) departing from the

U.S. airport not including those on route r.

$\mathbf{U.S.EndpointAirportAllianceMarketShare}_{cr}$: The sum of all passenger enplanements on flights by carrier c or alliance partners of carrier c departing from the U.S. endpoint airport not including those on route r divided by the sum of all passenger enplanements (by any carrier) departing from the U.S. airport not including those on route r.

Evidence Supporting the Identifying Assumption

The proposed approach for measuring directional price discrimination relies crucially on the assumption that airlines set the underlying prices of the outbound and inbound legs of a round trip independently of which leg they are paired with (Assumption 1). Fortunately, it is possible to provide empirical support for this assumption by testing one of its more immediate implications. As described in the previous section, the assumption implies that the price difference between two round-trip itineraries involving the same inbound flight but different outbound flights should be the same regardless of which inbound flight they are paired with.

An example of this price comparison is provided in Table 1. Itineraries 1 & 2 have different outbound flights but share the same inbound flight and differ in price by $30. Itineraries 3 & 4 have the same outbound flights as 1 & 2 but are paired with a different inbound flight. Despite being more expensive than 1 & 2, itineraries 3 & 4 still differ by the same $30, suggesting that the relative prices the two outbound flights do not change when paired with a different inbound flight.

To test this assumption more broadly in the data, I identify over 15 million pairs of domestic itineraries that share a common outbound or inbound flight (like the pairs 1 & 2 or 3 & 4 in Table 1). Then I determine whether the difference in price between the two itineraries ever changes when the common outbound or inbound flight is exchanged with another possible flight. If the price difference always remains the same whenever this common flight is replaced with another common flight, I consider these observations to be consistent with the identifying assumption. If the price difference for the itinerary pair does

Table 1: Example of Prices that Support the Identifying Assumption

American Airlines: Indianapolis, IN (IND) to Washington, DC (DCA)
Departing May 2, 2018, Returning: May 9, 2018

Itinerary	Outbound Departure Time	Inbound Departure Time	Airfare	Airfare Difference
1	9:20 AM	8:55 AM	$185	$30
2	5:15 PM	8:55 AM	$215	
3	9:20 AM	2:59 PM	$265	$30
4	5:15 PM	2:59 PM	$295	

change when the common flight is exchanged for another, then I designate the itinerary pair as exhibiting a violation of the identifying assumption. Across all itinerary pairs, 99 percent appear to be consistent with the identifying assumption. This remains true even when exchanging the common outbound or inbound flight with a flight from a different day. I interpret this evidence as a validation of the proposed approach for identifying directional price discrimination.

An identical test is performed using the 2.9 million pairs of international itineraries that share a common inbound or outbound flight. Across these pairs, 98 percent appear to be consistent with the identifying assumption.

With only 1 percent of domestic itinerary pairs (and 2 percent of international pairs) exhibiting prices that violate Assumption 1, the impact of including such flights in my analysis of price discrimination is likely to be relatively minor. However, these observations will generate nonzero values of my directional price difference measure even when no differences in prices actually exist. To be conservative, I have chosen to exclude from my remaining analysis all observations for an airline on a particular route if *any* of its prices on that route violate the identifying assumption in the tests above.

After eliminating potentially problematic observations, the domestic sample includes prices for 2271 unique airline-route pairs from 991 different routes representing 393 unique origin-destination pairs.[8] The international sample includes prices for 1192 airline-route

[8]Recall that *routes* have been defined as an origin-destination pair within a given data collection round, so many origin-destination pairs (though not all) will appear in several data collection rounds.

pairs from 917 different routes representing 287 unique origin-destination pairs.[9]

V Empirical Analysis and Results

Domestic Route Fares

Since the identification of the degree of directional price discrimination occurs entirely within the design of the data collection process, many of the most important findings can be revealed by simply summarizing the data. In particular, the most striking finding is a nearly complete lack of directional price discrimination on U.S. domestic routes. Less than 1% of observed directional price differences are nonzero! This pattern is highly consistent across data collection rounds which contain itineraries that vary in length of stay and the degree of advance purchase. Although my analysis focuses on economy fares which represent the vast majority of tickets sold, the observed first class and premium fares also exhibit zero directional price difference 99.9% of the time.

All major competitors in the U.S. domestic market (including American, Delta, United, JetBlue, and Alaska) seem to adopt the same strategy of avoiding directional discrimination.[10] The only airlines in the data that appear to exhibit directional price differences with any meaningful frequency are Frontier and Spirit Airlines who have nonzero price differences 11% and 15% of the time, respectively. However, upon closer inspection, these observed price differences seem more likely to have been misreported by the airfare aggregator website than to represent actual systematic directional price discrimination. Of the 115 routes served by Spirit Airlines in the data, directional price differences were observed for some itineraries on 44 routes, but only 9 routes exhibited directional differences in more than one of the 4 collection rounds and in 6 out of 9 cases the directional differences observed in different rounds were of opposite sign. Similarly, of the 57 routes served by Frontier, directional differences are observed on 15 routes but never appear in more than one of the 4 collection rounds. True price discrimination reflecting underlying directional

[9] Using the alternative approach of dropping only observations that exhibit clear violations to the identifying assumption but keeping other observations for those carriers on those routes does result in a significantly larger sample but has very little impact on the empirical findings.

[10] Recall that prices for Southwest Airlines are not observed in the data.

Table 2: Summary Statistics - Domestic Routes

	# of Obs.	Mean	Std. Dev.	Min	25th %tile	75th %tile	Max
Price	2,139,978	245	100	51	181	289	1461
DirectionalPriceDiff	2,139,978	0.2	3.5	-122	0	0	489
DirectionalPriceDiff%	2,139,978	0.001	0.015	-1.23	0	0	1.89
EndpointAirportMarketShare	2271	0.228	0.233	0	0.047	0.390	0.916
EndpointShareDiff	2271	0.022	0.347	-0.775	-0.144	0.266	0.810

differences in willingness to pay would likely have produced directional price differences that were more regularly observed and more consistent in sign across travel dates.

Table 2 presents summary statistics for the main variables described in Section IV.[11] As the results of Appendix A and Figure 5 suggest, the absence of directional price discrimination does not result from a lack of variation in airport presence across endpoints. Airlines often have much higher market share at one endpoint of a route than the other. The mean of the absolute value of $\text{EndpointShareDiff}_{cr}$ across all carrier-route pairs is 26 percentage points. Within each route there is also considerable variation in this measure across competing airlines. The mean absolute deviation from the route-level mean of $\text{EndpointShareDiff}_{cr}$ is 22 percentage points. On a route with two competing airlines, for example, this implies that the $\text{EndpointShareDiff}_{cr}$ of the competitors differ by 44 percentage points on average. Moreover, it is often the case that one competitor will have a higher market share at one endpoint while another competitor will have a higher market share at the other endpoint. Of routes served by more than one carrier, over 65% have at least one carrier with a negative value of $\text{EndpointShareDiff}_{cr}$ and another carrier with a positive value. In other words, to the extent that willingness-to-pay is influenced by an airline's market share at the originating airport, competing airlines on many routes are likely to disagree on whether A-B-A travelers are more elastic than B-A-B travelers.

While the systematic lack of directional pricing may reflect a concerted effort to avoid competition-enhancing price discrimination, other explanations also need to be con-

[11]Observation counts for market shares reflect the number of unique combination of airline, route, data collection round. Counts for MSA demographics reflect the number of unique combination of route and data collection round. Due to the elimination of potentially problematic observations described in Section IV, some routes (or airline-route pairs) do not appear in the sample for every data collection round.

sidered. It is possible that directional price differences are absent because airlines rely on heuristics to simplify a very complex optimal pricing problem. It has become increasingly common for round-trip tickets to be sold at the combined price of the two component one-way flights. This simpler pricing approach eliminates many potential forms of discrimination including directional price discrimination as well as round-trip discounts, minimum-stay requirements, and Saturday-night stayover discounts. Southwest pioneered this approach in the U.S. and carriers like American, United, and Delta appear to be increasingly adopting this strategy, possibly to better compete with carriers like Southwest. However, the major carriers still use more sophisticated fare rules and restrictions on many routes, and data reveal that the prices on many routes still exhibit discounts for round trip travel and for certain itineraries like those with Saturday-night stayovers. Based on price quotes from the top 1000 routes, 37% of round-trip prices are at least 5% cheaper than the sum of the one-way prices of their component flight legs.[12] Similarly, using a version of the identification strategy depicted in in Figure 2, I find that more than 19% of itineraries are priced at least 5% cheaper when they include a Saturday-night stayover than when they do not.[13] When such discounts are offered, they are usually quite substantial. The average observed discount for round-trip itineraries is around 29%, while the average discount for Saturday-night stayover is over 62%.[14]

Given the continued use of these other forms of itinerary-based price discrimination it seems implausible to conclude that the lack of directional price discrimination has resulted from an industry-wide shift to one-way route pricing. Moreover, the absence of directional discrimination does not appear to be a new phenomenon. Though data necessary to document this were unavailable at the time, Borenstein (1991) claimed several decades ago that the “fares on a route are virtually always offered without reference to the traveler’s point of origin.” Carriers appear to have long avoided directional price discrimination on

[12]Similar to the primary dataset, this comparison is based on price quotes contemporaneously collected from the airfare aggregator website for round-trips itineraries (departing Friday Oct 26th and returning Monday Oct. 29th) as well as the component one-way itineraries.

[13]This comparison is based on price quotes for itineraries involving travel on two of the following four dates: Wednesday Dec. 6th, Thursday Dec. 7th, Tuesday Dec. 12th, or Wednesday Dec. 13th. The magnitude of the discount for Saturday-night stayover can then be calculated as: *Discount* $= (P_{AB1\_BA2} + P_{AB3\_BA4} - P_{AB1\_BA4} - P_{BA2\_AB3})/2$.

[14]Reported averages are calculating using only itineraries that exhibit discounts of at least 5%.

domestic routes despite their consistent use of many other forms of itinerary-based price discrimination (that are unlikely to be competition-enhancing).

The ability to accurately document the lack of directional discrimination relies crucially on the customized fare quote data and empirical strategy utilized here. Attempting to measure directional price differences using average fares from the Department of Transportation's commonly used DB1B 10% ticket sample database introduces the potential for substantial bias resulting from an inability to control for unobserved factors. The DB1B only reveals the route and the quarter in which the flight was flown, so it is not possible to construct directional differences in prices for travel on the same exact flights. If travelers from one endpoint more frequently book tickets at the last minute or purchase tickets on higher-priced flights (peak travel times or days), directional differences in quarterly-average fares may primarily reflect differences in cost or scarcity-pricing rather than price discrimination.[15] In fact, misleading conclusions could easily be drawn from such an analysis. For comparison purposes, Appendix C examines directional price differences constructed using average airfares from the DB1B data and presents several regressions that relate these directional price differences to endpoint airport shares and endpoint city demographics. Substantial directional differences in average fares do arise in the DB1B data, and the resulting coefficient estimates on $\mathrm{EndpointShareDiff}_{cr}$ and on the demographic variables are economically sizable and strongly statistically significant despite the fact that true directional price differences are virtually non-existent in my disaggregated price quote data. The comparison reveals why Luttmann (forthcoming) concludes (using DB1B data) that airlines directionally discriminate based on demographics while I conclude that discrimination is absent in the U.S. domestic market.

International Route Fares

The empirical evidence above indicates an overwhelming lack of directional price differences on domestic airline routes. Before discussing why this might be the case, I present in this section contrasting evidence from international routes where directional pricing behav-

[15]Borenstein (1991) specifically discusses this issue in his original analysis of directional differences in consumer demand.

ior is entirely different.

Just as the absence of directional price differences was immediately apparent in the domestic route data, basic summary statistics from international routes reveal a pervasive use of directional discrimination. Over 95% of observed directional price differences are nonzero, and in most cases the magnitudes of these differences are substantial. Figure 3 displays a histogram of the percentage directional difference in fares (in absolute value). Over 67% of observed directional differences are larger than 10% of the fare on the route.

Closer examination confirms that these price differences can not have been generated by directional differences in taxes or fees. There is substantial variation in directional price differences across airlines and collection periods within a route as well as across routes connecting a particular foreign airport with different U.S. airports. The standard deviations of the observed directional price differences at each foreign endpoint airport have a median value of around $75. Unobserved country- or destination-specific differences in taxes or fees would likely have resulted in more uniform directional price differences across routes from different U.S. origins.

There are only a handful of routes on which airlines offer fares with no directional price difference. These mainly consist of flights (by airlines including American Airlines, Delta, JetBlue, and Bahamasair) in and out of Nassau, Havana, or Bermuda, as well as several American Airlines routes in the Caribbean to destinations including Grand Cayman and Curacao. Aside from Copa Airlines flights in and out of Panama City no other routes or airlines exhibit zero directional price differences with any meaningful frequency. Clearly, directional price discrimination is the norm on most international routes. Moreover, there is no meaningful difference across data collection rounds in the frequency or average magnitude of directional price discrimination, suggesting that these price differences arise regardless of length of stay and advance purchase.

To better understand what motivates carriers on international routes to directionally price discriminate, I next estimate several models that relate directional price differences to carrier-route characteristics. Following the logic of Borenstein (1991), passengers are likely to have a higher willingness to pay for airlines that have a large market share at their

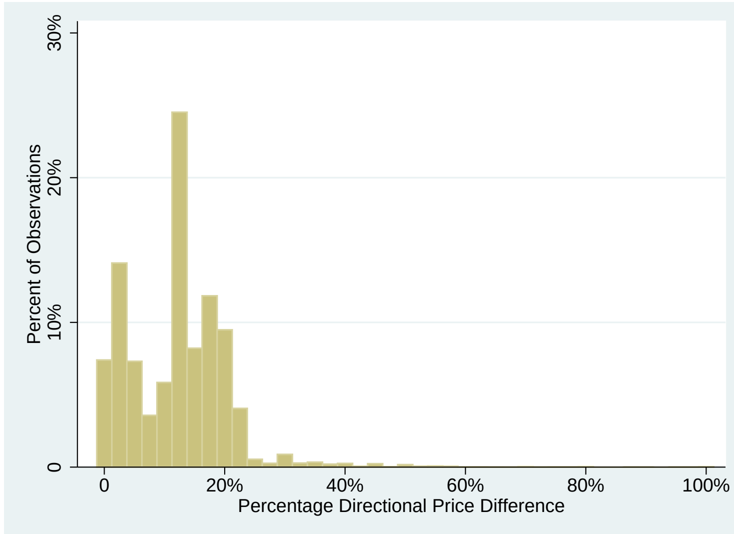

Figure 3: Histogram of the Absolute Percentage Directional Difference in Fares on International Routes

Notes: Values represent the absolute value of the directional price difference measured as a percent of the fare. Directional differences greater than 100% have been excluded (accounting for less than 0.1% of the sample).

origin airport, and therefore, one might expect directional price differences to be positively correlated with the difference in the carrier's airport shares at the two endpoint airports. On international routes, however, the potential importance of airport market share is not so straightforward. Frequent flyer programs give consumers an incentive to remain loyal to airlines that have a high market share in their local airport, but on international flights passengers can typically use and/or accrue frequent flier miles in an airline's program while traveling with the airline's alliance partners as well. Therefore, while Atlanta residents (who tend to be Delta frequent fliers because Delta dominates the Atlanta airport) may have a high value of taking a Delta flight when flying to London, London residents may also have a high value of taking Delta when flying to Atlanta if they are Virgin Atlantic (Delta's SkyTeam partner) frequent flier members. In this case, the total origin airport share of the airline alliance may be more important than the airport share of the actual ticketing carrier.

Unfortunately, for international routes I only observe airport market shares at each route's U.S. endpoint airport, so differences relative to the airport shares at the route's

Table 3: Summary Statistics - International Routes

	# of Obs.	Mean	Std. Dev.	Min	25th %tile	75th %tile	Max
Price	276,736	444	348	166	226	516	8567
DirectionalPriceDiff	276,736	31	159	-7532	-2	55	3490
DirectionalPriceDiff%	276,736	0.078	0.145	-6.23	-0.01	0.16	2.83
U.S.EndpointAirportAllianceShare	1192	0.361	0.283	0	0.123	0.680	0.910
RivalEndpointAirportAllianceShare	1192	0.173	0.201	0	0	0.246	0.853

foreign endpoint airport cannot be constructed. Nevertheless, routes on which an airline alliance has a large share at the U.S. airport may still exhibit higher directional price differences.[16] Table 4 reports the results of a simple regression of directional fare differences on the U.S. endpoint airport alliance market share and the average U.S. airport market share of other alliances competing on the route.[17] In Columns 2 & 3 foreign-endpoint-airport fixed effects are added to control for differences across countries that may impact the directional difference in airfares, and in Column 3 the sample is restricted to only include route-carrier combinations from routes that have more than one competing alliance (or non-alliance carrier). As described in Section IV, I avoid the complications of intra-alliance pricing decisions amongst alliance-partner carriers by limiting the sample to include only one observation per alliance per route.[18]

The coefficients on U.S. endpoint alliance airport share are positive and both statistically and economically significant in all specifications, suggesting that an increase in an airline alliance's presence at a U.S. endpoint airport may cause round trip prices for passengers originating from the U.S. airport to increase by more than prices on the equivalent round trip originating out of the foreign airport. More specifically, the results imply that a two-standard-deviation increase in the alliance airport share (i.e., an increase in share of 56 percentage points) would be associated with a 8.8 percentage point increase in the

[16]Recall that directional price differences on international routes are defined as the price of a round trip originating out of the U.S. endpoint airport minus the price of the equivalent round trip originating out of the foreign endpoint airport.

[17]The unobserved foreign endpoint airport shares of the carrier and it's rivals should be viewed as omitted explanatory variables here.

[18]The alliance operating the most flights on the route is selected. All non-alliance airlines are also included.

Table 4: Factors Associated with Directional Differences in International Fares

	All route-carrier pairs				Non-monopoly routes	
	(1)		(2)		(3)	
U.S.EndpointAirportAllianceShare	0.099	(0.058)	0.157	(0.052)	0.151	(0.071)
RivalU.S.EndpointAirportAllianceShare	0.086	(0.067)	0.119	(0.064)	0.138	(0.075)
Foreign Endpoint Airport Fixed Effects (Selected Cities Reported):						
Barcelona, Spain (BCN)			0.096	(0.034)	0.089	(0.038)
Beijing, China (PEK)			0.070	(0.139)	−0.154	(0.052)
Bombay, India (BOM)			−0.304	(0.035)	−0.307	(0.045)
Brussels, Belgium (BRU)			0.173	(0.020)	0.179	(0.028)
Cancun, Mexico (CUN)			−0.104	(0.039)	−0.095	(0.049)
Copenhagen, Denmark (CPH)			0.277	(0.001)		
Dubai, U.A.E. (DXB)			−0.290	(0.073)	−0.358	(0.055)
Frankfurt, Germany (FRA)			−0.038	(0.136)	−0.112	(0.310)
Hong Kong, Hong Kong (HKG)			−0.179	(0.106)	−0.065	(0.057)
Istambul, Turkey (IST)			0.349	(0.008)		
London, England (LHR)			0.045	(0.065)	0.006	(0.077)
Lima, Puru (LIM)			−0.307	(0.031)	−0.313	(0.031)
Mexico City, Mexico (MEX)			−0.064	(0.039)	−0.058	(0.044)
Osaka, Japan (KIX)			0.332	(0.028)	0.330	(0.036)
Panama City, Panama (PTY)			−0.101	(0.029)	−0.127	(0.053)
Paris, France (CDG)			0.027	(0.048)	0.041	(0.054)
Reykjavic, Iceland (KEF)			−0.173	(0.020)	−0.172	(0.020)
Rome, Italy (FCO)			0.252	(0.027)	0.250	(0.036)
San Jose, Costa Rica (SJO)			−0.099	(0.042)	−0.085	(0.048)
Seoul, South Korea (ICN)			0.169	(0.077)	0.158	(0.081)
Sydney, Australia (SYD)			0.688	(0.186)	0.689	(0.188)
Taipei, Taiwan (TPE)			−0.135	(0.067)	−0.137	(0.068)
Tel Aviv, Isreal (TLV)			−0.117	(0.109)	−0.120	(0.110)
Tokyo, Japan (NRT)			0.178	(0.124)	0.205	(0.138)
Toronto, Canada (YYZ)			0.019	(0.034)	0.009	(0.044)
Vancouver, Canada (YVR)			0.051	(0.032)	0.048	(0.038)
N	1192		1192		921	
R^2	0.011		0.245		0.236	

Notes: The dependent variable is the carrier-route average percent directional price difference. All specifications also include data collection round fixed effects. Standard errors are reported in parentheses and are clustered at the origin-destination pair level to control for correlation across dates and airlines within the same route. The variable RivalU.S.EndpointAirportAllianceShare is equal to zero when the alliance faces no competing carriers on the route. To conserve space, only a selection of the 83 foreign airport fixed effects are reported.

directional fare difference. Estimates of the coefficients on the U.S. airport share of other alliances competing on the route are also positive, suggesting that the fares of competing carriers are strategic complements as is typically the case in oligopoly pricing games. The R-squared measures from these specifications, however, indicate that alliance airport shares explain a very small fraction of the overall observed variation in directional price differences. In contrast, foreign airport fixed effects capture a substantial share of the variation in directional price differences, as indicated by the much larger R-squared values in Columns 2 & 3.

Estimated foreign-airport fixed effects (reported for some cities in Table 4) reveal that average directional price differences can vary substantially across foreign destinations. Destinations in Europe, Australia, Japan, and South Korea often have higher fares on average for travelers originating in the U.S. than for those traveling to the U.S. (i.e. average directional price differences are positive), while destinations in Central and South America, the Middle East, and the Far East (outside of Japan and Korea) often have lower fares on average for U.S. travelers. Exceptions to these generalizations do occur however, and a number of destinations exhibit substantial within-airport variation in directional price differences (across routes, carriers, and collection rounds) resulting in fixed effects with large standard error estimates.

To the extent that the foreign-airport fixed effects absorb any differences in the average level of alliance airport share at the foreign airport, which is unobservable due to data limitations, some of the variation in directional price differences that is explained by the fixed effects might be attributable to differences in alliance airport share between endpoints. The fixed effects will also capture any directional price differences that are associated with differences between the U.S. and the destination country in income or any other characteristics that influence willingness to pay for travel on the route. If the latter form of discrimination is substantial, then directional price differences should be symmetric across carriers so that all the carriers on the route charge lower prices at the same endpoint. Of the 265 routes on which prices are observed for more than one competing airline in the same collection round, only 65 (or 25%) contain one or more airlines with a positive directional price difference and also contain one or more with a negative directional price

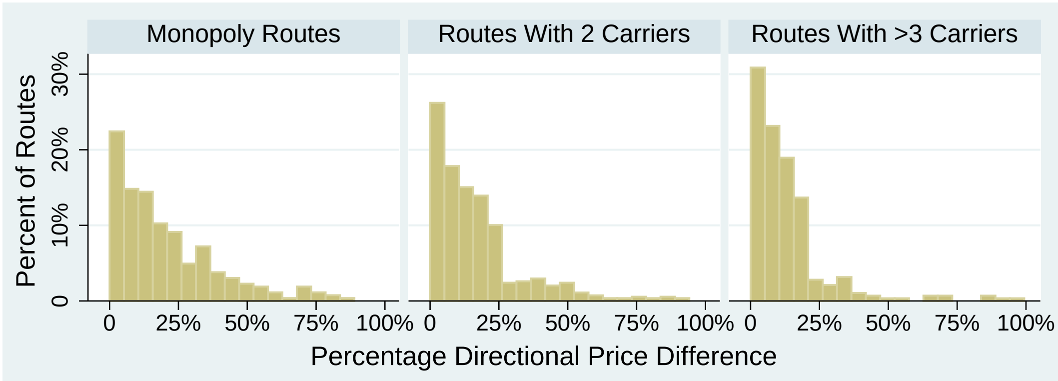

Figure 4: Histograms of Percentage Directional Differences in Fares on International Routes

Notes: Values represent the absolute value of the average directional price difference for each airline alliance (or non-alliance airline) on each route in each collection round, measured as a percent of the average fare. Directional differences greater than 100% have been excluded (accounting for less than 2% of routes).

difference. In other words, the direction of price discounting is symmetric on three-quarters of the non-monopoly routes.

Airlines also appear to charge larger directional price differences on routes where they face less competition. Figure 4 displays histograms of the percentage directional differences in average prices (in absolute value) charged by each airline alliance (or non-alliance airline) on each route flown. Separate histograms are reported for monopoly routes, routes with 2 competing alliances, and routes with 3 or more alliances. Directional price differences are concentrated closer to zero for more competitive routes while larger price differences occur more frequently on monopoly routes. Kolmogorov-Smirnov tests for the equality of any of these distributions are rejected at the 1% significance level, and tests for the equality of variances based on Brown and Forsythe (1974) are rejected at the 2% level. These findings are consistent with the theoretical predictions of Holmes (1989) and Stole (2007) on competition and third-degree price discrimination (with best response symmetry) as well as empirical evidence from Gerardi and Shapiro (2009) that general measures of airline price discrimination tend to decline with competition.

VI Discussion

Unlike other common forms of price discrimination in airline markets, directional price discrimination can be used in markets exhibiting best-response asymmetry where competing carriers disagree on which consumers should be targeted with discounted prices. This is notable because price discrimination under best-response asymmetry often leads to lower prices and lower profits for all firms than when discrimination is not possible.[19] In fact, such cases represent a prisoner's dilemma in which each firm has a dominant strategy to discriminate even though profits would be higher if all firms instead used uniform prices.

The underlying structure of the U.S. domestic airline industry likely generates situations of best-response asymmetry on most routes. Domestic carriers regularly maintain very high airport shares at a select number of airports (particularly at their *hub* airports), and they often fly on routes connecting an airport where they have a large share to an airport where a rival carrier has a large share. Moreover, the empirical evidence clearly shows that consumers strongly prefer carriers who maintain a large presence at their origin airport. Since competing carriers often have larger market shares at opposite endpoints, they may frequently have the incentive to engage in directional price discrimination under conditions of best-response asymmetry.

My empirical analysis of U.S. domestic airline routes, however, reveals that uniform prices (i.e., zero directional price differences) are maintained by every major carrier on virtually all routes. Directional discrimination is absent despite the fact that the costs and complexities of implementing directional differences appear to be minimal. These carriers are known to regularly use other forms of price discrimination, and I show empirically that they often charge different prices for one-way and round-trip itineraries and for itineraries that include a Saturday-night stayover. Moreover, these same U.S. carriers do use directional price differences on most international routes. The systematic and nearly universal avoidance of directional differences in the domestic market instead seems consistent with a coordinated effort to overcome the prisoner's dilemma created by best-response asymmetry

[19] This conclusion holds whenever firms earn most of their sales from their strong market, which I have shown to be the case for U.S. directional airline markets and which is likely to be true in most applied settings.

and avoid the more aggressive form of competition that would result.

The setting exhibits many of the factors known to facilitate tacit coordination. The domestic market contains a relatively small and stable set of carriers, and these carriers repeatedly interact with each other over time and across a large number of different routes. Prices are highly visible and carriers can easily detect and respond to the use of directional price differences by rival carriers. Perhaps most importantly, the elimination of directional discrimination may be much easier to establish and sustain than other tacit arrangements such as the coordination of prices or capacity levels. Prices and capacity differ across routes and often need to be adjusted independently to account for cost and demand fluctuations, whereas a comprehensive strategy to avoid directional price differences would require no adjustment.

International markets provide an interesting setting for comparison. A much larger collection of carriers are present, and routes to different countries or regions are usually served by a varying collection of competitors. Moreover, the populations served at each endpoint are likely to exhibit greater heterogeneity than those on domestic routes. International airline alliances also change the nature of loyalty programs and may impact the extent to which preferences for specific airlines differ between endpoints.

The pervasive use of directional price differences on these routes demonstrates the feasibility of this form of discrimination, but also indicates that coordination by carriers to avoid such directional discrimination may either be more difficult to support or less beneficial in the international setting. Due to international alliances, differences in alliance airport shares between endpoints may be smaller and less-frequently asymmetric.[20] When the overall elasticity of demand for travel on the route (regardless of carrier) is sufficiently lower at one endpoint than the other, the incentive for all airlines to charge higher prices at the same less-elastic endpoint may overcome any asymmetry in preferences generated by asymmetric endpoint market shares. In that case best-response symmetry will result, and directional price discrimination used to offer low prices to travelers originating form

[20]On many international routes alliances will have smaller endpoint airport share differences because it is common for alliances to fly from the U.S. hub of one alliance member airline to the foreign hub of another alliance member airline.

more elastic countries would generate higher prices and profits for all firms (just as in the monopoly setting). The prisoner's dilemma does not arise and coordination to avoid directional price differences would be counterproductive.

The empirical evidence from international routes reveals pricing behavior that is consistent with best-response symmetry on most routes. Directional price differences are symmetric across carriers on over three-quarters of the multi-carrier routes. Moreover, international destination fixed effects explain a much greater share of directional price differences than do the U.S. endpoint airport market shares, suggesting potentially large differences in elasticity for travel between U.S. and foreign endpoints.

Although it is not possible to provide direct proof of tacit coordination by domestic carriers to avoid directional discrimination, the collection of evidence revealed here is consistent with this conclusion. The findings show that directional price differences are almost entirely absent on domestic routes—where they would have likely intensified competition—despite being quite common on international routes where they are more likely to enhance profits. It is important to point out that while this form of coordination could increase industry profits it does not represent a violation of U.S. antitrust law unless carriers explicitly discuss implementing such strategies. If anything, the straightforward nature of this form of coordinated behavior raises the likelihood that such pricing strategies arise and are maintained tacitly from repeated interactions in the marketplace making explicit communication unnecessary.

The novel empirical strategy utilized here is also new to the literature and contributes a more robust approach to identifying price discrimination while controlling for unobserved cost differences. Similar specialized data collection approaches may be useful for investigating other pricing strategies in the airline market and perhaps other similar quoted-price markets.

VII Conclusion

The theoretical literature has established that price discrimination can intensify competition in oligopoly markets under conditions of best-response asymmetry, yet few empirical studies have examined the extent to which this occurs or the ability of firms to avoid this lower-profit equilibrium. This study analyzes the role of competition-enhancing discrimination within the airline industry, where previous work has extensively examined discrimination but has focused entirely on forms of discrimination that reflect best-response symmetry. I identify a form of price discrimination that is very likely to exhibit best-response asymmetry within the U.S. domestic market and show that carriers have successfully avoided engaging in this form of competition-enhancing discrimination despite using similar pricing strategies frequently on international routes. The findings highlight the importance of considering potential coordination in firms' strategic decisions of whether to utilize certain forms of price discrimination. Hopefully, future work will provide additional evidence from other markets to improve our understanding the conditions under which firms engage in (or avoid) competition-enhancing price discrimination.

References

Airline Tarriff Publishing Company, "Advances in Airline Pricing, Revenue Management, and Distribution," Octobert 2017. Discussion Paper.

Asplund, Marcus, Rickard Eriksson, and Niklas Strand, "Price Discrimination in Oligopoly: Evidence from Regional Newspapers," *The Journal of Industrial Economics*, 2008, *56* (2), 333–346.

Azar, Ofer H., "A Linear City Model with Asymmetric Consumer Distribution," *PLOS ONE*, 2015, *10* (6), 1–13.

Borenstein, Severin, "Price Discrimination in Free-Entry Markets," *The RAND Journal of Economics*, 1985, *16* (3), 380–397.

_ , "The Dominant-Firm Advantage in Multiproduct Industries: Evidence from the U. S. Airlines," *The Quarterly Journal of Economics*, 1991, *106* (4), 1237–1266.

_ **and Nancy Rose**, "Competition and Price Dispersion in the U.S. Airline Industry," *The Journal of Political Economy*, August 1994, *112* (4), 653–683.

Brown, Morton B. and Alan B. Forsythe, "Robust Tests for the Equality of Variances," *Journal of the American Statistical Association*, 1974, *69* (346), 364–367.

Chandra, Ambarish and Mara Lederman, "Revisiting the Relationship between Competition and Price Discrimination," *American Economic Journal: Microeconomics*, 2018, *10* (2), 190–224.

Chen, Yongmin, "Paying Customers to Switch," *Journal of Economics & Management Strategy*, 1997, *6* (4), 877–897.

Chen, Yuxin, Chakravarthi Narasimhan, and Z. John Zhang, "Consumer Heterogeneity and Competitive Price-Matching Guarantees," *Marketing Science*, 2001, *20* (3), 300–314.

Corts, Kenneth S., "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment," *The RAND Journal of Economics*, 1998, *29* (2), 306–323.

Dana, James D., "Advance-Purchase Discounts and Price Discrimination in Competitive Markets," *Journal of Political Economy*, 1998, *106* (2), 395–422.

Fudenberg, Drew and Jean Tirole, "Customer Poaching adn Brand Switching," *RAND Journal of Economics*, Winter 2000, *31* (4), 634–657.

Gerardi, Kristopher S. and Adam Hale Shapiro, "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry," *Journal of Political Economy*, 2009, *117* (1), 1–37.

Grennan, Matthew, "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, February 2013, *103* (1), 145–77.

Holmes, Thomas J., "The Effects of Third-Degree Price Discrimination in Oligopoly," *The American Economic Review*, 1989, *79* (1), 244–250.

Luttmann, Alexander, "Evidence of Directional Price Discrimination in the U.S. Airline Industry," *International Journal of Industrial Organization*, forthcoming.

Miller, Nathan H. and Matthew Osborne, "Spatial differentiation and price discrimination in the cement industry: evidence from a structural model," *The RAND Journal of Economics*, 2014, *45* (2), 221–247.

Nevo, Aviv and Catherine Wolfram, "Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals," *The RAND Journal of Economics*, 2002, *33* (2), 319–339.

Puller, Steven L. and Lisa M. Taylor, "Price discrimination by day-of-week of purchase: Evidence from the U.S. airline industry," *Journal of Economic Behavior & Organization*, 2012, *84* (3), 801 – 812.

_ , Anirban Sengupta, and Steven N. Wiggins, "Testing Theories of Scarcity Pricing in the Airline Industry," Working Paper 15555, National Bureau of Economic Research December 2009.

Shaffer, Greg and Z. John Zhang, "Competitive Coupon Targeting," *Marketing Science*, 1995, *14* (4), 395–416.

_ and _ , "Competitive One-to-One Promotions," *Management Science*, 2002, *48* (9), 1143–1160.

Stavins, Joanna, "Price Discrimination in the Airline Market: The Effect of Market Concentration," *The Review of Economics and Statistics*, 2001, *83* (1), 200–202.

Stole, Lars A., "Price Discrimination and Competition," in Mark Armstrong and Robert H. Porter, eds., *Handbook of Industrial Organization, Volume 3*, Elsevier, 2007, chapter 34, pp. 2221–2299.

Thisse, Jacques-Francois and Xavier Vives, "On The Strategic Choice of Spatial Price Policy," *The American Economic Review*, 1988, *78* (1), 122–137.

Villas-Boas, Sofia Berto, "An empirical investigation of the welfare effects of banning wholesale price discrimination," *The RAND Journal of Economics*, 2009, *40* (1), 20–46.

Appendix A: Directional Differences in Domestic Route Shares

As described in Section II, directional pricing on a route is likely to exhibit best response asymmetry and generate lower profits because directional market shares on specific routes are positively correlated with market shares at the origin airport of the round trip. In other words, carriers have greater market shares on routes where they also sell predominantly to frequent flyer members with less elastic demand. Borenstein (1991) provides the most direct evidence of this, showing that the market share of an airline on the round-trip route from A to B to A tends to be larger than its market share on the route from B to A to B when that airline has a greater presence at airport A than at airport B or when its rival carriers on the route have a smaller presence at airport A than at airport B. However, his finding is based on data from 1986. The structure of the airline market and the nature of loyalty programs have both changed dramatically in the last three decades, so I begin my analysis by estimating a basic version of the model in Borenstein (1991) using data from 2016.

The first specification uses a simple directional difference in market share, $\text{RouteShareDiff}_{cr}$, as the dependent variable. Since this dependent variable is bounded within the range $[-1, 1]$, Borenstein instead chooses to use a logistic transformation of the route share variable defined as

$$\text{LogRouteShare}_{cra} = ln\left(\frac{\text{RouteShare}_{cra}}{1 - \text{RouteShare}_{cra}}\right),$$

where the directional difference, $\text{LogRouteShareDiff}_{cr}$, is simply the difference in the values of $\text{LogRouteShare}_{cra}$ for origin airports A and B of route r. A second specification is estimated using this alternative dependent variable.

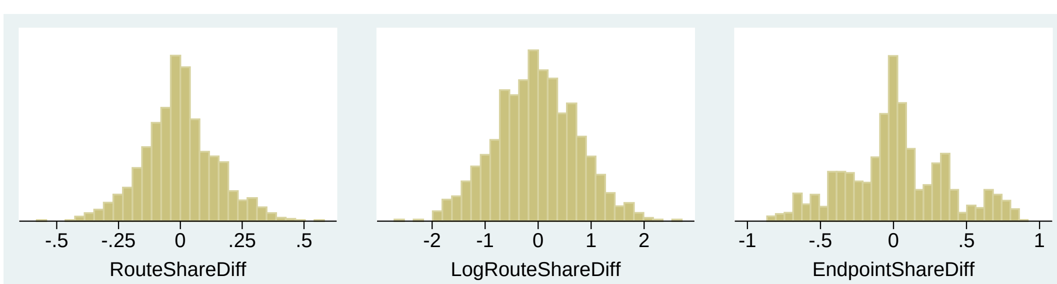

Figure 5: Histograms of Directional Differences in Route Market Shares and Airport Shares

Table 5: Impact of Airline's Origin Airport Dominance on Route Market Share

Dependent Variable:	$\text{RouteShareDiff}_{cr}$ (1)	$\text{LogRouteShareDiff}_{cr}$ (2)
EndpointShareDiff	0.205	1.11
	(0.007)	(0.04)
RivalsEndpointShareDiff	-0.165	-0.85
	(0.007)	(0.04)
constant	-0.002	-0.02
	(0.001)	(0.01)
N	1583	1583
R^2	0.61	0.51

Notes: Standard errors are reported in parentheses and are clustered at the origin-destination pair level to control for correlation airlines within the same route.

I focus my analysis on the top 1000 non-stop domestic routes according to U.S. DOT's DB1B ticket sales data. Any routes that have only one carrier are excluded as route shares in both directions would be 100%. In addition, since the DB1B data is a 10% sample of tickets sold, observations for carrier-route pairs that have less than 500 sampled tickets (implying less than 5000 travelers annually) are dropped to ensure that observed carriers represent legitimate competitors on the route.[21] The data reveal that directional differences in route market shares ($\text{RouteShareDiff}_{cr}$ or $\text{LogRouteShareDiff}_{cr}$) as well as directional differences in endpoint airport market shares ($\text{EndpointShareDiff}_{cr}$) both vary substantially across routes and carriers, as is clearly evident in the histograms displayed in Figure 5. Moreover, the estimates reported in Table 5 from a simple model relating route market shares and airport shares confirm that the patterns identified by Borenstein are still present today. In both specifications, the coefficient estimates on $\text{EndpointShareDiff}_{cr}$ are positive and large and the coefficient estimates on $\text{RivalsEndpointShareDiff}_{cr}$ are negative and only slightly smaller in magnitude.[22] Based on the specification in Column 1, a 10 percentage point increase in origin airport share would yield an increase in route share of 2 percentage points. If that increase in origin airport share comes at the expense of an airline that

[21] This restriction eliminates less than 16% of carrier-route observations. Ticket sales for these carriers are still included in route share calculations for other carriers.

[22] The variable $\text{RivalsEndpointShareDiff}_{cr}$ is calculated using route-share weighted averages of rivals' airport shares. Therefore, it can be viewed as exogenous to carrier c's route share as long as the relative shares of rivals are determined independently from their collective share on the route.

also competes on this route, then the effect on route share would be significantly larger. The logistic specification in Column 2 is more directly comparable to the specifications in Borenstein (1991), and, interestingly, the coefficient estimates are quite similar in size to those obtained by Borenstein roughly three decades ago. Clearly travelers continue to prefer flying on carriers that have a dominant position at their origin airport.

Appendix B: Directional Price Discrimination in a Hotelling Spatial Competition Setting

To more clearly illustrate a specific setting in which best-response asymmetry and competition-enhancing directional discrimination will arise, suppose the competition between airlines can be represented using a Hotelling spatial differentiation model. Denoting A to B to A itineraries as Market A and B to A to B as Market B, let consumers in each market be distributed along the unit interval with measure normalized to one where firm 1 is located at $x = 0$ and firm 2 is located at $x = 1$. To reflect the fact that more consumers will prefer to fly with the carrier that has a higher origin airport presence, the distribution of customers along the interval is assumed to be asymmetric. More specifically, suppose that the population density function of consumers between 0 and 1 in Market A is $f(x)$ and in Market B is $f(1-x)$ so that each firm's position in Market A is identical to their rival's position in Market B. Additionally, assume that $f(x)$ is everywhere decreasing, $f'(x) < 0 \;\; \forall \; x$, and let $F(x)$ represent the cumulative distribution function associated with $f(x)$. If consumers located at x in market m purchase from firm 1 at price p_m^1 they will receive utility $U_m^1 = V - p_m^1 - tx^2$ and if they purchase from firm 2 they receive $U_m^2 = V - p_m^2 - t(1-x)^2$. To ease exposition, assume that both firms face an identical constant marginal cost of c.

Azar (2015) derives a number of properties of the Hotelling model with asymmetric consumer distribution. In particular, he shows that in equilibrium firm 1 will charge a higher price than firm 2 if and only if more than half of the customers are located to the left of the center of the interval (i.e., $F(\frac{1}{2}) > \frac{1}{2}$). In my setting, more than half the consumers will be to the left of center in Market A but the opposite will be true in Market B. As a result, it can be shown based on Proposition 2 of Azar (2015) that best-response asymmetry will result: $b^1_A(p^2) > b^1_B(p^2) \;\; \forall\, p^2$ and $b^2_A(p^1) < b^2_B(p^1) \;\; \forall\, p^1$. Hence, directional price discrimination arises in equilibrium in this setting and the equilibrium price differences between the markets are asymmetric across firms.

When firms are not able to directionally price discriminate and must charge a uniform price across Markets A and B, the density of consumers in the aggregate A+B market becomes symmetric by assumption. Since exactly half the customers are now located on each side of the aggregate market interval, the firms will set equal prices and earn equal profits. Moreover, based on the profit function derived by Azar (2015), the sum of profits earned in Market A and in Market B under directional discrimination will always be smaller than the profits earned under uniform pricing.[23] In other words, this Hotelling spatial model with asymmetrically distributed consumers provides an example of a setting in which the ability to directionally discriminate generates a prisoner's dilemma that results in lower profits for both firms.

While the example above assumes constant marginal costs, actual airlines frequently face binding capacity constraints in the short run as they get close to selling all available seats on their scheduled flights. It is straightforward to show in the Hotelling setting that directional dispersion will remain even when capacity constraints are binding, and this is often true in more general settings. Suppose a firm with constant marginal costs faces residual demand curves $P_1(q_1)$ and $P_2(q_2)$ in Markets 1 and 2 such that $p^*_1 = P_1(q^*_1) > p^*_2 = P_2(q^*_2)$, where q^*_1 and q^*_2 represent the firm's best-responses when facing no capacity constraints. If the firm where to face a capacity constraint of $\overline{q} \leq q^*_1$ it might seem optimal to abandon selling in Market 2 and instead sell its entire capacity in the less elastic market at a price

[23] Since Market A and B are assumed to be symmetric (with inverse consumer density), the following will always hold for market m and its counterpart market $-m$ in equilibrium: $p^{*1}_m = p^{*2}_{-m}$, $q^{*1}_m = q^{*2}_{-m}$, and $\Pi^{*1}_m = \Pi^{*2}_{-m}$. Hence, based on the profit functions derived in Proposition 1 of Azar (2015), the total profit earned by firm i across both markets can be expressed as:

$$\Pi^{*i} \equiv \Pi^{*i}_A + \Pi^{*i}_B = \Pi^{*i}_A + \Pi^{*-i}_A = \frac{2t\left[(q^{*i}_A)^2 + (q^{*-i}_A)^2\right]}{f(x^{m*}_A)}$$

where x^{m*}_A represents the location of the consumer that is indifferent between buying from firm 1 and firm 2 given the equilibrium prices. Since $f(x)$ is decreasing and $x^{m*}_A < \frac{1}{2}$, it must be that $f(x^{m*}_A) > f(\frac{1}{2})$.

If instead firms are restricted to charging uniform prices across Markets A and B, then equilibrium prices and quantities sold will be identical across firms and markets, and the marginal consumer will be located at $x = \frac{1}{2}$. Since $f(x^{m*}_A) > f(\frac{1}{2})$ and also $(q^{*1}_A)^2 + (q^{*2}_A)^2$ is maximized at $q^{*1}_A = q^{*2}_A = \frac{1}{2}$, both firms will earn lower profits Π^{*i} in the directional price discrimination equilibrium than when price discrimination is not possible.

of $P_1(\overline{q}) \geq p_1^*$. However, such a strategy ignores the fact that the capacity constraint increases the marginal opportunity cost of selling in Market 1 to be equal to forgone marginal revenue of selling in Market 2. As a result, selling all output in Market 1 is optimal only when the marginal revenue from Market 1 at $\overline{q}$ as least as large as the willingness to pay of the highest value consumer in Market 2, i.e., if $P_1(\overline{q}) + \overline{q}P_1'(\overline{q}) \geq P_2(0)$.[24] Moreover, if any units are being sold in Market 2, the firm will set q_1 and q_2 to equate the marginal revenues across the two markets, which will result in differential pricing across markets except in the special case where $q_1 P_1'(q_1) = q_2 P_2'(q_2)$. Therefore, if an airline is selling tickets in both directional markets of a given route then it will have the incentive to directionally price discriminate (even if capacity constrained) as long as there are directional differences in residual demand elasticity.

Appendix C: Directional Fare Differences Using DB1B Data

Analysis based on disaggregated fare quote data clearly reveals a lack of directional price discrimination in the U.S. domestic airline market. However, a similar investigation performed using quarterly-average prices from the DOT's DB1B 10% ticket sample database has the potential to yield misleading results due to aggregation and the inability to control for unobserved heterogeneity in costs across the flights of a particular carrier and route within a given quarter. In this appendix, I explore this potential by examining directional differences in average fares computed form DB1B data and estimating whether they vary systematically with differences in endpoint airport market shares.

If airlines were discriminating based on consumers' preference to fly with an airline that has a large presence at their airport of origin, we would expect the $\text{EndpointShareDiff}_{cr}$ variable to be positively correlated with $\text{DirectionalPriceDiff}_{cr}$. Directional price differences could also have been used to discriminate on routes where travelers originating from one of the endpoints have a more elastic demand for travel (regardless of airline) than travelers originating from the other endpoint. For example, willingness-to-pay might vary with

[24]In the Hotelling example proposed above it will always be optimal to set different price levels in Markets 1 & 2 because the marginal revenue in the strong market will always be below V which also represents the maximum willingness to pay in the weak market.

the average income level of the city or the size of the metropolitan area of the traveler. Luttmann (forthcoming) focuses on this second form of discrimination, relating directional price differences observed in DB1B data to endpoint airport demographics. Here I estimate regressions similar to Luttmann (forthcoming), while also including additional endpoint airport share variables that have the potential to indicate competition-enhancing discrimination.

An observation is a unique route-carrier pair, and the sample includes the top 1000 routes as measured by passenger enplanements.[25] The new directional price difference measure is defined as:

$$\textbf{DirectionalPriceDiffDB1B\%}_{cr} = \frac{\tilde{P}^A_{cr} - \tilde{P}^B_{cr}}{\tilde{P}^B_{cr}},$$

where $\tilde{P}^A_{cr}$ and $\tilde{P}^B_{cr}$ represent the average fares of tickets sold on route r by carrier c during the year 2016 that originate out of airport A and airport B respectively. In addition to the DB1B route share and endpoint-airport share variables defined in the paper, I have also collected demographic information for metropolitan areas surrounding endpoint airports from the U.S. Census 2016 American Community Survey. The demographic variables are defined for route r and airport a as follows:

$\textbf{EndpointMSAMedianIncome}_a$**:** The 2010 median income in the MSA surrounding endpoint airport a.

$\textbf{EndpointMSAPopulation}_a$**:** The 2010 population in the MSA surrounding endpoint airport a.

$\textbf{LogMedianIncomeDiff}_r$**:** The difference in the values of $\textbf{ln}(\textbf{EndpointMSAMedianIncome}_a)$ for endpoint airports A and B of route r.

$\textbf{LogPopulationDiff}_r$**:** The difference in the values of $\textbf{ln}(\textbf{EndpointMSAPopulation}_a)$ for endpoint airports A and B of route r.

[25] Carrier-route observations are only included if the carrier has more than 500 tickets appearing in the DB1B 10% sample (implying more than 5000 travelers annually) on that route.

Table 6: Summary Statistics - Endpoint Airport Shares and Demographic Characteristics for top 500 U.S. Domestic Routes

	# of Obs.	Mean	Std. Dev.	Min	Max
EndpointAirportMarketShare	1160	0.284	0.271	0	1
EndpointShareDiff	1160	0.033	0.373	-0.921	0.951
EndpointMSAMedianIncome (in $1,000s)	1160	64.5	9.9	48.3	100.5
EndpointMSAPopulation (in 1,000,000s)	1160	6.37	5.34	0.16	20.03

Directional differences in the DB1B-based price measure are, in fact, substantially different than those observed in the disaggregate fare-quote data. The median size (in absolute value) of the directional differences in average fares is around 6.7% as compared with a median of zero in the fare quote data. Nearly 20% of carrier-route observations exhibit directional fare differences of over 15%.

Results from the regressions relating these DB1B prices to endpoint airport characteristics are reported in Table 7. Columns 3 thorough 6 limit the sample to only the largest 500 routes to more closely match the selection of routes in the fare quote data. All specifications indicate large and statistically significant positive relationships between endpoint airport share differences and directional price differences. Similarly, prices appear to be higher for travelers departing out of larger cites and cities with higher median income. For example, a one standard deviation increase in carrier endpoint airport share is associated with an increase in price on the order of 5 percent, and a 15 percent (i.e., roughly 1 standard deviation) increase in median income is associated with an increase in price of about 2.5 percent. Without fare quote data one might be tempted to conclude from these results that airlines do engage in directional price discrimination. However, given that actual fare quotes did not exhibit any meaningful directional differences, observed differences in the DB1B average ticket prices most likely indicate that travelers from large high-income cities and travelers buying tickets from a carrier that dominates the origin airport more often purchase tickets on relatively expensive flights or purchase closer to the date of travel. Differences in prices across flights operated by a carrier on a route may result in part from price discrimination but are also likely to reflect differences in the (opportunity) cost of

Table 7: Analysis of Directional Differences in DB1B Average Fares

	Top 1000 routes			Top 500 routes		
	(1)	(2)	(3)	(4)	(5)	(6)
EndpointShareDiff	0.124		0.077	0.162		0.113
	(0.007)		(0.007)	(0.012)		(0.009)
RivalEndpointShareDiff	-0.004		-0.053	0.016		-0.037
	(0.012)		(0.011)	(0.017)		(0.013)
LogMedianIncomeDiff		0.163	0.162		0.184	0.182
		(0.015)	(0.015)		(0.019)	(0.019)
LogPopulationDiff		0.025	0.020		0.034	0.030
		(0.002)	(0.003)		(0.003)	(0.003)
constant	0.003	0.002	-0.0003	0.001	-0.002	-0.004
	(0.004)	(0.003)	(0.003)	(0.005)	(0.004)	(0.004)
N	1791	1791	1791	1160	1160	1160
R^2	0.180	0.236	0.361	0.211	0.283	0.442

Notes: The dependent variable is the percent directional difference in DB1B average ticket prices for a carrier on a route. Standard errors are reported in parentheses and are clustered at the origin-destination pair level to control for correlation across airlines within the same route.

operating the flight. As a result, it appears that drawing conclusion about directional price discrimination from DB1B data is likely to be highly misleading.