On the Absence of Directional Price Discrimination in the U.S. Airline Industry*

Matthew S. Lewis†

Draft Date: September 25, 2019

(forthcoming in the Journal of Industrial Economics)

Abstract

Certain forms of price discrimination in oligopoly markets can lead to more aggressive competition and lower profits, yet few empirical studies examine how extensively such strategies are used. I consider one such strategy, testing whether airlines charge different prices on the same flights to passengers that originate from different endpoints. Using fare quote data I formulate a new approach to measure discrimination while controlling for cost heterogeneity and find that carriers within the U.S. domestic market do not engage in directional price discrimination despite frequently using other similar pricing strategies that are unlikely to enhance competition.

*I thank Severin Borenstein, Babur De Los Santos, Nicolas de Roos, Andrew Hanssen, Charles Thomas, John Turner, and seminar participants at the University of Georgia, the University of Indiana, the Clemson IO Workshop, and the International Industrial Organization Conference for helpful comments and feedback. Xiaosong Andy Wu provided invaluable research assistance.

†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 greater 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. Despite an extensive theoretical literature, few empirical studies have directly explored the types of discriminatory pricing strategies that have the potential to intensify competition.[1]

Airlines are known for their extensive use of sophisticated pricing strategies that 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 (e.g., Stavins, 2001; Puller and Taylor, 2012; Puller et al., 2009; Dana, 1998; Borenstein and Rose, 1994; Gerardi and Shapiro, 2009; Chandra and Lederman, 2018). For example, airlines frequently charge more for tickets on the same flights to passengers that purchase 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. Such strategies, however, do not segment consumers based on the likelihood that they value one particular airline over another and, therefore, cannot be used to target discounts to customers who

[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.

prefer a rival carrier.

One strategy that could be used by airlines to target rivals' customers is to price discriminate based on the airport from which a passenger's itinerary originates. Due to frequent flyer and other incentive programs, many travelers prefer to concentrate their trips on one airline. However, maintaining loyalty to an airline is easier when 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 *directionally 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). If competing airlines have a larger presence at opposite endpoints of the route, this form of directional discrimination may result in airlines offering lower prices to their rival's most inelastic customers.

A recent study by Luttmann (2019) presents empirical evidence that directional price discrimination is used on U.S. domestic routes. However, that analysis focuses on directional price differences that arise in response to differences in endpoint city demographics. It also relies on less detailed data in which flight dates and times are unobserved, complicating the identification of directional discrimination relative to other sources of price variation. In this study I empirically reexamine the use of directional price discrimination to investigate whether it is used to target customers that prefer rival carriers.

Utilizing data on airline price quotes for specific flights I implement a new empirical design that isolates fare differences resulting from directional price discrimination. Previous studies of airline price discrimination (including Luttmann, 2019) 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]

In stark contrast to the findings of Luttmann (2019), my analysis reveals that carriers serving domestic routes within the United States do not directionally price discriminate. Over 99% of itineraries exhibit identical prices for passengers originating at different endpoints. Importantly, I confirm that the absence of directional discrimination is not a consequence of airlines moving to a simpler one-way pricing approach. Southwest and other low-cost carriers have long priced round trips as the sum of their one-way fares, and legacy carriers like United, American, and Delta have adopted this approach on some routes. Nevertheless, legacy carriers are shown to avoid directional price discrimination despite utilizing other forms of itinerary-based pricing, like round-trip discounts, on over 45% of observed fares.[3]

Unlike the flight level price quote data used in my analysis, Luttmann uses the U.S. Department of Transportation's Origin and Destination Survey (DB1B) where fares are reported only by the quarter in which travel takes place. With these data it is not possible to identify whether different prices are being charged for the same seats or whether customers in some cities are simply more likely to purchase higher priced flights or tickets (with fewer restrictions, for example) than consumers in other cities. While Luttmann concludes from his findings that airlines do engage in directional price discrimination, I am able to confirm with more detailed data that this does not appear to be the case.

The lack of directional discrimination is interesting given the airlines' frequent use of other forms of itinerary-based price discrimination on these same U.S. routes. It is possible that the potential additional revenue generated by such a strategy is simply not enough to compensate for the costs of implementation. Alternatively, since this particular form of discrimination is unique in that it could have to potential to intensify competition, airlines may avoid directional discrimination in order to prevent the aggressive price competition that can result from targeting rivals' customers. Without detailed data on ticket sales for specific

[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.

[3] This result holds even on routes with more than one legacy carrier and on routes where competing legacy carriers have hubs at opposite endpoints of the route. See Section V. for additional discussion.

flights it is not possible to measure potential revenue gains or determine whether directional discrimination would, in fact, be competition enhancing given the demand characteristics of the market. Nevertheless, the paper concludes with a discussion of additional evidence that may provide clues as to why directional discrimination does not occur in domestic markets. Interestingly, a similar analysis of international routes (from U.S. airports) reveals a very different pattern, with these same airlines engaging in directional price discrimination over 95% of the time. However, the directional price differences observed on international routes appear to be predominantly targeted at large differences (across countries) in the overall elasticity of demand for travel and do not generally exhibit the asymmetry across carriers necessary to intensify competition.

II. Oligopoly Price Discrimination and Airline Pricing

Theoretical studies including Borenstein (1985), Holmes (1989), and Corts (1998) establish that price discrimination in non-monopoly markets is sustainable and can generate higher profits when competing firms agree on which types of consumers are more elastic. On the other hand, 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, Corts (1998) shows that 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. 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 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. Borenstein (1991) 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 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. Lederman (2008) provides additional empirical evidence that frequent-flyer programs specifically are responsible for a significant portion of the price premia earned by hub airlines.

Airlines may be able to capitalize on this geographic heterogeneity in preferences by implementing 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 the other forms of price discrimination commonly studied in airline markets, directional price discrimination has the potential to enhance competition and reduce static equilibrium profits, and firms that repeatedly interact may earn higher profits by sustaining an equilibrium that does not involve directional discrimination.[4]

It is important to note that directional price discrimination can raise profits 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.

[4]A more detailed theoretical example in the spirit of Corts (1998) is presented in Appendix B.

III. Empirical Strategy

The central goal of my empirical analysis is to examine whether and how airlines utilize directional price discrimination. In this case, 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, under specific conditions (which are shown to be strongly supported in the data) it is possible to recreate this comparison.

The necessary assumption (as stated in Assumption 1) is that observed round trip fares on each directional route behave as if the component outbound and inbound legs were priced independent of which leg they are paired with.

Assumption 1 *For each directional route ($A \to B \to A$) there exists a unique latent price ($\tilde{P}$) for each potential flight (X) in the outbound direction ($A \to B$) and for each potential* *flight (Y) in the inbound direction ($B \to A$), such that the observed price of any round trip itinerary (X_Y) in the sample will be the sum of the latent prices of the associated outbound and inbound flights:*

$$P_{X\_Y}^{A \to B \to A} = \tilde{P}_X^{A \to B \to A} + \tilde{P}_Y^{A \to B \to A}. \tag{1}$$

Moreover, if the assumption holds, then for any two outbound flights X and X' and inbound flights Y and Y' on a given directional route it follows that:

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

It is important to highlight that this assumption does not itself rule out directional price discrimination because the latent prices for specific flights are allowed to vary *across* the directional routes that connect two endpoints. On the directional route from A to B to A, the same latent price $\tilde{P}_X^{A \to B \to A}$ must be reflected in the fare for any itinerary involving the specific flight X (from A to B), but a different latent price $\tilde{P}_X^{B \to A \to B}$ may hold for itineraries in the B to A to B route that utilize the same flight X.[5] This fundamental insight drives the identification of directional discrimination.

Note also that this assumption only needs to hold for the sampled itineraries used to identify directional discrimination and does not reflect an assumption about airline pricing behavior more broadly. For example, this condition will be violated when airlines offer discounted fares for itineraries that satisfy minimum-stay or Saturday-night-stay requirements, as the (latent) price associated with a particular outbound flight may change depending on whether it is paired with a qualifying inbound flight. As a result, my empirical investigation specifically utilizes itineraries that are all likely to satisfy the same rules and restrictions (e.g., itineraries that all include a Saturday-night stayover). Fortunately, by comparing the observed itinerary prices of different pairs of outbound and inbound flights, it is possible to demonstrate that the condition in Equation 2 holds in the selected data, providing support for the main identifying assumption.

Given the identifying 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

[5] In fact, this is exactly what is observed on international routes, where I demonstrate that Assumption 1 holds and then reveal extensive directional price discrimination.

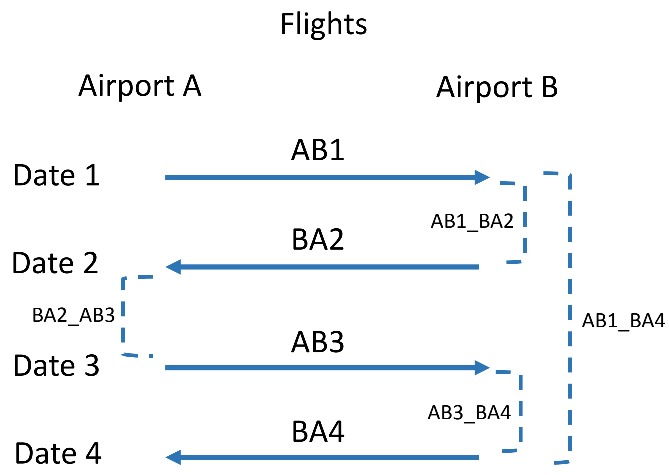

Figure 1: Flights and Itineraries Used to Identify Directional Price Discrimination

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 1 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, following Assumption 1, $P^A_{AB1\_BA2} + P^A_{AB3\_BA4} - P^A_{AB1\_BA4}$ 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 route between A and B can be measured as:

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

By design, this represents a difference in the prices quoted 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 collected from a major airfare aggregator website. The aggregator itself obtains 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[6] from the top 500 most heavily-traveled U.S. domestic routes.[7] 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.

[6]"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.

[7]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.

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 identical to those listed on 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.

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 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 1. 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, maximizing the odds that all within-route price comparisons are appropriately based on fares that were offered concurrently.[8] 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.[9]

A route is defined by its unique paring of endpoint airports A and B, and is not directional. 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.[10] Equation 3 is used to construct 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 be helpful to measure the price difference in percentage terms, which is similarly constructed as follows:

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

[8] Since fares change frequently it is still possible that fares may have occasionally changed during this several-minute collection. It is also possible that the airfare aggregator website does not always update fares instantly, instead relying on a cache of recently accessed fares. In manual observation, fares reported on the aggregator site have always matched those reported on the airline's website. Nevertheless, such deviations are possible and have the potential to artificially generate directional price differences when they are absent. In fact, in Section V. I present evidence suggesting that the few apparent occurrences of directional discrimination observed on domestic routes seem more likely to reflect fare reporting error than actual discrimination.

[9] 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.

[10] In this context, itinerary i actually refers to a combination of four specific round-trip itineraries (with specific flight times) as in Figure 1 that are used to generate one observed P^A and P^B pair.

IV. .1 Evidence Supporting the Identifying Assumption

The proposed approach for measuring directional price discrimination relies crucially on the assumption that, for each directional route, 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 of 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 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.

With only 1 percent of domestic itinerary pairs exhibiting prices that violate As-

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	

sumption 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.[11] 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.

V. Empirical Analysis and Results

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

[11]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.

Table 2: Summary Statistics - Prices and Directional Price Differences

	# 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

data collection rounds which contain itineraries that vary in length of stay and the degree of advance purchase. Table 2 presents summary statistics for prices as well as directional price differences. 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.

Directional price discrimination does not appear to be used by any of the major competitors in the U.S. domestic market (including American, Delta, United, JetBlue, and Alaska).[12] The only airlines in the data that seem to potentially 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 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.

In recent years 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

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

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 legacy carriers like American, United, and Delta have increasingly adopted this strategy, possibly to better compete with carriers like Southwest. However, one-way pricing has not been adopted by all airlines on all routes, and, therefore, does not appear to provide a satisfactory explanation for the industry-wide absence of directional price discrimination.

Additional round-trip and one-way fare quote data from the airfare aggregator website confirm that legacy carriers continue to use more sophisticated fare rules and restrictions to discount round trip travel on many routes. Based on the sampled itineraries, 45% of the round-trip prices offered by legacy carriers are at least 5% cheaper than the sum of the one-way prices of their component flight legs.[13] This fraction is virtually unchanged when focusing specifically on routes with two or more legacy airline competitors or on routes with competing carriers operating hubs at opposite endpoints. Similarly, using a version of the identification strategy depicted in in Figure 1, I find that more than 18% of itineraries are priced at least 5% cheaper when they include a Saturday-night stayover than when they do not.[14] 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 43%.[15] The frequency with which discounts are observed may fluctuate for tickets purchased closer to or further in advance of the date of travel, but these results nevertheless reveal that round trip and Saturday-night stayover discounts still play an important role in some domestic markets.

Although my data do not include direct information on the restrictions associated with each observed fare, the price comparisons above provide evidence that itinerary-based

[13]Similar to the primary dataset, this comparison is based on price quotes contemporaneously collected from the airfare aggregator website during the second week of October for round-trip itineraries departing Friday Oct 26th and returning Monday Oct. 29th 2018 as well as the component one-way itineraries on each of the top 1000 routes.

[14]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 2007. The magnitude of the discount for Saturday-night stayover can then be calculated as: *Discount* $= (P_{AB1\_BA2} + P_{AB3\_BA4} - P_{AB1\_BA4} - P_{BA2\_AB3})/2$.

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

pricing practices such as round-trip discounts are relatively common. To confirm that these price differences actually reflect the frequent use of such restrictions I have collected a limited amount of data from an alternative, more-detailed source. The airfare search platform ITA Matrix (https://matrix.itasoftware.com/) provides a complete breakdown of each quoted fare as well as additional information on fare rules including whether the fare is available only for round-trip travel. For each of the 1000 largest non-stop routes I have collected information on the lowest fare quoted by each airline serving the route for travel on a single set of round-trip travel dates.[16] In these data, 38% of the observed fares offered by legacy airlines were only valid for round-trip itineraries, and 42% of routes have at least one carrier offering a round-trip restricted fare.

In Table 3 data from ITA Matrix is used to provide an example of a route where round-trip restrictions are present and yet directional discrimination does not occur. To calculate directional price differences, fares are collected for a set of four different round-trip itineraries offered by each of the two competing airlines on the route between Houston, TX (IAH) and Minneapolis, MN (MSP).[17] In addition to the round-trip fares, the final column also reports the sum of the one-way fares the airlines were offering at the time for each leg of the round trip. The directional price differences constructed using the 4 round trip fares (following to Equation 3) are zero for both airlines. However, this lack of directional discrimination did not result from the airlines using a simplified one-way pricing structure. The round-trip prices for both airlines are lower than the sum of the available one-way fares on every itinerary.[18] An examination of the fare rules accompanying each of these itineraries confirms that in every case the quoted fares were only available as part of a round trip.

Based on the evidence above, the lack of directional price discrimination does not seem to be a consequence of an industry-wide shift to one-way route pricing. In fact,

[16]Fares were collected on March 30th, 2019 for travel departing on Wednesday April 10th and returning on Wednesday April 17th.

[17]Note that this is another example of a route on which best-response asymmetry is likely to arise. United commands a 78% airport market share at IAH but only 4% at MSP, whereas Delta holds a 70% airport market share at MSP but only 5% at IAH.

[18]These data also include a separate breakdown of all taxes and fees, making it straightforward to confirm that two one-way tickets include the exact same taxes and fees as one round trip.

Table 3: Example of Round-Trip Fares with No Directional Discrimination

Houston, TX (IAH) and Minneapolis-St. Paul, MN (MSP)				
Itinerary	Outbound Departure Time	Inbound Departure Time	Round-Trip Airfare	Combined One-Way Airfares
United Airlines:				
1: IAH to MSP	Apr 5 7:50 AM	Apr 12 3:45 PM	$472	$507
2: IAH to MSP	Apr 19 7:50 AM	Apr 26 5:45 PM	$287	$427
3: IAH to MSP	Apr 5 7:50 AM	Apr 26 5:45 PM	$402	$507
4: MSP to IAH	Apr 12 3:45 PM	Apr 19 7:50 AM	$357	$427
	Implied Directional Difference:		$0	
Delta Airlines:				
1: IAH to MSP	Apr 5 7:45 AM	Apr 12 2:20 PM	$517	$552
2: IAH to MSP	Apr 19 7:45 AM	Apr 26 5:55 PM	$287	$427
3: IAH to MSP	Apr 5 7:45 AM	Apr 26 5:55 PM	$402	$507
4: MSP to IAH	Apr 12 2:20 PM	Apr 19 7:45 AM	$402	$472
	Implied Directional Difference:		$0	

Notes: Following Equation 3, directional differences are calculated as $P_1 + P_2 - P_3 - P_4$. Prices include all taxes and fees. All flights dates are Fridays during 2019.

carriers may never have used directional discrimination on domestic routes. Though data necessary to document this were unavailable at the time, Borenstein (1991) claimed three decades ago that the "fares on a route are virtually always offered without reference to the traveler's point of origin." Prior to the Airline Deregulation Act of 1978, prices for each (non-directional) city-pair were regulated according to a distance-based cost formula, and exceptions were rarely granted. Following deregulation, common practices such as advance purchase requirements and Saturday-night stayover restrictions were rapidly adopted[19], but there is no evidence that directional pricing strategies were ever used. It appears that carriers have long avoided directional discrimination despite consistently using 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. The Department of Transportation's DB1B data used by Luttmann (2019) only reveals the route and the quarter in which the flight was flown, so it is not possible to construct directional dif-

[19] See Schwieterman (1985).

ferences 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.[20]

Using the DB1B data it is straightforward to demonstrate how misleading conclusions could easily be drawn from such an analysis. Adopting the approach of Luttmann (2019) I compute directional differences in the average fares charged by each carrier on each of the top 1000 routes during the year 2016. Unlike in my more detailed fare-quote data, large directional differences in this DB1B-based price measure are common. 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%.

In Appendix C I also estimate regressions similar to those in Luttmann (2019) that relate directional differences in average prices to endpoint city demographics and endpoint airport shares. The results suggest that, within a route, fares tend to be significantly higher for round trips originating from the endpoint city with the highest median income or largest population. Similarly, for each carrier on the route, average fares tend to be higher in the direction originating out of the endpoint where that carrier has a larger airport presence. Based on such results, it would be easy to mistakenly conclude, as Luttmann (2019) does, that airlines offer higher prices to customers originating out of larger or higher-income cities and cities where they have more frequent flyer customers. However, my analysis of more detailed data demonstrates that this is not the case. Directional differences in average prices, like those identified by Luttmann (2019), may arise because customers differ across endpoints in how they respond to the airlines' other discriminatory strategies (based, for example, on advance purchase or travel time) but they do not reflect true directional discrimination.

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

VI. Discussion

My analysis demonstrates that airlines frequently use other forms of itinerary-based discrimination but do not directionally price discriminate. There appears to be something different about this particular form of price discrimination that makes it less attractive.

In theory, there are a variety of cost- and demand-related factors that could cause directional discrimination to be unprofitable. One possibility is that the costs of implementation would be larger than the additional revenue generated. Given the systems in place to execute other itinerary-based discounts, however, substantial technical implementation costs seem unlikely. In addition to round-trip and Saturday-night-stayover discounts, airlines also use "hidden city" discounts which specifically discriminate based on passengers' origin and destination airports by charing more for a direct flight than for some connecting itineraries utilizing the same flight.[21] Directional discrimination could encourage frequent travelers on a route to attempt to construct round trips in one direction by piecing together flights from ticketed round trips in the opposite direction. This has the potential to annoy travelers and lead to increased flight cancelations. However, Saturday-night stayover restrictions generate similar incentives for travelers and yet airlines have used them extensively.

There are other price components on which airlines have chosen not to price discriminate. U.S. carriers charging baggage fees or cancelation fees generally apply the same fee across all routes (within a given ticket class) even though the willingness to pay for such services almost certainly varies. However, potential discrimination involving secondary pricing components (ancillary fees) would likely frustrate consumers by forcing them to check the different fee levels for every flight they are considering. In contrast, directional discrimination works more like other frequently used discriminatory strategies, generating variation in the primary fare component that is readily observable to consumers when considering flights.

The oligopoly price discrimination models of Corts (1998) and others (discussed in Section 2) offer an alternative explanation for why this particular form of discrimination

[21] See Glusac (2019) for a discussion of "hidden city" discounts.

is not used. If best-response asymmetry arises, directional discrimination could potentially result in lower prices and profits for all competitors. In a repeated game environment, there may be an equilibrium in which all competitors earn higher profits by choosing not to directionally discriminate.

Best-response dynamics are determined by the residual demand elasticities that each airline faces for travel out of each endpoint. A variety of factors influence these elasticities. The elasticity of overall demand for travel from each endpoint city may be different, as hypothesized by Luttmann (2019), or the elasticity of demand may vary across airlines as a result of directional differences in origin airport presence, as suggested by Borenstein (1991). Existing evidence suggests that both types of differences are likely to occur at least on some routes.

Luttmann's (2019) analysis identifies sizable directional differences in average prices. Despite being mistakenly characterized as evidence of directional price discrimination, the patterns reflected in his findings (and the replication I present in Appendix C) nevertheless suggest that consumers originating from larger or higher income cities have a significantly less elastic demand for travel and are more willing to purchase more expensive tickets with fewer restrictions.

On the other hand, Borenstein (1991) demonstrates that carriers often have a much larger airport presence at one endpoint of the route than the other and that consumers appear to place significant value on greater origin airport presence. Appendix A confirms this result using updated data from the U.S. DOT's 2016 DB1B dataset. Of particular interest is the $\text{EndpointShareDiff}_{cr}$ variable which captures the difference in a carrier's airport presence across the routes's two endpoint airports.[22] The mean of the absolute value of $\text{EndpointShareDiff}_{cr}$ across all carrier-route pairs is 28 percentage points. In addition, 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 26 percentage points. On a route with two competing airlines, for example, this implies that the $\text{EndpointShareDiff}_{cr}$ of the competitors differ by 52 percentage points on average.

[22]An airline's airport presence is measured by its share of the total passenger enplanements occurring at the airport across all routes. See Appendix A for details.

It is also common for one competitor to have a higher market share at one endpoint while another competitor has a higher market share at the other endpoint. Of routes served by more than one carrier, over 78% have at least one carrier with a negative value of $\text{EndpointShareDiff}_{cr}$ and another carrier with a positive value.[23] In other words, to the extent that willingness-to-pay is influenced by an airline's market share at the originating airport, many routes exhibit the conditions necessary for competing airlines to disagree on whether A-B-A travelers are more elastic than B-A-B travelers. When the effects of directional differences in endpoint airport presence outweigh the effects of directional differences in the overall elasticity of demand for travel, best-response asymmetry will arise, creating the potential for *all-out competition* (in the language of Corts, 1998) as well as the possibility that competing airlines refrain from directionally discriminating in order to avoid earning lower profits.

Although the potential for best-response asymmetry with respect to directional pricing appears to be present on many domestic routes, it is important to note that there are also many routes (including monopoly routes) where this will not occur. The absence of directional discrimination on these routes could suggest that best-response asymmetry and the incentive to avoid competition-enhancing discrimination may not be the primary motivation. However, because airlines frequently enter and exit routes and best-response dynamics are likely to fluctuate over time, there may be a substantial commitment advantage to adopting a network-wide practice of not using directional price differences. Given the repeated interaction of competitors over time, it may be simpler for each airline to make one (easily observable) network-wide decision rather than thousands of route-level decisions on whether to use directional pricing. In addition, many routes have only one airline offering non-stop service, but there are almost always competing airlines offering connecting service, so some degree of competition is always present and best-response asymmetry may still arise.

Since directional discrimination is not observed in the domestic market, direct empirical tests of its impact on prices or profits are infeasible. Moreover, without detailed data

[23] This asymmetry in $\text{EndpointShareDiff}_{cr}$ is often quite sizable. Around 43% of routes having one carrier with an $\text{EndpointShareDiff} > 0.20$ and another carrier with an $\text{EndpointShareDiff} < -0.20$.

on flight-specific ticket sales, demand can not be estimated precisely enough to determine whether directional discrimination would, in fact, be competition enhancing. Nevertheless, some indications can be obtained by examining nondiscriminatory pricing behavior. Appendix D presents a comparison of prices across routes with different market structures and provides evidence suggesting that directional discrimination would intensify competition under the observed market conditions. In particular, the estimates reveal that, when possible, airlines do find it profitable to charge substantially higher prices to consumers originating from the airline's dominant airport.

An examination of pricing patterns on international routes between the U.S. and other countries also provides an interesting comparison that can aid in the evaluation of domestic market pricing behavior. A much larger collection of carriers are present in international markets, 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.

Following the data collection and empirical approach used for domestic routes, I obtain price quotes on international routes for four different sets of travel dates and construct directional price difference measures. A more detailed discussion of this data collection is described in Appendix E. All observed fares are quoted by the airlines in U.S. dollars and are inclusive of all taxes and fees.

Just as the absence of directional price differences was immediately apparent in the domestic route data, basic summary statistics from international routes reveal the exact opposite—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 2 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 gener-

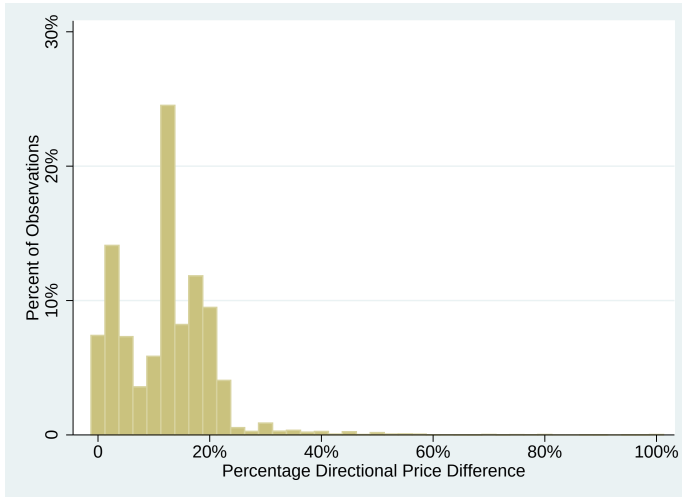

Figure 2: Histogram of the Absolute Percentage Directional Difference in Fares on International Routes

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).

ated by directional differences in taxes or fees. Taxes and fees vary substantially across airports and countries, but are always applied separately to each flight segment and are not determined by the origin of the round trip itinerary.[24] Moreover, 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.

The observed directional price differences from international routes can be used to test whether discrimination is used in a manner more consistent with best-response symmetry or best-response asymmetry. If best response symmetry holds and discrimination is targeted at differences between the U.S. and the destination country in the overall elasticity

[24] Arghyrou et al. (2011) show that airfares quoted in different currencies (for the same flight) do not always match current exchange rates. The round-trip itineraries I study do originate from different countries, but since all collected fares are quoted in U.S. dollars, differences due to exchange rate uncertainty should not be a factor.

of demand for travel on the route, then all carriers on the route would be expected to offer lower prices to passengers traveling in the same direction. In contrast, if competing carriers are offering lower fares in opposing directions, perhaps due to differences in endpoint airport presence, this is more consistent with best-response asymmetry. Of the 265 routes on which prices are observed for more than one competing airline in the same collection round, only 65 routes (or 25%) contain one or more airlines with a positive directional price difference and also contain one or more with a negative directional price difference. In other words, the direction of price discounting is symmetric on three-quarters of the non-monopoly routes.

The frequent use of directional price differences on international routes demonstrates the feasibility of this form of discrimination, however the evidence suggests that it is predominantly used to take advantage of large directional differences in the overall elasticity of demand for travel between U.S. and foreign endpoints. It is possible that directional discrimination is absent in the domestic market simply because such large differences between endpoint cities in the elasticity of demand for travel are much less common. Alternatively, the fact that directional price discrimination is generally best-response symmetric on international routes may explain why it is observed there but not in the domestic market where the likelihood of best-response asymmetry may be higher. Due to international alliances, differences in alliance airport shares between endpoints are likely to be smaller and less-frequently asymmetric than on domestic routes where flights between many major cities connect a hub or concentration airport for one airline to a hub or concentration airport for a competing airline.[25]

VII. Conclusion

The theoretical literature has established that price discrimination can intensify competition in oligopoly markets under conditions of best-response asymmetry and that repeated game equilibria may exist in which firms earn higher profits by not discriminating. Despite

[25] 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. Since alliance airlines have compatible frequent flyer programs, consumer willingness to pay is probably more closely tied to the airport presence of the alliance than the specific airline.

this, few empirical studies have examined the extent to which such discrimination occurs or the ability of firms refrain from competition-enhancing discrimination. In this study I investigate a directional price discrimination strategy that has the potential to exhibit best-response asymmetry within the U.S. domestic airline market and show that carriers do not engage in this form of discrimination despite frequently using other similar (best-response symmetric) pricing strategies. This finding is particularly notable in light of a recent study by Luttmann (2019) who concludes based on more aggregate data that domestic airlines do engage in directional discrimination. The contrasting findings emphasize the importance of the novel empirical strategy I employ, which utilizes price quote data to more completely control for unobserved cost differences.

While I am not able to directly test between possible explanations for the lack of directional discrimination, supplementary evidence suggests that it may be the result of an attempt to avoid engaging in competition-enhancing discrimination. The costs of implementing directional discrimination appear to be low and past evidence has demonstrated that consumers prefer to use airlines that have a large presence at their origin airport, creating the potential for best-response asymmetry to arise on many routes. This possibility is also consistent with the evidence that directional discrimination is used extensively on international routes where best-response asymmetry is less likely.

My findings highlight the importance of more carefully considering 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) different discriminatory pricing strategies. In addition, the specialized data collection strategy introduced here provides a more robust approach to identifying price discrimination and may be useful for investigating other pricing strategies within the airline market and other similar quoted-price markets.

References

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

Arghyrou, Michael G., Andros Gregoriou, and Panayiotis M. Pourpourides, "Risk aversion, exchange-rate uncertainty, and the law of one price: insights from the market for online air-travel tickets," *Canadian Journal of Economics*, 8 2011, *44* (3), 880–906.

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.

_ , "Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry," *The RAND Journal of Economics*, 1989, *20* (3), 344–365.

_ , "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.

Brueckner, Jan K., Darin Lee, and Ethan S. Singer, "Airline competition and domestic US airfares: A comprehensive reappraisal," *Economics of Transportation*, 2013, *2* (1), 1 – 17.

_ **, Nichola J. Dyer, and Pablo T. Spiller**, "Fare Determination in Airline Hub-and-Spoke Networks," *The RAND Journal of Economics*, 1992, *23* (3), 309–333.

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.

Glusac, Elaine, "Airline Booking Hacks: What's Dandy, What's Dicey: Travel Desk," *New York Times*, 2019, *March 31*.

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.

Lederman, Mara, "Are Frequent-Flyer Programs a Cause of the "Hub Premium"?," *Journal of Economics & Management Strategy*, Spring 2008, *17* (1), 35–66.

Luttmann, Alexander, "Evidence of Directional Price Discrimination in the U.S. Airline Industry," *International Journal of Industrial Organization*, January 2019, *62*, 291 – 329.

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," December 2009. National Bureau of Economic Research, Working Paper 15555.

Schwieterman, Joseph P., "Fare is Fair in Airline Deregulation: The Decline of Price Discrimination," *Regulation*, 1985, *9*, 32.

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.

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 has the potential to exhibit best response asymmetry and generate lower profits because preferences for airlines on specific routes are positively correlated with origin airport market shares. 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 it is useful to re-explore this relationship by estimating a basic version of Borenstein's (1991) model using data from 2016.

Data on 2016 aggregate ticket sales and passenger enplanements are collected 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:

$\mathbf{EndpointAirportMarketShare}_{cra}$**:** 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.

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

$\mathbf{RouteShare}_{cra}$**:** 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.

Table 4: Summary Statistics - Endpoint Airport Market Shares

	# of Obs.	Mean	Std. Dev.	Min	25th %tile	75th %tile	Max
EndpointAirportMarketShare	3104	0.255	0.236	0	0.059	0.404	1
RouteShare	3104	0.410	0.255	0.008	0.195	0.601	0.995
EndpointShareDiff	1552	0.024	0.366	-0.916	-0.214	0.272	0.919
RouteShareDiff	1552	0.0005	0.153	-0.581	-0.092	0.086	0.581

Since empirical identification is based on studying directional differences, the values at the two endpoints airports of the route r are differenced for each of the variables above 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.

Summary statistics are reported in Table 4.

The first specification uses the 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.

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

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

Dependent Variable:	$\text{RouteShareDiff}_{cr}$ (1)	$\text{LogRouteShareDiff}_{cr}$ (2)
EndpointShareDiff	0.219	1.10
	(0.007)	(0.04)
RivalsEndpointShareDiff	-0.162	-0.84
	(0.007)	(0.04)
constant	-0.0004	-0.02
	(0.001)	(0.01)
N	1548	1548
R^2	0.63	0.54

Notes: 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.

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.[26] 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 3. 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

[26]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.

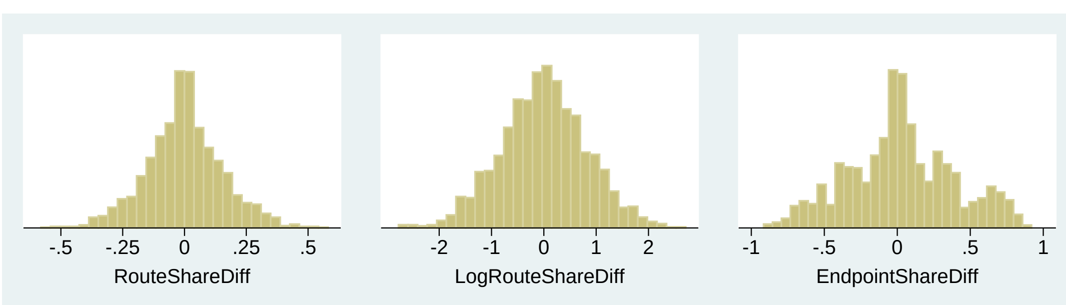

Figure 3: Histograms of Directional Differences in Route Market Shares and Airport Shares

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.[27] 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 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: A Theoretical Example of Directional Price Discrimination

A simple theoretical example in the spirit of Corts (1998) illustrates how directional price discrimination can reduce static equilibrium profits. 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.[28] 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 B to A to B market, and the reverse is true for firm 2:

$$\begin{aligned} b^1_{A\_B\_A}(p^2) &> b^1_{B\_A\_B}(p^2) \quad \forall\, p^2 \\ b^2_{B\_A\_B}(p^1) &> b^2_{A\_B\_A}(p^1) \quad \forall\, p^1. \end{aligned} \tag{4}$$

[27] 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.

[28] 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).

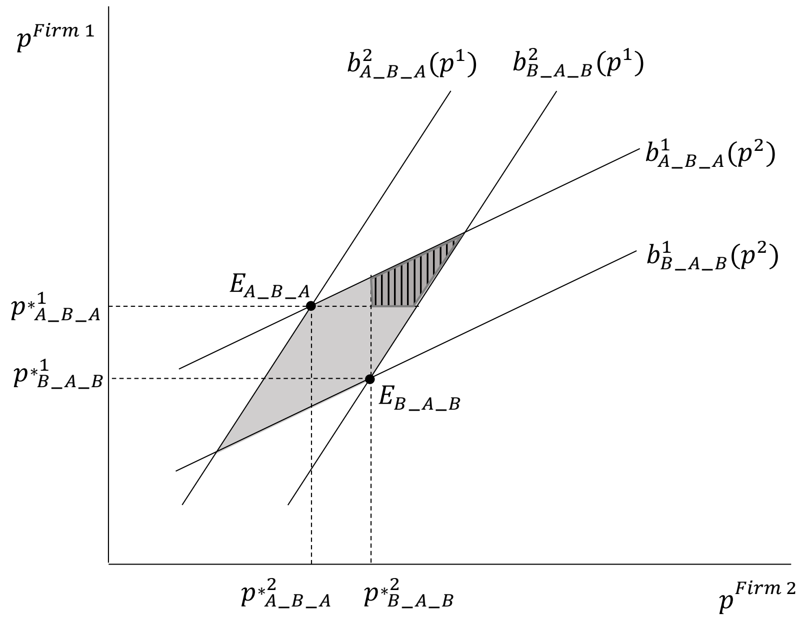

Figure 4: Directional Price Discrimination with Best-Response Asymmetry

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 4. 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 4 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.

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^1_m they will receive utility $U^1_m = V - p^1_m - tx^2$ and if they purchase from firm 2 they receive $U^2_m = V - p^2_m - t(1-x)^2$. 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.[29] 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 of

[29] 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.

$P_1(\overline{q}) \geq p_1^*$. 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 $\overline{q}$ is at least as large as the willingness to pay of the highest value consumer in Market 2, i.e., if $P_1(\overline{q}) + \overline{q}P_1'(\overline{q}) \geq P_2(0)$.[30] 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 from the 2016 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 (defined in Appendix A 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-

[30]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.

pay might vary with the average income level of the city or the size of the metropolitan area of the traveler. Luttmann (2019) 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 (2019), 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.[31] The new directional price difference measure is defined as:

$$\mathbf{DirectionalPriceDiffDB1B\%}_{cr} = \frac{\tilde{P}^{A}_{cr} - \tilde{P}^{B}_{cr}}{\tilde{P}^{B}_{cr}},$$

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:

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

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

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

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

Summary statistics for the price, airport share, and demographic variables are reported in Table 6.

[31] 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
DirectionalPriceDiffDB1B%	1190	0.016	0.129	-0.378	0.491
EndpointShareDiff	1190	0.033	0.373	-0.921	0.951
EndpointMSAMedianIncome (in $1,000s)	1190	64.5	9.9	48.3	100.5
EndpointMSAPopulation (in 1,000,000s)	1190	6.37	5.34	0.16	20.03

Table 7: Analysis of Directional Differences in DB1B Average Fares

	Top 1000 routes			Top 500 routes		
	(1)	(2)	(3)	(4)	(5)	(6)
EndpointShareDiff	0.119		0.077	0.153		0.108
	(0.008)		(0.007)	(0.013)		(0.010)
RivalEndpointShareDiff	-0.009		-0.053	0.006		-0.042
	(0.013)		(0.011)	(0.018)		(0.014)
LogMedianIncomeDiff		0.153	0.162		0.180	0.177
		(0.014)	(0.015)		(0.018)	(0.018)
LogPopulationDiff		0.030	0.020		0.038	0.035
		(0.002)	(0.003)		(0.003)	(0.004)
constant	0.010	0.009	-0.0003	0.010	0.005	0.004
	(0.004)	(0.003)	(0.003)	(0.005)	(0.003)	(0.003)
N	1879	1879	1879	1190	1190	1190
R^2	0.158	0.269	0.375	0.189	0.305	0.450

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.

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 4 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 4 percent, and a 15 percent (i.e., roughly 1 standard deviation) increase in median income is associated with an increase in price of about 2.7 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 operating the flight. As a result, it appears that drawing conclusion about directional price discrimination from DB1B data is likely to be highly misleading.

Appendix D: Competitive Effects of Directional Discrimination

Since directional discrimination is never observed in the domestic market it is difficult to directly evaluate the extent to which competition-enhancing discrimination lowers prices. Nevertheless, comparisons across routes with different market structures may provide some indication of the potential competitive effects.

First, consider a hypothetical route between airports A and B that is served by two carriers, and suppose both carriers have a much higher airport share at A than at B. In this case best-response symmetry is likely to arise across directional markets, as both airlines face a less-elastic residual demand at airport A than at B, but neither airline will have an advantage in either directional market (A to B or B to A). So, if the airlines do not engage in directional price discrimination and both sell fairly equal quantities in the A to B market and B to A market, then the equilibrium non-discriminatory prices will lie roughly halfway between prices that would be observed in A to B and B to A under directional discrimination.

Next, consider the case where one carrier has a much higher market share at A while the other has a higher much market share at B. Here best-response asymmetry is likely to arise across the directional markets, as one airline will face a less-elastic residual demand at airport A and the other at B. In this case, the first airline will sell most of the tickets in the A to B route due to its large presence at airport A, and the second airline will sell most of the tickets in the B to A route. As a result, if directional price discrimination is not used, the first airline will likely set a price relatively close to what it would have charged in the A to B route under discrimination, because most of the tickets it sells will be in the A to B market. Similarly, the second airline will set price close to what it would have charged in the B to A market under discrimination.

In the first case, when the airlines have larger presence in the same airport, they compete more directly and are not able to raise prices to take advantage of inelastic (loyal) travelers out of their dominant airport. In the second case, each airline predominantly sells to consumers originating from their own dominant airport, so they compete less directly and are able to charge higher prices. This prediction can be tested empirically by comparing prices observed on duopoly routes where the competing airlines have hubs at the same endpoint to prices observed on routes where the competing airlines have their hubs at opposite endpoints.

I perform this comparison using fares reported in the 2016 DB1B (described in Section IV.)[32] for nonstop routes served by exactly two of the following carriers: American, United, Delta, Alaska, Southwest, and JetBlue.[33] Since directional discrimination does not occur in the domestic market, prices are averaged to the level of the airport pair (i.e., routes are bi-directional) for each airline. Similar to Borenstein (1989), Brueckner et al. (1992), Brueckner et al. (2013), and others, I estimate the following reduced form model of equilibrium fares that includes a variety of cost- and demand-related route characteristics

[32] Data from the DB1B are used because they offers a more accurate measure of average prices to be compared across routes. In contrast, my fare quote data only include the prices for travel on a few select days.

[33] Smaller airlines are not considered because their demand and pricing strategies may be substantially different from the major airlines. Routes where smaller airlines are also present are not considered in this sample.

as control variables:

$$ln(PricePerMile_{cr}) = \delta SameHub_r + \mathbf{X}_{cr}\beta + \lambda_c + \epsilon_{cr}.$$

For carrier c on route r, λ represents the carrier fixed effect, $\mathbf{X}$ is the vector of route characteristics, and $SameHub$ is and indicator equal to 1 if the carriers on the route have a hub at the same endpoint and 0 if their hubs are at opposite endpoints.

The vector $\mathbf{X}$ includes the distance and squared distance of the route since cost per mile typically declines with distance at a decreasing rate. The total number of enplanements on the route is also included because cost may be lower on higher-volume routes. While there is some potential for enplanements to be endogenous, the effect is likely to be small. The price measured in the dependent variable is only faced by non-stop travelers, while enplanements include all travelers connecting to or from other destinations. In addition, cross-route variation in enplanements likely results more from variation in the underlying demand for travel between the endpoints than from responses to variation in fares. Other demographic measures such as the average median income and population at the endpoints and the absolute difference across endpoints in median income and population are also included in some specifications to examine robustness.

The sample of routes used for estimation is also selected to improve comparability. Routes are only used if they serve over 20,000 passengers per year (i.e., 2,000 tickets observed in the DB1B 10% sample) and the share of route passengers originating out of each endpoint is at least 30%. The second condition excludes routes involving popular tourist destinations like Las Vegas or Orlando where most passengers are flying in and relatively few are flying out.

The OLS coefficient estimates are reported in Table 8. The coefficient on $SameHub$ is negative and significant, indicating that average prices are about 8% lower on duopoly routes where the airlines both have a hub at the same endpoint airport than on routes where the airlines have hubs at opposite endpoints. This estimate is robust to the inclusion of carrier fixed effects and additional income and population control variables. While the evidence is certainly not definitive, higher profit margins on routes with hubs at opposite endpoints suggests that airlines find it profitable to charge more when pricing only (or

Table 8: Impact of Hub Location on Prices for Duopoly Routes

Dependent Variable: ln(Average Ticket Price Per Mile)	(1)	(2)	(3)
Hubs at Same Endpoint	-0.089	-0.086	-0.078
	(0.037)	(0.032)	(0.034)
Route Distance (in 1,000s of miles)	-1.214	-1.223	-1.251
	(0.078)	(0.062)	(0.058)
(Route Distance)^2	0.177	0.184	0.185
	(0.018)	(0.014)	(0.013)
Total Enplanements on Route (in 1,000s)	-0.068	-0.054	-0.059
	(0.011)	(0.010)	(0.010)
(Total Enplanements on Route)^2	0.002	0.002	0.002
	(0.0004)	(0.0004)	(0.0003)
Average Median Income (in $10,000s)			-0.001
Across Enpoints			(0.026)
Average Population (in millions)			0.027
Across Enpoints			(0.011)
Absolute Difference in Median			-0.040
Income Across Enpoints			(0.017)
Absolute Difference in Population			-0.015
Across Enpoints			(0.006)
Carrier FEs (omitted = American Airlines)			
Alaska		-0.236	-0.190
		(0.038)	(0.037)
JetBlue		-0.322	-0.276
		(0.036)	(0.038)
Delta		0.035	0.043
		(0.030)	(0.029)
United		0.003	-0.001
		(0.026)	(0.024)
Southwest		-0.195	-0.187
		(0.029)	(0.027)
N	286	286	284
R^2	0.861	0.904	0.914

Notes: 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.

mainly) to customers flying out of their hubs. The pricing patterns appear to reveal substantial differences in the residual elasticity of demand between travelers depending on whether they are originating from the airline's hub. In this setting, directional discrimination would be expected to produce substantially lower prices.

Appendix E: International Route Data

Price quotes for international routes were obtained from the same airfare aggregator website 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. 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.

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 (or code share) serve as both ticketing and operating carriers on their flights, so all quotes by these airlines are also included.

Fares from these international routes appear to largely satisfy the main identifying assumption (Assumption 1). To confirm this, I repeat the same assumption test that was performed for domestic routes 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. As was done with domestic routes, I exclude 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 international sample includes prices for 1192 airline-route pairs from 917 different routes representing 287 unique origin-destination pairs.

Directional price difference measures are constructed exactly as they are for domestic routes. As discussed in Section VI. , over 95% of observed directional price differences are nonzero, and in most cases the magnitudes of these differences are substantial. 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.