Endogenous and Selective Service Choices After Airline Mergers

Sophia Li
Cornerstone Research

Joe Mazur
Purdue University

Yongjoon Park
University of Maryland

James Roberts
Duke University and NBER

Andrew Sweeting*
University of Maryland and NBER

Jun Zhang
University of Maryland

January 2018

Abstract

We estimate a model of service choice and price competition in airline markets, allowing for the carriers that provide nonstop service to be a selected subset of the carriers competing in the market. Our model can be estimated without an excessive computational burden and we use the estimated model to illustrate the effects of selection on equilibrium market structure and to show how accounting for selection can change predictions about post-merger market power and repositioning, in ways that are consistent with what has been observed after actual mergers, and possible merger remedies.

Keywords: endogenous market entry, selection, horizontal merger analysis, static games, airlines

JEL Codes: C31, C35, C54, L4, L13, L93

*Corresponding author, sweeting@econ.umd.edu, who conducted some of this research during a five-month spell as an Academic Visitor at the US Department of Justice, whose hospitality is warmly acknowledged, although the paper does not reflect the views, opinions or practice of the Department. We thank a number of seminar participants and discussants for useful comments. The research has been supported by NSF Grant SES-1260876. An earlier version of this paper was circulated as "Airline Mergers and the Potential Entry Defense". Peichun Wang provided excellent research assistance during an early phase of the project. The usual disclaimer applies.

1 Introduction

When mergers are proposed in differentiated product markets, the antitrust authorities need to evaluate not only how much market power might be created holding fixed the set of available products, but also whether the merger might lead other firms to enter or to reposition their products in a way that would be "timely, likely and sufficient" (Section 9 of the 2010 *Horizontal Merger Guidelines*) to prevent increased market power from being exercised. While equilibrium models that assume static Bertrand Nash pricing, in the spirit of Nevo (2000), are widely used to guide the first part of the evaluation, assessments of repositioning, especially by rivals, are typically based on less formal analyses of historical repositioning and rivals' likely business plans. While the lack of formal modeling may seem surprising given the large literature on discrete choice "entry games" in Industrial Organization, it reflects the fact that most of this literature has failed to link entry and post-entry competition in a way that allows the likelihood of repositioning and its sufficiency in constraining market power to be convincingly quantified.

In this paper, we develop and estimate an integrated model of positioning and price competition and use it to analyze endogenous service choices and competition after mergers in the airline industry. Our service choice involves a carrier deciding whether to offer nonstop or connecting service on a particular route. Our model has a standard two-stage structure where carriers choose their type of service and then choose equilibrium prices. The distinction between nonstop and connecting service has been important in the analysis of airline mergers (Dunn (2008))[1], even though it has often been ignored in the academic literature. We assume that carriers have complete information about the qualities and costs associated with different service choices of all carriers throughout the game. This implies that the carriers that choose nonstop service will be a selected subset of the carriers competing in the market, and, in particular, carriers that choose connecting service will tend to be less effective nonstop competitors (lower quality or higher cost) if, for some reason, they had to change their service type.

Our paper makes two major contributions. First, we use our estimated model to illustrate how selection affects equilibrium market structure and how considering selection can impact the analysis of mergers and potential merger remedies. When there is no selection market structure

[1] See also the Department of Justice's 2013 Competitive Impact statement on the American Airlines/US Airways merger, https://www.justice.gov/atr/case-document/competitive-impact-statement-219 (accessed June 26, 2017), and the US Government Accountability Office's 2010 report on the United Airlines/Continental Airlines merger, http://www.gao.gov/new.items/d10778t.pdf (accessed June 26, 2017).

is tightly linked to the level of demand in the market. As a result, the elimination of a nonstop carrier when nonstop duopolists merge will likely induce another carrier to initiate nonstop service, and, because there is no selection, the new nonstop carrier is likely to be viewed as an effective competitor in the sense of being an effective constraint on the prices of the merged firm. However, when we account for the selection implied by both carriers' observed characteristics and their pre-merger service choices, we predict that new nonstop service is less likely and, if it occurs, it will tend to be less effective at preventing price increases. In fact, for a set of nonstop duopoly mergers, we predict post-merger price increases that are quite similar to those in a model with fixed service types. This is partly because we predict that nonstop service would be initiated in only 20% of markets, a rate which is very similar to the observed rate following mergers that took place after our sample period. However, it is also because new nonstop carriers will tend to be less effective nonstop competitors than carriers that choose to provide nonstop service prior to the merger. This is illustrated by considering a remedy, where American Airlines offered to commit to initiate nonstop service on several routes, which was proposed when United and US Airways attempted to merge in 2000. Under this remedy the number of nonstop competitors would not have fallen, and, when selection is completely ignored, this remedy appears effective as a way for preventing prices from increasing. However, when we account for selection on both observables and unobservables, we predict that the merged carrier would increase its prices by 6.5%, which is similar to the 7.8% price increase predicted without the remedy (where the probability that American or any other carrier would initiate nonstop service is low).

The second contribution comes from the fact that we estimate our selective entry model without an excessive computational burden. With selection, the estimation of demand and marginal cost functions cannot be separated from the estimation of the discrete service choice model. A nested fixed point routine, of the type typically used to estimate discrete choice games, would require repeatedly solving games where firms make both discrete and continuous choices. Estimation would be further complicated by the possible existence of multiple equilibria and the discontinuity of simulated objective functions resulting from the discrete nature of service choices. Taken together, these issues create an excessive computational burden unless the number of players is constrained to be very small and very simple demand and cost specifications are used. Instead, we approximate a set of moments using importance sampling, following Ackerberg (2009). To do so, we set up a model that allows for rich, and plausible, cross-carrier and cross-market heterogeneity and then solve a large number of simulated games with different demand and cost draws for different firms. During estimation of the structural parameters, we approximate moments by re-weighting the outcomes of interest from the simulated games, which only involves multiplying a set of probability density functions.[2] The resulting objective function is smooth, which allows the use of standard minimization routines. While we focus on a model where service choices are made in a known sequential order to avoid multiplicity of equilibria, we show that our parameter estimates are robust to allowing for simultaneous moves or an unknown sequential move order.[3]

Before discussing related literature, we identify two broad limitations of our analysis. First, our model is static rather than dynamic. One way in which this matters is that we do not allow for carriers who are not active in a market at all to begin operations once a merger has taken place, or for the merged or non-merging carriers to significantly re-configure their networks.[4] While these responses could have economically important effects on market power and welfare in the long-run, a static model, which enables us to use richer specifications, is more consistent with the short-run focus of most merger analysis.[5] Our static approach also rules out the possibility that carriers engage in any form of dynamic limit pricing to deter entry or changes in service types. While Sweeting, Roberts, and Gedge (2017) provide evidence of dynamic limit pricing on a subset of routes with a dominant incumbent carrier, in this paper we are focused on routes where mergers may significantly reduce competition. Second, we do not model choices of route-level capacity or schedules, which means that we may attribute some differences in carrier market shares to unobserved quality and costs when they really reflect strategic capacity or flight

[2]Approximation will entail some loss of efficiency and importance sampling approximations will only be consistent under some conditions (Geweke (1989)), which we test in Appendix B.

[3]Here we make a small innovation. The current literature that allows for multiplicity in the estimation of static discrete choice games (e.g., Ciliberto and Tamer (2009), Sweeting (2009), Wollmann (2016)) has assumed that the equilibrium played will be one of the pure strategy equilibria in a simultaneous move game. We allow for the equilibrium to be either one of these equilibria or an equilibrium in a sequential game where the order is unknown. While the set of equilibrium outcomes from simultaneous and sequential games are often identical, this is not always the case.

[4]In an earlier version of this paper (Li, Mazur, Roberts, and Sweeting (2015)) we estimated a model where carriers made trinomial choices to provide connecting service, to provide nonstop service or to not serve the market at all, whereas in this paper we focus on the decision of carriers who do serve the market to provide connecting or nonstop service. The richer model had a greater computational burden, and the decision to provide connecting service, rather than no service, was estimated to be quite random, which is likely explained by the fact that the definition of whether carriers are connecting or not serving a market often depends on arbitrary thresholds for carrying enough traffic to be considered a competitor (see discussion in Section 2). As a result, the estimates and the counterfactuals were harder to interpret than when we use a binary connecting/nonstop service decision.

[5]Aguirregabiria and Ho (2010), Aguirregabiria and Ho (2012) and Benkard, Bodoh-Creed, and Lazarev (2010) consider long-run dynamic models of the airline industry.

scheduling choices. We hope to extend our model to allow for these choices in future work, and a computationally-light approach to estimation will be even more important when we do so.

The rest of the Introduction briefly discusses the related literature. Section 2 outlines the data and explains how we define several important variables. Section 3 describes the model, while Section 4 describes estimation and discusses identification. Section 5 presents the parameter estimates both with and without a known order of entry, and assesses the fit of the model. Section 6 quantifies the extent of selection implied by our estimates and the implications of selection for market structure. Section 7 presents our analysis of merger counterfactuals under different selection assumptions. Section 8 concludes. The Appendices, which contain more details of the data and estimation, are available online.

Related Literature

Ashenfelter, Hosken, and Weinberg (2014) summarize the literature on the effects of consummated airline mergers on route-level prices. Prior to 1989, mergers were regulated by the Department of Transportation, which allowed all proposed mergers partly based on the theory that the threat of new entry or service changes would constrain post-merger prices increases (Werden, Joskow, and Johnson (1991)). Several papers have estimated that prices increased after mergers during this period, although magnitudes vary depending on the chosen time-window and control group.[6] Analysis of more recent Department of Justice-approved mergers has provided more mixed results. Hüschelrath and Müller (2014) and Hüschelrath and Müller (2015) identify short-run price increases of as much as 10% after recent mergers, suggesting significant increases in market power, although Israel, Keating, Rubinfeld, and Willig (2013) suggest that the expansion of the merged carriers' networks may have increased consumers' willingness to pay. The assessment of recent mergers may be complicated by allegations of price collusion or coordination between the largest carriers (Ciliberto and Williams (2014), Azar, Schmalz, and Tecu (forthcoming)) from 2008 onwards. Our model assumes non-cooperative behavior so we

[6]For example, several papers have measured the effects of the 1986 Northwest/Republic and TWA/Ozark mergers, both of which involved mergers of carriers that had hubs at the same airports. Borenstein (1990) estimated that these mergers increased prices, on routes where both carriers had provided service and no other carriers were active, by 6.7% and -5.8% (i.e., a decrease) respectively. Werden, Joskow, and Johnson (1991) provide evidence that prices rose after both mergers, although only slightly in the case of TWA/Ozark. Peters (2009) finds that prices increased after both mergers, but by more after TWA/Ozark. Morrison (1996) finds that prices fell after Northwest/Republic in the short-run but increased in the long-run, with the opposite effect in TWA/Ozark.

estimate our model using data from 2006.

The second closely related literature concerns the estimation of entry games. Most of the early literature (*inter alia* Bresnahan and Reiss (1991), Bresnahan and Reiss (1990), Mazzeo (2002), Seim (2006) and, considering airline markets, Berry (1992) and Ciliberto and Tamer (2009)) estimated reduced-form payoff functions without a clear link to prices or consumer surplus. Subsequent work has tried to integrate models of entry and competition, introducing the challenges outlined above. A common approach, for example Draganska, Mazzeo, and Seim (2009), Eizenberg (2014), Wollmann (2016) and Fan and Yang (2016), excludes selection by assuming that firms have no information on unobserved demand or marginal cost shocks when entry or service choices are made.[7] This assumption allows demand and marginal cost functions to be estimated separately from the entry game. However, it means that some firms may regret their first-stage choices ex-post, which is unsatisfactory if the data is to be interpreted as reflecting an industry in steady-state equilibrium, and, as our results suggest, it may lead to merger analysis to generate misleading results if selection is actually present.[8]

The airline entry papers of Reiss and Spiller (1989) and the working paper by Ciliberto, Murry, and Tamer (2016) (CMT, hereafter) are especially closely related. Reiss and Spiller estimate a model of service choice and subsequent price competition in airline markets, and they distinguish between nonstop and connecting service for reasons that are very similar to ours. They create a manageable computational burden, and side-step the issue of multiple equilibria, by making carriers symmetric, conditional on service choice, and assuming that only one carrier can provide nonstop service. In this paper we make more flexible assumptions about both carrier heterogeneity and service choices, which is possible because of thirty years of advances in computing technology.[9] CMT and the current paper were developed contemporaneously. CMT also estimate a complete information model of entry and competition in route-level airline markets

[7] Related work, most notably Fan (2013), has examined how mergers may affect continuous measures of quality, as well as price. An advantage of analyzing continuous choices is that equilibrium choices will be determined by a set of first-order conditions, and responses to changes in the environment may be quite small, so that the implicit assumption that unobservable terms in the first-order conditions will remain the same when the environment changes may be more realistic.

[8] In dynamic games, Sweeting (2013) and Jeziorski (2013) also separate estimation into stages, by making timing assumptions about when innovations in product qualities occur. The issue of selection has been addressed head-on in the empirical analysis of auctions by Bhattacharya, Roberts, and Sweeting (2014) and Roberts and Sweeting (2013), using incomplete information games where potential bidders may have noisy information about their true values when deciding to enter the auction.

[9] Reiss and Spiller noted that entry models "must recognize that entry introduces a selection bias in equations explaining fares or quantities." (p. S201).

with selection and they also consider applications to mergers. There are, however, significant differences between the papers that are informative about the trade-offs involved. CMT use a Nested Fixed Point estimation procedure where they repeatedly solve for all (pure strategy) Nash equilibria in many simulated games, which they use to construct an objective function based on inequalities. The resulting objective function is discontinuous and the computational burden is addressed by limiting markets to six players and by using a simulated annealing minimization algorithm on a supercomputer. The computational burden of our approach is much lower and it should therefore be accessible to more researchers. That said, our method will have lower econometric efficiency for a similar number of simulations. Substantively, CMT, following Ciliberto and Tamer (2009) and Berry (1992), focus on the decision of carriers to enter a market, without making a distinction between nonstop and connecting service. We focus on the decision to provide nonstop service, as competition been nonstop carriers has been central to the antitrust analysis of airline mergers and the data suggests that the fixed costs of providing connecting service, when a carrier already serves both airport endpoints, may be small, whereas the fixed costs of providing nonstop service, which requires a commitment of aircraft and gates, may be much more substantial.[10]

2 Data and Empirical Setting

In this section we highlight some of the most relevant features of our sample and describe how we define players, service types and several key variables. Full details are in Appendix A.

We estimate our model using a cross-section of publicly available data, taken from the Department of Transportation's DB1 10% ticket sample and its T100 Origin and Destination database, which provides data on flights between pairs of airports, from the second quarter of 2006.[11] We choose relatively old data for two reasons. First, our model is best viewed as a representation of an industry that is roughly in steady-state with firms behaving non-cooperatively. Subsequent years were associated with the after-effects of the financial crisis, several large mergers, and allegations of cooperative pricing behavior among major carriers.[12] Second, we want to see whether

[10]Dunn (2008) and Berry and Jia (2010) show that nonstop service is perceived to be a significantly higher quality product by at least some consumers.

[11]These data are widely used, but lack some information, such as details of when tickets are purchased, which would be required to build a model that included important industry practices such as revenue management.

[12]Of course, the industry was experiencing some changes in Q2 2006, following the 2005 US Airways/America

Figure 1: Empirical Cumulative Distribution Functions for the Number of Passengers Recorded in DB1 for Two Types of Carrier in the Sample Markets

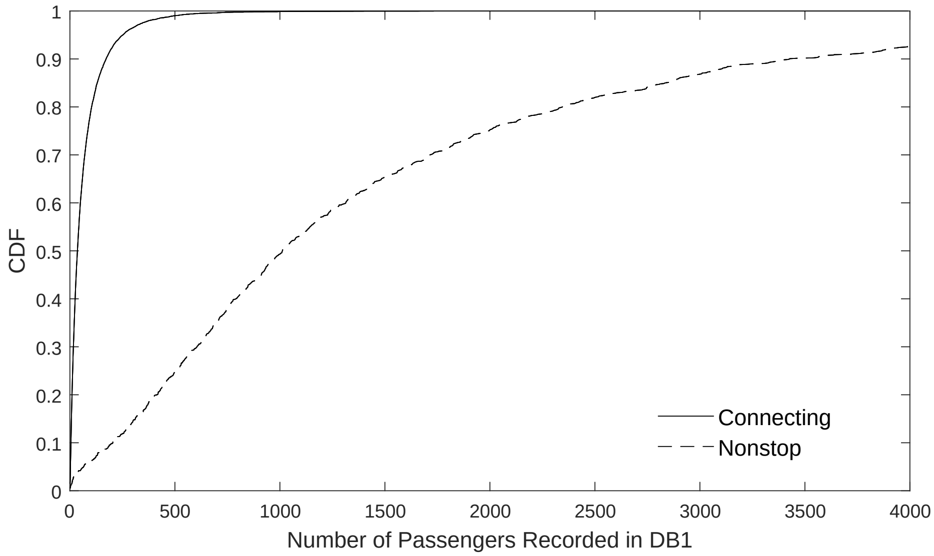

our model can predict observed changes in service types after subsequent mergers.

Market Selection, Carriers, Service Types, Market Shares and Prices. We use a sample of 2,028 airport-pair markets taken from the set of routes linking the 79 busiest US airports in the lower 48 states. Appendix A explains the selection criteria. After deleting itineraries with unusual prices, we aggregate itineraries to the level of the ticketing carrier. In this paper we will focus on seven named carriers, American, Continental, Delta, Northwest, Southwest, United and US Airways, and two composite carriers, which aggregate the other carriers that we observe in the data: "Other Legacy" (primarily Alaska) and "Other Low-Cost Carrier (LCC)". Our classification of carrier types follows Berry and Jia (2010).

A feature of the DB1 data is that, in many markets, a number of carriers are reported as carrying a very small number of passengers via connections. Figure 1 shows, for the markets and named carriers in our sample, the empirical cumulative distribution functions for the number of passengers recorded in DB1 for carriers who have no scheduled flights on a route ("connecting", 9,246 observations) and carriers who have at least ten scheduled nonstop flights (from T100) during the quarter (which, for the purpose of constructing this figure only, we call "nonstop", 1,256).[13] The median number of recorded passengers for passengers for the first (connecting)

West merger and the Q1 2006 closure of Independence Air.

[13]Throughout the paper, a return passenger counts as one, and a one-way passenger as one-half. In constructing this figure we sum up across the number of passengers and flights originating from both endpoints, and include flights by regional affiliates. 53 carrier-market observations with between one and ten nonstop flights are not

type is only 34, compared with 1,013 for the second (nonstop) type. The low connecting median suggests that the fixed costs of providing connecting service must typically be small, which motivates our focus on the choice of nonstop service, and it also suggests that many of the connecting carriers listed in DB1 may provide only weak competitive constraints on the pricing of nonstop flights. As the computational burden increases in the number of players, we use thresholds to define players and service types.

We define the actual players in a market as those carriers who achieve at least a 1% share of travelers, regardless of originating endpoint, and, based on the assumption that DB1 is a 10% sample, have no less than 200 return passengers per quarter.[14] We define a carrier as providing nonstop service on a route if, in T100, it is recorded as having at least 64 nonstop flights in each direction and at least 50% of the DB1 passengers that it carries are recorded as not making connections. The remaining players are defined as providing connecting service. Our service classification is not sensitive to the 64 flight and 50% nonstop thresholds as almost all nonstop carriers exceed these thresholds. For example, less than 10% of DB1 passengers make connections for more than 80% of our nonstop carriers. For this reason, we also feel comfortable ignoring the fact that carriers may provide both nonstop and connecting products in the same market. However, consistent with Figure 1, the 1% share/200 passenger thresholds do affect the number of connecting carriers.

We model demand and pricing in each direction on each route.[15] We use the average price in DB1 to measure a carrier's price. A carrier's market share in a particular direction is defined by the total number of passengers that it carries, regardless of service type, divided by a measure

shown.

[14]This approach assumes that it is relevant to focus on the carriers that were already serving a market when trying to predict competition after a merger using a counterfactual. This can be rationalized by the fact that the set of competing carriers is fairly stable, at least in the short-run. For example, of the 1,172 carriers that we define as providing nonstop service in Q2 2006, 1,027 of them were providing nonstop service on the same route in Q2 2005 and only 26 of them were present at the endpoints but not serving the market at all (given our definitions).

[15]Carriers may choose a similar set of ticket prices that they can use in each direction but revenue management techniques mean that average prices can be systematically different. Differences in market shares across directions can depend on carrier endpoint presence, because frequent-flyer programs or marketing may mean that departing passengers prefer to travel on a carrier that has greater local presence even if prices and frequencies are similar. A reduced-form analysis indicates that these effects can be large. For example, a route fixed effects regression where the difference in market shares across directions is regressed on the difference in presence indicates that a one standard deviation increase in the difference in presence increases the expected difference in market shares by 1.3 percentage points, which is large given that average market shares are 7.1%. The difference in presence also has statistically significant effects on differences in average prices across directions, although the percentage magnitudes are much smaller.

of market size. Appendix A describes how we define market size using the predicted values from a gravity model. We prefer this approach to using the geometric average of endpoint city populations, the most common approach in the literature, because that approach produces implausible variation in market shares across routes and across directions on the same route.

Explanatory Variables. We construct a number of variables that we include in the demand and/or cost equations of our model. The composite "Other Legacy" and all of the named carriers except Southwest are defined as legacy carriers. Carrier presence at an airport is defined by the number of domestic routes that the carrier, or its regional affiliates, serve nonstop from the airport divided by the total number of different routes served nonstop by all carriers out of the airport. We distinguish between presence at the origin and the destination of a directional route. Nonstop distance is defined as the great circle distance, in miles, of a return trip. We define Reagan National, LaGuardia and JFK as slot-constrained airports. We allow for the price sensitivity of demand to vary with a measure of the proportion of business travelers on the route based on data provided to us by Severin Borenstein (Borenstein (2010)).[16] For named carriers, we allow the marginal costs of connecting service to depend on the distance flown via the carrier's domestic hub that involves the shortest total journey distance.

The legacy carriers in our data operate hub-and-spoke networks, and nonstop service is likely profitable on many medium-sized routes out of hubs only because of the amount of traffic that nonstop service generates for other routes on the network. While our structural model only captures price competition for passengers traveling the route itself, we allow for connecting traffic to reduce the effective fixed cost of providing nonstop service by including three carrier-specific variables in our specification of fixed costs. Two variables are indicators for the principal domestic and international hubs of the non-composite carriers (these are listed in Appendix A).

We also include a continuous measure of the potential connecting traffic that will be served if nonstop service is provided on routes involving a domestic hub. The value of the variable is the prediction from a reduced-form regression model, estimated using Q2 2005 data, where we use a Heckman selection approach to correct for the fact that routes may have nonstop service only when the carrier can serve unusually large amounts of connecting traffic. The model, and the exclusion restrictions, are detailed in Appendix A.1. We acknowledge that this measure will not

[16] This is based on measures of business traveler usage at the airport-level. For this reason we treat it as exogenous to prices and service decisions at the route-level, while recognizing that it is likely to be an imperfect measure of how many business travelers want to travel on a particular route.

Figure 2: Number of Carriers Offering Nonstop Service Between Selected Airports

be completely consistent with the structural model that we are estimating, because it is based on a model where a hub carrier's service decisions do not depend on the outcome of a multi-carrier game. However we have found that including this variable can help to explain patterns of service in the data and we view it as an approximation of the type of non-game theoretic models that carriers may use to predict flows of connecting passengers.

Summary Statistics. Table 1 contains market-level and market-carrier-level summary statistics for the primary variables in our data. On average, there are four carriers in each market, with more carriers in long-distance markets where there tend to be more plausible connections. For example, Seattle to Baltimore and Seattle to Orlando have the maximum nine carriers (one nonstop carrier). 53% of routes have no nonstop service, but larger markets and routes connecting the hubs of multiple carriers have as many as four nonstop carriers. To illustrate how market structure varies, Figure 2 shows the number of nonstop carriers for the routes in our sample between ten airports with varying hub status serving metropolitan areas of different sizes. Most nonstop service involves a hub airport: for example, Salt Lake City, a Delta hub, has more nonstop service than non-hub airports in larger MSAs such as San Diego and San Antonio. Smaller, non-hub airports, such as Greensboro's Piedmont-Triad, only have nonstop service to nearby hubs.

Fares vary systematically with distance (an increase in the return distance of 1,000 miles

Table 1: Summary Statistics for the Estimation Sample

	Obs.	Mean	Std. Dev.	10th pctile	90th pctile
Market Variables					
Market Size (directional)	4,056	24,327.4	34,827.37	2,794	62,454
Num. of Carriers	2,028	3.98	1.74	2	6.2
Num. of Nonstop	2,028	0.668	0.827	0	2
Total Passengers (directional)	4,056	6970.90	10830.06	625	17,545
Nonstop Distance (miles, return)	2,028	2,444	1,234	986	4,384
Business Index	2,028	0.41	0.09	0.30	0.52
Market-Carrier Variables					
Nonstop	8,065	0.17	0.37	0	1
Price (directional, return $s)	16,130	436	111	304	581
Share (directional)	16,130	0.071	0.085	0.007	0.208
Airport Presence (endpoint-specific)	16,130	0.208	0.240	0.038	0.529
Low Cost Status	8,065	0.22	0.41	0	1
≥ 1 Endpoint is a Domestic Hub	8,065	0.13	0.33	0	1
≥ 1 Endpoint is an International Hub	8,065	0.10	0.30	0	1
Connecting Distance (miles, return)	7,270	3,161	1,370	1,486	4,996
Log(Predicted Connecting Traffic)	1,036	6.44	0.81	5.31	7.47

increases average fares by $30), whether service is nonstop (nonstop service fares are $43 higher than connecting fares), whether the carrier is low-cost (low-cost carrier fares are $70 lower than legacy fares) and the degree of competition, and especially the number of nonstop carriers. Controlling for route distance and the identity of the carrier, the first nonstop carrier is associated with connecting fares falling by $10, while a second nonstop carrier is associated with a $40 reduction in nonstop fares and a $30 reduction in connecting fares. This pattern motivates our focus on what determines the number of carriers providing nonstop service in equilibrium.[17]

Changes in Service Choices After Actual Mergers Given our focus on service choices, it is natural to ask what service changes are observed after actual mergers. To do this, we have examined 17 routes where in the quarter that a merger between legacy carriers (Delta/Northwest, United/Continental, American/US Airways) closed financially, the merging parties were both providing nonstop service and no other carriers were doing so, as these are the routes where both

[17]The distance, nonstop service and competition estimates come from regressions of a carrier's weighted (across directions) average fare on a route on nonstop distance, carrier dummies, a dummy for whether the carrier provides nonstop service and interactions between whether a carrier provides nonstop service and the number of nonstop carriers on a route. To estimate the effect of low-cost status we replace carrier dummies with a dummy for the low-cost status of the carrier.

intuition and our estimates suggest that there may be the largest anti-competitive effects. For this exercise we define nonstop service using only T100 and treat a route as being served by a carrier nonstop when the carrier itself or its regional affiliates fly at least 130 flights (in either direction) in each quarter.

For all of these routes, the merged firm continued providing nonstop service for the four years after the merger was financially completed. After one, two and three years the number of routes where at least one other carrier had initiated nonstop service were two, four and six respectively, out of 17 routes, so less than one-quarter of markets had experienced new entry within the two-year window that is often viewed as being relevant for evaluating supply-side substitution in merger analysis. We show below that we can only match this rate of nonstop initiation in our counterfactuals when we allow for selection on both observables and unobservables.[18]

3 Model

Consistent with the majority of the airline literature we focus on carriers' strategic decisions at the route-market level (see Mazur (2016) for an exception). Consider a particular market, m, connecting two airports A and B. Denote the players by $i = 1, ..., I_m$. The carriers play a two-stage, complete information game. In the first stage they decide whether to provide nonstop or connecting service (i.e., they make a binary choice as in most of the entry literature, but both alternatives involve some level of service). This choice is non-directional. Nonstop service implies paying a fixed cost, F_{im}, whereas we assume that there is no fixed cost associated with providing connecting service. Our model does not allow for the possibility that a carrier provides both nonstop and connecting service on the same route, motivated by the fact that when nonstop service is offered almost all passengers travel nonstop (see Section 2). As a baseline assumption, we assume that carriers decide what type of service to provide in a sequential order, with the carriers with the highest average presence moving first. In the second stage, they choose prices.

[18]Two of the routes where nonstop service was initiated involved an airport (Newark or Reagan National) where the merging parties had to divest slots as part of the merger approval process. For example, Southwest was able to enter Newark, and the Denver-Newark route that had been a United and Continental nonstop duopoly, through an approved purchase of slots from United/Continental. It is obviously possible that a carrier receiving slots would choose to serve some of the routes that the Department of Justice was most concerned about in order to encourage the Department to pursue this type of remedy in future airline mergers. We have also looked at what happened after the Southwest/Airtran merger. In this case there was an even lower rate of entry on nonstop duopoly routes.

3.1 Second Stage: Post-Entry Price Competition

We assume that, given service choices, carriers play two static Bertrand Nash pricing games for passengers originating at each endpoint. We model consumer demand from each endpoint separately and, in each case, demand is described by a nested logit model. For example, for customer k originating at endpoint A, the indirect utility for a return-trip on carrier i is

$$u_{kim}^{A\to B} = \beta_{im}^{A\to B} - \alpha_m p_{im}^{A\to B} + \nu_m + \tau_m \zeta_{km}^{A\to B} + (1-\tau_m)\varepsilon_{kim}^{A\to B} \tag{1}$$

where $p_{im}^{A\to B}$ is the directional price charged by carrier i, given the type of service that it offers. The first term represents carrier quality associated with the type of service that it offers,

$$\beta_{im}^{A\to B} = \beta_{im}^{CON,A\to B} + \beta_{im}^{NS} \text{ x } \mathcal{I}(i \text{ is nonstop})$$

where

$$\beta_{im}^{CON,A\to B} \sim N(X_{im}^{CON}\beta_{CON}, \sigma_{CON}^2)$$

and

$$\beta_{im}^{NS} \sim TRN(X_{im}^{NS}\beta_{NS}, \sigma_{NS}^2, 0, \infty)$$

so that quality can depend on observed characteristics, such as the type of carrier (legacy vs. LCC) and route characteristics, but it also depends on a random component that is unobserved to the researcher. TRN denotes a truncated normal distribution and the lower truncation of β_{im}^{NS} at zero implies that the perceived quality of nonstop service will always be greater than that of connecting service on the same carrier. To apply our estimation procedure we will impose some additional restrictions on supports, described below. We also allow the price coefficient and nesting parameters to be heterogenous across markets, with $\alpha_m \sim N(X^\alpha\beta_\alpha, \sigma_\alpha^2)$, where X^α will include the business index for the route, and $\tau_m \sim N(\beta_\tau, \sigma_\tau^2)$. We assume that α_m and τ_m are the same across directions for the same route.[19]

ν_m is a market-specific random effect that is designed to capture the fact that in some markets there are more travelers in both directions, relative to our chosen definition of market size, than can be rationalized with independent quality heterogeneity across carriers. We assume that ν_m

[19]This helps us to fit the pattern that the differences in carrier prices across directions are usually small.

is normally distributed with mean zero and standard deviation σ_{RE}. $\varepsilon_{kim}^{A\to B}$ is a standard logit error for consumer k and carrier i.

Each carrier has a marginal cost of carrying a passenger. Specifically we assume that

$$c_{im} \sim N(X_{im}^{MC}\beta_{MC}, \sigma_{MC}^2)$$

where the expected cost can depend on the type of carrier, the type of service and the distance traveled through the parameters β_{MC}. For nonstop service the distance is simply the nonstop distance between A and B. For a connecting carrier the distance is the distance from A to the carrier's nearest major domestic hub or focus city plus the distance from that same hub or focus city to B.[20] As we assume that travelers are making return trips we treat the marginal cost as non-directional.

This specification is restrictive in two ways. First, the random component of marginal costs does not vary with the service choice, which is different to what we assumed about quality. Second, our data gives us two directional average prices and two directional market shares for each carrier, while here we are allowing for two directional quality unobservables and a single marginal cost unobservable so we cannot rationalize every realization of market shares and prices in the data. We have adopted these restrictions based on the fact that we have found that models with independent cost shocks across either directions or service choices have fit the data less well (for example, implying more variable prices and market shares across directions than is actually observed).

Given Bertrand Nash equilibrium pricing choices (which will be unique given that we assume nested logit demand, linear marginal costs and single product firms), we can calculate variable profits in each direction, $\pi_m^{A\to B}(s)$, as a function of a vector of service types, s, and realized draws for cost and quality. We define market-level variable profits as $\pi_m(s) = \pi_m^{A\to B}(s) + \pi_m^{B\to A}(s)$, as service choices are assumed to be the same in both directions.

[20]For the composite Other Legacy and Other Low Cost carriers it is not straightforward to assign connecting routes. Therefore we use the nonstop distance for these carriers, but include additional dummies in the connecting marginal cost specification to provide more flexibility.

3.2 First Stage: Service Type Selection

In the first stage of the game carriers choose whether to commit to the fixed costs associated with nonstop service. If not, they provide connecting service. For our baseline estimation, we model carriers as making their service choice sequentially in an order that is known to both the firm and the researcher, so there is an extensive form game where the payoff of a carrier i is defined as

$$\pi_{im}(s) - F_{im} \text{ x } \mathcal{I}(i \text{ is nonstop in } m) \tag{2}$$

where F_{im} is a fixed cost draw associated with providing nonstop service. We assume that

$$F_{im} \sim TRN(X_{im}^F \beta_F, \sigma_F^2, 0, \infty).$$

where X_{im}^F includes several airport and carrier network characteristics. We assume that all of the market-level and carrier-level demand and cost draws are known, by all carriers, when service choices are made. We assume that the move order is determined by the average presence of the carriers across the market endpoints, with the highest average presence carrier moving first.[21] We also consider the robustness of our estimates when we allow for the equilibrium played to be any of the pure strategy Nash equilibria in the simultaneous service choice game[22] *or* a subgame perfect Nash equilibrium in a sequential move game with any order of moves.

3.3 Solving the Model

Conditional on s, we solve for equilibrium prices, market shares and profits by solving the system of pricing first-order conditions in the usual way. The natural way to solve for the subgame perfect Nash equilibrium in the sequential first stage of the model is by backwards induction. However, rather than solving for equilibrium profits at all branches of the game tree, we reduce the game tree by selectively growing it *forward*. To be precise, we first calculate whether it would be profitable for the first mover to operate as a nonstop carrier if it was the only carrier in the

[21] Berry (1992) has previously estimated a model of sequential entry for airline markets, assuming that profitability and incumbency affect the order.

[22] Given the assumed form of competition, there will be at least one pure strategy equilibrium in the simultaneous move game.

market, given its F.[23] If not, then we do not even need to consider any of the branches where it provides nonstop service, immediately eliminating half of the game tree from consideration. If it is profitable, then we need to keep both of the initial branches. We then turn to the second carrier, and ask the same question, for each of the remaining first carrier branches under consideration, and we only keep the nonstop branch for the second carrier if nonstop service yields positive profits. Once this has been done for all firms we can solve backwards to find the unique subgame perfect equilibrium using the resulting tree.

In our game the benefits from this selective growing of the game tree are useful but not necessary for our approach to be feasible. Indeed, we use a more standard approach when we calculate all of the pure strategy Nash equilibria in a simultaneous move game. However, if we were to allow for more choices or more carriers then this type of approach may be necessary for estimation to be feasible.

4 Estimation and Identification

Nested fixed point estimation procedures are computationally expensive because each time a parameter is changed the entry and pricing models need to be solved for every market. We view this approach as being infeasible for our model, where there are up to nine players, directional demand and directional pricing, without access to massive computational resources.

Instead, we use an estimation approach that has two steps. In the first step, we *solve* a large number of games where carrier qualities, marginal costs and fixed costs are drawn from *importance densities* chosen by us as researchers. In the second step we estimate the structural parameters (the βs and the σs from the model description in Section 3) using a method-of-moments estimator where we approximate the moments implied by the parameters by re-weighting the outcomes from the games solved in the first step. The key feature of the second step is that we only need to calculate a large number of probability density functions, not re-solve the economic model. The first step can be spread across a number of machines as each game is solved independently.[24]

[23] To be clear, here we are testing whether the profits from providing nonstop service are positive, which is a necessary condition for this service choice ever to be optimal, not whether it is more profitable than providing connecting service.

[24] An additional advantage is that alternative specifications that only involve changing the explanatory variables that affect the conditional means of different draws can be estimated without repeating the first step.

In this section we outline the estimation procedure and our selection of moments, and discuss the possible problems that are known to exist with this type of approach. Appendix B describes additional details and a Monte Carlo experiment that evaluates how well the procedure works both with a known sequential order of entry and a more agnostic equilibrium selection assumption.

4.1 Importance Sampling

Our method is based on Ackerberg (2009), who describes the potential advantage of importance sampling as a method for approximating an objective function when estimating a rich economic model. In our setting, suppose that we want to calculate the expected value, $E_m(y)$, of a particular outcome, y (e.g., whether American provides nonstop service), in market m. Denote a realization of the quality and cost draws for each carrier as θ_m, and the parameters that describe the distribution of these draws, which are the parameters that we want to estimate, as Γ. Denoting the density of the θ draws as $f(\theta_m|X_m,\Gamma)$,

$$E_m(y|\Gamma) = \int y(\theta_m, X_m) f(\theta_m|X_m,\Gamma) d\theta_m$$

where, because our baseline model generates a unique equilibrium, $y(\theta_m, X_m)$ is the unique outcome given θ_m and observed X_m. This integral cannot, in practice, be calculated analytically, but we can exploit the fact that

$$\int y(\theta_m, X_m) f(\theta_m|X_m,\Gamma) d\theta_m = \int y(\theta_m, X_m) \frac{f(\theta_m|X_m,\Gamma)}{g(\theta_m|X_m)} g(\theta_m|X_m) d\theta_m$$

where $g(\theta_m|X_m)$ is an importance density chosen by the researcher.

An important assumption is that $g(\theta_m|X_m)$ and $f(\theta_m|X_m,\Gamma)$ have the same support, and that this support does not depend on Γ. We specify the supports for all of the demand and cost draws prior to estimation, trying to include the full range of values that we believe to be plausible.[25] For a given set of S draws from g we can then approximate $E_m(y)$ using

$$\widehat{E_m(y|\Gamma)} \approx \frac{1}{S}\sum_{s=1}^{S} y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$$

where we calculate $y(\theta_{ms}, X_m)$ once for each draw before estimation, and then re-weight the outcomes from each of these draws using $\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$, which only requires calculating a pdf, during estimation of Γ.[26] A major benefit is that $\widehat{E_m(y|\Gamma)}$ will be a smooth function of Γ even when the outcome of interest, such as a service choice, is discrete.

4.2 Moments, Supports, Starting Values and Weighting Matrix

We minimize a standard simulated method of moments objective function in the second step

$$m(\Gamma)'Wm(\Gamma)$$

where W is a weighting matrix. $m(\Gamma)$ is a vector of moments where each element has the form $\frac{1}{2{,}028}\sum_{m=1}^{m=2{,}028}\left(y_m^{data} - \widehat{E_m(y|\Gamma)}\right) Z_m$, where subscript ms represent markets. We use a large number (1,384) of moments in estimation, based on a range of price, share and service-type outcomes, y_m, defined either at the carrier-level or the market-level, and observed variables that are treated as exogenous. Appendix B provides additional details.

To apply importance sampling we need to specify the support of each of the θ draws and to choose the importance density g. To generate the reported results, we use the supports and truncated densities listed in Table 2. The supports were chosen to be broad in the sense that they contain all of the values that were likely to be relevant, with the exception of the support for the nesting parameter which was restricted because we found, when using broader supports, some local minima with implausibly high or low values of τ. The assumed range of τ is consistent with most values in the literature (for example, Berry and Jia (2010) and Ciliberto and Williams (2014) report estimates between 0.62 and 0.77, albeit with a different definition of market size) and with values of τ that are estimated if demand is estimated separately (i.e., selection is not

[25] There is a trade-off here. When we use wider supports we will be taking more demand and cost draws that will likely be irrelevant given the estimated parameters. For a given number of draws, this reduces efficiency. However, choosing supports that are too small may limit our ability to match important patterns in the data.

[26] As g does not depend on Γ, g can be calculated once at the beginning of the estimation procedure.

Table 2: Description of g For the Final Round of Estimation

Market Draw	Symbol	Support	g
Market Random Effect	υ_m	[-2,2]	$N(0, 0.411^2)$
Market Nesting Parameter	τ_m	[0.5,0.9]	$N(0.634, 0.028^2)$
Market Demand Slope (price in \$00s)	α_m	[-0.75,-0.15]	$N(X_m^{\alpha}\beta_{\alpha}, 0.022^2)$
Carrier Draw			
Carrier Connecting Quality	$\beta_{im}^{CON,A\rightarrow B}$	[-2,10]	$N(X_{im}^{CON}\beta_{CON}, 0.219^2)$
Carrier Incremental Nonstop Quality	β_{im}^{NS}	[0,5]	$N(X_{im}^{NS}\beta_{NS}, 0.257^2)$
Carrier Marginal Cost (\$00s)	c_{im}	[0,6]	$N(X_{im}^{MC}\beta_{MC}, 0.173^2)$
Carrier Fixed Cost (\$m)	F_{im}	[0,5]	$N(X_{im}^{F}\beta_{F}, 0.234^2)$

Notes: where the covariates in the Xs are the same as those in the estimated model, and the values of the βs for the final (initial) round of draws are as follows: β_{α}.constant= -0.668 (-0.700), β_{α}.bizindex=0.493 (0.600), β_{α}.tourist= 0.097 (0.2), β_{CON}.legacy= 0.432 (0.400), β_{CON}.LCC= 0.296 (0.300), β_{CON}.presence= 0.570 (0.560), β_{NS}.constant= 0.374 (0.500), β_{MC}.legacy= 1.802 (1.600), β_{MC}.LCC= 1.408 (1.400), β_{MC}.nonstop_distance= 0.533 (0.600), β_{MC}.nonstop_distance^2 = -0.005 (-0.01), β_{MC}.conn_distance= 0.597 (0.700), β_{MC}.conn_distance^2 = -0.007 (-0.020), the remaining marginal cost interactions are set equal to zero, β_F.constant= 0.902 (0.750), β_F.dom_hub= 0.169 (-0.25), β_F.conn_traffic= -0.764 (-0.01), β_F.intl_hub= -0.297 (-0.55), β_F.slot_constr= 0.556 (0.700). In the initial round the standard deviations of the draws were as follows: random effect 0.5, nesting parameter 0.1, slope parameter 0.1, connecting quality 0.2, nonstop quality premium 0.5, marginal cost 0.15, fixed cost 0.25.

accounted for). Draws from the gs are taken independently for each market, carrier and type of draw.

To get to the parameters used to form the g densities, we initially attempted to match (by eye) a small number of price, market share and entry moments to make sure that our model could capture the main patterns in the data. This led to the "initial" parameterization reported in the notes to the table, where we tried to allow for sufficiently large standard deviations that, during estimation, there would be enough draws covering a wide range of qualities and costs that the mean coefficients could move significantly if this allowed the estimated model to achieve a better fit. We then ran a couple of rounds of our estimation routine to identify the parameters that we use to create the draws for the final round of estimation whose results we report. While the estimator can be consistent for any set of gs that give finite variances, Ackerberg (2009) recommends using a multi-round estimation procedure to improve efficiency.[27] We take 2,000

[27] A formal iterated procedure was used by Roberts and Sweeting (2013) in estimating a model of selective entry for auctions, where the standard errors were bootstrapped to account for this multi-stage estimation procedure. To implement this bootstrapping approach, to account for what happens in the early iterations, in the current setting would create a large computational burden, so we instead present our results as being conditional on the

sets of draws from the gs for each market. 1,000 sets are used in the estimation (i.e., $S = 1,000$), with the full sets of 2,000 being used as a pool of draws that we use when performing a non-parametric bootstrap to calculate standard errors.

We use a diagonal weighting matrix with equal weight on the price, share and service-type moments, and, within each of these groups, the weight on a particular moment is based on the reciprocal of the variance based on some initial estimates.[28] We choose not to use the inverse of the full covariance matrix of the moments because, with a large number of moments relative to the number of markets, we cannot claim that we can estimate the full variance-covariance matrix consistently, and, in practice, the coefficient estimates are less stable if an estimate of the full-covariance matrix is used.

4.3 Identification

As shown by CMT, the complete information assumption, and the selection that it implies, means that the demand and marginal cost equations cannot be consistently estimated without an explicit model of entry/service choices. To see why, consider the linear estimating equation for a logit-based demand model with aggregate data (Berry (1994)) in a setting where single product firms choose whether or not to enter the market. Selection implies that the unobserved product characteristic will be correlated with observed characteristics and will no longer have mean zero. The second property implies that the use of instruments will not be sufficient for consistent estimation. We are looking at a service choice within an exogenous set of active carriers, but a similar problem arises because the unobserved component of the incremental quality of nonstop service will not have mean zero for the carriers that choose to enter nonstop. This problem can only be addressed with a model of the non-linear form of selection that emerges from the first-stage game.

The intuition for identification is that we are imposing exclusion restrictions on the equations defining the mean values of demand, marginal costs and fixed costs. For example, carrier endpoint presence is assumed to only affect the preferences of consumers originating at that endpoint, with no direct effect on marginal costs or fixed costs (although fixed costs can be affected by some non-

final round g, while acknowledging that the choice of g was informed by our initial attempts at estimation. See Appendix B.3 for a discussion of how using different gs affect the estimates in a Monte Carlo.

[28]The sum of the values on the diagonal of the weighting matrix equals 1 for each of the three groups of moments.

directional measures of a carrier's network). Route distance, which can vary across routes and across carriers depending on the location of their domestic hubs, is allowed to affect marginal costs, but not demand (our gravity based definition of market size accounts for the effect of distance on demand prior to estimation) or fixed costs. Domestic and international hub status, slot constraints and our continuous measure of generated connecting traffic affect the fixed cost of nonstop service but not demand or marginal costs. Our measure of connecting traffic may be especially valuable because when it is very large, nonstop service may be close to certain so that the draw of incremental nonstop quality should be almost uncensored (not selected).

Of course, our parametric assumptions and the assumed order of entry will also contribute to identification. We assume that, for a given carrier, the quality, marginal cost and fixed cost residuals are uncorrelated. This assumption could be relaxed (and CMT do so), although our estimates suggest that observables account for most of the variation in marginal and fixed costs so that the gains from allowing this type of correlation may be limited, and we have found that the objective function is more likely to have multiple local minima when we allow for unrestricted correlations. Unlike CMT, we allow for correlation in demand across carriers, in the form of a market-level random effect.

5 Parameter Estimates

In this section we discuss the parameter estimates and assess the model's fit and the performance of our estimation method. We analyze the extent and effects of selection and counterfactuals in Sections 6 and 7.

5.1 Estimates with Known Order of Entry

The parameter estimates given our assumed order of entry are presented in column (1) of Table 3. The standard errors, in parentheses, are based on 100 bootstrap replications where 2,028 markets are sampled with replacement, and we draw a new set of 1,000 simulation draws (taken from our original 2,000 draw sets) for each drawn market.

Demand Parameters. The estimated standard deviation for the market random effect indicates that there is unobserved heterogeneity in the level of demand for air travel across markets, whereas there is very little unobserved heterogeneity in the nesting and demand slope

Table 3: Coefficient Estimates (bootstrapped standard errors in parentheses)

				(1) Assumed Order of Entry		(2) No Eqm. Selection
Demand: Market Parameters						
Random Effect	Std. Dev.	σ_{RE}	Constant	0.311	(0.138)	0.350
Nesting Parameter	Mean	β_{τ}	Constant	0.645	(0.012)	0.647
	Std. Dev.	σ_{τ}	Constant	0.042	(0.010)	0.040
Demand Slope	Mean	β_{α}	Constant	-0.567	(0.040)	-0.568
(price in $100 units)			Business Index	0.349	(0.110)	0.345
	Std. Dev.	σ_{α}	Constant	0.015	(0.010)	0.017
Demand: Carrier Qualities						
Carrier Quality for	Mean	β_{CON}	Legacy Constant	0.376	(0.054)	0.368
Connecting Service			LCC Constant	0.237	(0.094)	0.250
			Presence	0.845	(0.130)	0.824
	Std. Dev.	σ_{CON}	Constant	0.195	(0.025)	0.193
Incremental Quality	Mean	β_{NS}	Constant	0.258	(0.235)	0.366
of Nonstop Service			Distance	-0.025	(0.034)	-0.041
			Business Index	0.247	(0.494)	0.227
	Std. Dev.	σ_{NS}	Constant	0.278	(0.038)	0.261
Costs						
Carrier Marginal Cost	Mean	β_{MC}	Legacy Constant	1.802	(0.168)	1.792
(units are $100)			LCC Constant	1.383	(0.194)	1.331
			Conn. X Legacy	0.100	(0.229)	0.134
			Conn. X LCC	-0.165	(0.291)	-0.077
			Conn. X Other Leg.	-0.270	(0.680)	0.197
			Conn. X Other LCC	0.124	(0.156)	0.164
			Nonstop Dist.	0.579	(0.117)	0.589
			Nonstop Dist.^{2}	-0.010	(0.018)	-0.012
			Connecting Distance	0.681	(0.083)	0.654
			Connecting Distance^2	-0.028	(0.012)	-0.024
	Std. Dev.	σ_{MC}	Constant	0.164	(0.021)	0.159
Carrier Fixed Cost	Mean	β_{F}	Legacy Constant	0.887	(0.061)	0.913
(units are $1 million)			LCC Constant	0.957	(0.109)	1.015
			Dom. Hub Dummy	-0.058	(0.127)	-0.140
			$\widehat{\text{Connecting Traffic}}$	-0.871	(0.227)	-0.713
			International Hub	-0.118	(0.120)	-0.168
			Slot Const. Airport	0.568	(0.094)	0.602
	Std. Dev.	σ_{F}	Constant	0.215	(0.035)	0.198
			Run Time	29 CPU-hours		47 CPU-hours

parameters. All else equal, demand on business routes is less elastic, consistent with estimates from richer demand models that allow for multiple types of customers (Berry and Jia (2010) and Ciliberto and Williams (2014)). The expected price parameter for the market with the highest business index (Dayton to Dallas-Fort Worth) is -0.34 compared to the cross-market average of -0.57. The point estimates imply an average (absolute value) own-price demand elasticity of 4.25. The nesting parameter implies that if a carrier's price rises, most substitution is to other carriers rather than the outside good. The average elasticity of demand for air travel (i.e., the change in the total number of travelers when all prices are increased) is 1.29.[29]

The remaining demand parameters indicate that customers prefer carriers with a higher presence at their originating airport, which is also consistent with the earlier literature. The point estimates imply that preference for nonstop service is stronger on shorter routes and routes with a higher business index, although these coefficients are not statistically significant. Legacy carriers are estimated to give higher utility, all else equal, than low-cost carriers.

Marginal Cost Parameters. We allow a fairly rich specification for observable marginal costs, in order to try to capture some of the differences in prices across routes and carriers. The coefficients indicate that legacy carriers have higher marginal costs for both nonstop and connecting service, and that distance increases nonstop and connecting costs in a similar way. For a legacy carrier, the average marginal cost of providing nonstop service on a roughly 3,000 mile round-trip route, Miami to Minneapolis, is \$345, compared to \$367 for connecting service. Marginal costs for Southwest are lower and, for this route, its nonstop and connecting (via Chicago Midway) costs are almost identical (\$303 and \$298 respectively). Estimated unobserved heterogeneity in marginal costs is quite small (estimated standard deviation is \$16).

Fixed Cost Parameters. The expected fixed cost for nonstop service is around \$841,000, although the expected value for the carriers that choose nonstop service is around \$610,000. We estimate higher fixed costs for routes out of slot-constrained airports, reflecting the opportunity cost of using a slot for a specific route. The remaining parameters have the expected signs and allow us to explain why carriers serve many routes nonstop from their domestic and international hubs, and especially those that will generate significant connecting traffic for other destinations.[30]

[29]This estimate is consistent with the existing literature: for example, Gillen, Morrison, and Stewart (2003) report a median elasticity of 1.33 across 85 airline demand studies, and Berry and Jia (2010) estimate an elasticity of 1.67 using a much more disaggregated demand model and data from 2006.

[30]The connecting traffic prediction variable is scaled, and for routes out of domestic hubs its mean is 0.52 and its standard deviation is 0.34. This implies that the mean of the untruncated fixed cost distribution can be

Estimated unobserved heterogeneity in fixed costs is relatively small, reflecting the fact that observables are able to explain which carriers provide nonstop service.

5.2 Estimates Without a Known Order of Entry

As explained above, the assumption that there is a known (to the researcher), sequential order of entry is helpful in allowing the model to be solved quickly and it also implies that the model will generate a unique predicted outcome. However, the assumption is stronger than is necessary, and earlier papers in this literature have found that imposed equilibrium selection assumptions can be restrictive.[31] Column (2) of Table 3 reports the point estimates when we minimize an objective function based on moment inequalities, allowing for the observed outcome to be the outcome associated with any pure strategy Nash equilibrium in the simultaneous move game or the subgame perfect Nash equilibrium in a sequential game with any order of moves.

This method, described in more detail in Appendix B.2, is implemented by collecting together all of the possible equilibrium outcomes of the game, for a given set of draws θ_{ms}, and calculating the maximum and minimum values of each predicted outcome. Importance sampling can be used, as before, to calculate the expected values of the maximum and minimum outcomes for a given set of parameters Γ, and we can form moments under the assumption that, on average, observed outcomes should be less than the expected maximum and above the expected minimum.

Our analysis indicates that the objective function is minimized for a unique set of parameters. The value of the objective function for this estimator is 77.59. We can also evaluate the objective function of this inequality estimator at the parameters reported in column (1): in this case the value is 85.70. Obviously a natural step would be to evaluate whether the coefficients in column (1) are within the confidence sets that can be constructed for the inequality estimator.[32] This is not a straightforward task when the number of moments is large, which is inconsistent with the assumptions in most of the literature. We note, however, that when we apply the tests proposed

negative for some of the routes with the most connecting traffic, but because the distribution is truncated at zero, realized fixed costs will still be positive.

[31] For example, the estimated reduced-form profit function in Berry (1992) is sensitive to the assumed order of entry and Ciliberto and Tamer (2009) calculate that their estimates imply multiple equilibria in over 90% of market simulations.

[32] While the baseline move order is not inconsistent with the more general assumptions, it is not necessarily the case that the restricted parameters should lie within the confidence set of the more general estimates. In particular, when multiplicity is quite rare the parameters may be point-identified from outcomes that are predicted uniquely, so the assumed baseline order is an over-identifying restriction that can be rejected.

Table 4: Number of Outcomes Supported as Pure Strategy, Simultaneous Move Nash Equilibria or Subgame Perfect Nash Equilibria in a Sequential Move Game Given the Estimated Parameters

Number of Carriers	1	2	3	4	5	6	7	8	9	All
Number of Markets	141	304	416	413	342	228	136	46	2	2,028
Average Numb. of Eqm. Outcomes Per Simulation Draw	1	1.004	1.014	1.019	1.025	1.022	1.028	1.037	1.042	1.017

by Chernozhukov, Chetverikov, and Kato (2016) (CCK) to test whether moment inequalities assumed for our estimator are violated we actually have a lower test statistic for the parameters in column (1) than for those in column (2).[33]

More importantly, the coefficients in column (2) are very similar to those in column (1), with the exception of some of the interaction coefficients in the marginal cost function. These coefficients also have large standard errors in column (1). The estimates are very similar partly because multiple equilibrium outcomes are rare. Table 4 shows the average number of distinct equilibrium outcomes when we solve for all pure strategy Nash equilibria and all subgame perfect Nash equilibria in any sequential move game using 2,000 draws for each market based on the estimated column (1) parameters. The average number of equilibrium outcomes is only 1.017, and 98.4% of draws support only a single equilibrium outcome.

5.3 Performance of the Estimation Algorithm

We have claimed that our estimation algorithm has desirable practical properties. We now briefly discuss this claim in the context of our estimates. As reported in Table 3, the reported

[33]For our estimators, the value of the CCK test statistic is 10.75 for the column (1) estimates and 12.58 for the column (2) estimates, when their critical values for a 5% significance test lie between 4.1 and 4.3 depending on the method used to construct them. The CCK approach is potentially very useful in our application because it is valid when the number of moments is large and the calculation of the critical values does not require the minimization of an objective function. However, it cannot be adjusted to account for the fact that some inequalities are violated quite significantly when the parameters are estimated, as is the case in Ciliberto and Tamer (2009), Ciliberto, Murry, and Tamer (2016) and the current paper. In Ciliberto and Tamer (2009) and Ciliberto, Murry, and Tamer (2016) this problem is dealt with by deducting the value of the minimized objective function from the value calculated at other parameters. However, in CCK the test statistic is not calculated using the objective function so this fix cannot be implemented.

time taken to estimate the parameters in the second step of the routine is under 30 hours.[34] Optimization is performed on a single processor without using parallelization or even analytic derivatives, although for specifications where we have provided derivatives the estimation time was significantly lower. The resources required to estimate several specifications of the model should therefore be available to a wide range of researchers and practitioners.

In Appendix B.4 we plot the objective function when we vary the parameters one-at-a-time around their estimated column (1) values. In almost all cases the objective function has a simple convex shape. While these plots do not imply that the objective function is convex in multiple dimensions, they provide some optimism that we have found a global, as well as a local, minimum. We also use a graphical test of the assumption that the variance of the moments is finite following Koopman, Shephard, and Creal (2009), by examining how the volatility of the sample variance changes as the number of simulation draws is increased. We observe much lower volatility when the number of simulations, S, is above 500. In our application we use $S = 1,000$.

5.4 Model Fit

We now discuss the fit of the model. To do so, we simulate 20 new sets of demand and cost draws from the distributions implied by the parameter estimates and solve for equilibrium outcomes. The standard errors in parentheses are the standard deviation in the reported means when we perform further simulations based on the estimates from our bootstrap samples. Table 5 compares average prices and market shares in the data with those predicted by the model for different types of service and different market sizes.

We match average *differences* in market shares and prices across service types very accurately, although we overpredict the levels of prices and market shares.[35] For our counterfactuals it is particularly important to be able to predict service choices. Using our 20 simulated outcomes per market, our success rate at predicting a carrier's service is 87.5% (standard error 1.1%). This involves correctly predicting 91.7% (1.0%) of decisions to provide connecting service and

[34] On a medium-sized cluster, the first step can be performed in a couple of days without requiring any parallelization for a given market.

[35] The difference in the level of predicted average prices partly reflects the fact that here we are analyzing fit using a new set of simulation draws, not the importance sample draws that we used to predict prices during estimation, where, on average, we match average prices almost perfectly. These cross-carrier averages mask some differences at the carrier-level. For example, the observed and (predicted) prices for United's direct and connecting services are \$479 (\$472) and \$436 (\$445) so the match is very close, whereas for Delta the comparisons are \$498 (\$453) and \$448 (\$466).

Table 5: Model Fit: Average Market Shares and Prices (bootstrapped standard errors in parentheses)

			Data	Model Prediction
Average	All Markets	Any Service	$436	$455 (5)
Prices		Nonstop	$415	$436 (8)
(directions weighted		Connecting	$440	$458 (5)
by market shares)				
	Market Size Groups			
	1st Tercile	Any Service	$460	$465 (5)
	2nd Tercile	Any Service	$442	$460 (5)
	3rd Tercile	Any Service	$412	$441 (5)
Average	All Markets	Any Service	7.1%	8.4% (0.3%)
Carrier Market		Nonstop	17.9%	20.5% (0.9%)
Share		Connecting	4.9%	5.8% (0.3%)
	Market Size Groups			
	1st Tercile	Nonstop	25.6%	29.8% (2.4%)
		Connecting	8.6%	8.0% (0.4%)
	2nd Tercile	Nonstop	23.1%	26.6% (1.5%)
		Connecting	4.3%	5.5% (0.3%)
	3rd Tercile	Nonstop	15.9%	18.7% (0.8%)
		Connecting	1.8%	3.4% (0.3%)

Table 6: Model Fit: Prediction of Service Choices by Carriers at a Selection of Domestic Hubs

		Number of	% Nonstop	
Airport	Carrier	Routes	Data	Simulation
Atlanta	Delta	57	96.5%	92.5% (2.3%)
Salt Lake City	Delta	65	73.8%	52.9% (4.3%)
Chicago O'Hare	American	53	96.2%	90.2% (2.7%)
Chicago O'Hare	United	57	94.7%	92.4% (2.7%)
Charlotte	US Airways	46	84.7%	77.9% (2.7%)
Denver	United	58	72.4%	73.4% (4.2%)
Newark	Continental	43	86.0%	61.6% (5.0%)
Houston Intercontinental	Continental	55	90.9%	85.4% (4.3%)
Minneapolis	Northwest	62	85.4%	77.7% (6.3%)
Chicago Midway	Southwest	44	72.7%	64.5% (6.0%)

Table 7: Model Fit: Predictions of Service Decisions at Raleigh-Durham

	Number of Routes	Mean Presence Endpoints	% Nonstop Data	% Nonstop Simulation
American	44	0.29	22.7%	22.8% (1.6%)
Continental	30	0.14	10.0%	10.0% (1.0%)
Delta	57	0.24	8.7%	14.8% (1.9%)
Northwest	22	0.18	9.1%	11.0% (1.2%)
United	25	0.12	4%	14.4% (1.9%)
US Airways	54	0.12	5.6%	9.4% (2.7%)
Southwest	48	0.30	12.5%	14.5% (4.3%)
Other Low Cost	25	0.08	4%	13.4% (4.9%)

67.1% (2.8%) of decisions to provide nonstop service. However, if, for a given carrier-route, 11 or more of our 20 simulations predict nonstop service, this is what we observe in the data for 82.6% (2.2%) of market-carrier observations. Table 6 shows the performance of our model at predicting service at a number of hubs for the hub carrier. While we predict less nonstop service by Delta in Salt Lake City and Continental at Newark than they actually provide, the fit is generally impressive. We do even better at many non-hub airports. Table 7 shows the percentage of routes out of Raleigh-Durham served nonstop by each carrier (the number of routes varies across carriers depending on the airports that they serve, including via connections). Both the percentage (reported in the table) and the identity of routes served nonstop is predicted very accurately for the largest nonstop carriers, American and Southwest. The largest difference between the prediction and the data is for United, as most simulations predict that United would serve its hubs in Denver and San Francisco nonstop. These routes were added by United after 2006. Delta, whose service is also overpredicted, has also subsequently increased its nonstop service at RDU.

6 The Extent and Effects of Selection

We use our estimated model in two different ways. In Section 7 we present merger counterfactuals. In this section we quantify the extent of selection implied by our model and examine how selection affects market structure and consumer surplus.

6.1 The Extent of Selection Implied by the Estimates

In our model observed and unobserved variation in market demand, carrier quality, carrier marginal costs and carrier fixed costs all affect whether a carrier will provide nonstop service, whereas in the previous literature unobserved variation in market demand, quality and marginal costs was only revealed to firms once entry decisions or service choices had been made. To quantify how different variables affect service choices we estimate a set of linear probability models using the 20 sets of draws for each market that we used to characterize the fit of the model.[36] The dependent variable is a dummy for whether a carrier provides nonstop service and the observed and unobserved components of demand and cost are regressors. We rescale the continuous explanatory variables to have mean zero and standard deviation one, so that it is easier to compare the coefficients when variables have different units.

Table 8 shows seven specifications. Column (1) has only market-level regressors. Higher and less elastic demand make nonstop service more likely, and a one standard deviation change in (observed) market size has a much larger effect on service choices than variation in the random effect, nesting or price parameters. In column (2) we include the components of the carrier's own qualities and costs that are based on observed variables. These five variables increase the adjusted R^2 from 0.23 to 0.43, and they indicate that higher quality, lower nonstop marginal costs and, especially, lower fixed costs raise the probability of nonstop service.[37] Column (3) adds the unobserved components of carriers' quality and cost draws. This raises the adjusted R^2 by a further 0.1 (23%). Column (4) adds dummies for the carrier's position in the move order and the number of carriers that are players in the game. Consistent with almost all simulations having a single equilibrium outcome regardless of the move order, the adjusted R^2 and many coefficients change only slightly. The remaining specifications include either market, market-simulation or market-carrier fixed effects so that the coefficients are identified from cross-carrier or cross-simulation variation within markets and coefficients on variables that only vary across markets are not identified.

There are two important patterns in the coefficients. First, observable variation in demand

[36] Our model implies non-linear relationship between characteristics and service choices, but we view the estimated coefficients in this linear specification as being informative about the relative importance of different characteristics. The results are very similar using probit or logit models.

[37] Recall that an increase in connecting quality also increases nonstop quality, so that coefficients on nonstop quality measure the effects of incremental nonstop quality. The coefficient on observed nonstop quality has an unexpected negative sign and this may reflect a non-linear effect.

Table 8: Determinants of Nonstop Service at the Estimated Parameters

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Market Size	0.183***	0.152***	0.154***	0.160***			
	(0.006)	(0.005)	(0.005)	(0.005)			
Market Random	0.021***	0.021***	0.021***	0.021***	0.022***		0.022***
Effect	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)		(0.001)
Nesting	-0.003***	-0.002***	-0.003***	-0.003***	-0.003***		-0.003***
Parameter	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)		(0.001)
Price	0.033***	0.023**	0.025***	0.042***	-645.709		-639.643
Parameter	(0.003)	(0.009)	(0.009)	(0.009)	(1,271.633)		(1,293.485)
Components of Carrier Quality							
Obs. Connecting		0.080***	0.078***	0.091***	0.019***	0.019***	
		(0.020)	(0.019)	(0.019)	(0.003)	(0.004)	
Unobs. Connecting			-0.049***	-0.049***	-0.049***	-0.050***	-0.049***
			(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Obs. Nonstop		-0.021	-0.018	-0.052***			
		(0.019)	(0.018)	(0.018)			
Unobs. Nonstop			0.094***	0.094***	0.094***	0.105***	0.094***
			(0.002)	(0.002)	(0.002)	(0.003)	(0.002)
Components of Carrier Marginal Cost							
Obs. Connecting		0.011	0.010	0.149***	0.135***	0.133***	
		(0.030)	(0.029)	(0.031)	(0.029)	(0.033)	
Obs. Nonstop		-0.086**	-0.089**	-0.238***	-0.220***	-0.218***	
		(0.042)	(0.041)	(0.042)	(0.040)	(0.045)	
Unobs. MC			-0.027***	-0.027***	-0.027***	-0.030***	-0.027***
			(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Components of Carrier Fixed Cost							
Observed		-0.124***	-0.134***	-0.126***	-0.158***	-0.158***	
		(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	
Unobserved			-0.048***	-0.048***	-0.049***	-0.051***	-0.044***
			(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Carrier Position in Move Order							
2nd				-0.054***	-0.048***	-0.047***	
				(0.005)	(0.005)	(0.005)	
3rd				-0.091***	-0.085***	-0.084***	
				(0.006)	(0.006)	(0.007)	
4th				-0.105***	-0.102***	-0.101***	
				(0.007)	(0.006)	(0.007)	
5th				-0.107***	-0.106***	-0.106***	
				(0.007)	(0.007)	(0.008)	
6th				-0.102***	-0.102***	-0.102***	
				(0.008)	(0.007)	(0.008)	
7th				-0.105***	-0.108***	-0.107***	
				(0.009)	(0.008)	(0.009)	
8th				-0.110***	-0.115***	-0.114***	
				(0.014)	(0.013)	(0.015)	
9th				-0.096***	-0.102***	-0.099***	
				(0.008)	(0.007)	(0.009)	
Fixed Effects	-	-	-	Number of Carriers	Market	Market-Simulation	Market-Carrier
Observations	161,300	161,300	161,300	161,300	161,300	161,300	161,300
Adjusted R^2	0.230	0.427	0.521	0.528	0.588	0.550	0.632

and costs explains much of the variation in service choices. This highlights one of the benefits of our estimation method which allows many covariates to be included. It also helps to explain one pattern in the merger counterfactuals that we examine in Section 7, where we find that controlling for selection on observables leads to predicted post-merger price increases that are much closer to those made using a model that allows for selection on both observables and unobservables than a model that ignores selection entirely. Second, among the unobservables that we allow for, a robust pattern is that carrier-level variation in the incremental quality of nonstop service has a larger effect on service choices than market-level variation in demand or carrier-level variation in marginal costs or fixed costs.

6.2 Information and Market Structure

We now examine how our assumption of full information, which implies selection, affects the distribution of equilibrium outcomes. To do so, we focus on a single market, Austin-Los Angeles (LAX). In our data, this market has six carriers (Southwest, American, United, Other Legacy, US and Continental with this order of moves), with American and Southwest providing nonstop service. We choose this market because no carrier has overwhelming presence at either endpoint (the highest values are 39% for United at LAX and 37% for Southwest at Austin) so that we may plausibly see a variety of service choice outcomes for different draws. The market sizes in both directions are similar (63,231 and 69,891), and in our experiments we set them equal to the average and vary them in steps of 5,000 from 16,561 and 306,561 (which is below the market size of the two largest markets in our data).[38]

To analyze the effect of full information, we solve for equilibrium outcomes in two different sequential service choice models using the estimated parameters and simulate outcomes for 50,000 sets of the quality and cost draws for each market size in each case. The first model, "full information", is the model that we have estimated. The alternative "limited information" model assumes that carriers do not know the value of any unobserved marginal cost or quality shocks when they take service decisions but that they do know the fixed costs of nonstop service for all carriers.[39] This model is similar to the ones used in most of the literature that has combined

[38]We also fix the small standard deviation price and nesting coefficients equal to their mean values for this market and the demand random effect equal to zero. The price coefficient for this business-oriented market is -0.43.

[39]For each market configuration we approximate the expected profits of each carrier in every possible market

models of entry and competition. While information can change outcomes in many different ways, we focus on the relationship between market size, the number and identity of nonstop carriers and consumer surplus. Figure 3 shows the expected number of nonstop carriers and expected consumer surplus as a function of market size and Figure 4 shows the proportion of simulations resulting in different equilibrium numbers of nonstop carriers for different market sizes.

Under limited information there are fewer nonstop carriers in markets smaller than 46,561 and more nonstop carriers in larger markets. The intuition for this result is that, in small markets, nonstop service may only be profitable if marginal costs are unusually low or quality is unusually high. Knowledge of unobserved draws can therefore make nonstop service more likely. On the other hand, in large markets, nonstop service may only be unprofitable for high presence carriers if their quality or marginal cost draws are very unfavorable, so knowledge can make nonstop service less likely. This logic also affects which carriers provide nonstop service. For example, when market size equals 106,561, the modal number of nonstop carriers is 3 for both models, but the probability that Southwest, the highest average presence/first mover carrier, provides nonstop service is 0.85 under limited information and only 0.51 under full information. While larger markets have more nonstop carriers on average, consumer surplus tends to be higher in the full information model until market size exceeds 126,561, as knowledge of quality and marginal cost draws leads to the highest quality and lowest marginal cost carriers being selected into the product type that consumers prefer.

The distributions of the number of nonstop carriers in Figure 4 are also relevant for understanding the effects of mergers. Under limited information, the number of nonstop carriers is closely tied to the size of the market, so that if one nonstop carrier is eliminated and carriers re-optimize their service choices then it will be likely that a carrier will initiate nonstop service. On the other hand, with full information, the number of nonstop carriers has greater variance as it depends on the particular quality and marginal cost draws that carriers receive. Intuitively, knowledge of these draws can make it more likely that connecting carriers will find it unprofitable to initiate nonstop service if it was not profitable before the merger. Service upgrading

configuration by taking 1,000 draws of marginal costs and qualities. We then solve the sequential, limited information service choice game for each of the 50,000 draws of fixed costs before simulating realizations of the marginal cost and quality draws to compute expected consumer surplus. Note that the limited information model still has a complete information service choice game in the sense that all carriers have the same information.

Figure 3: The Relationship Between Market Size, Expected Consumer Surplus and the Expected Number of Nonstop Carriers Under Different Informational Assumptions

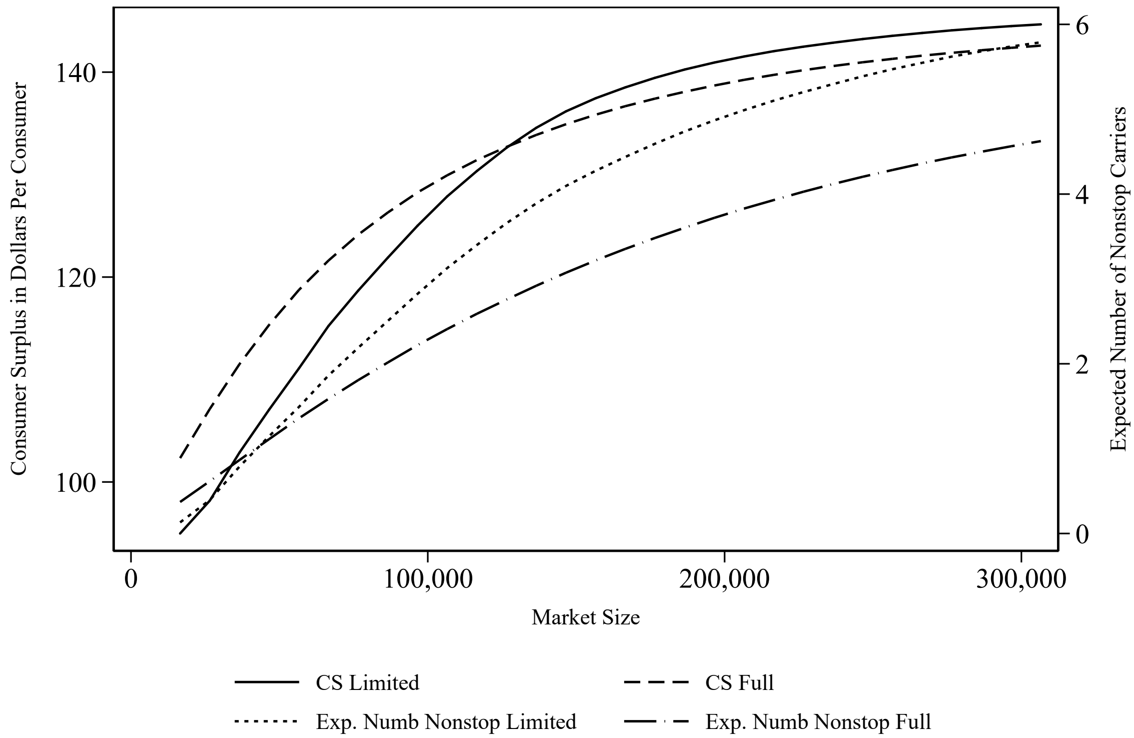

Figure 4: The Relationship Between Market Size and Equilibrium Market Structure Under Different Informational Assumptions

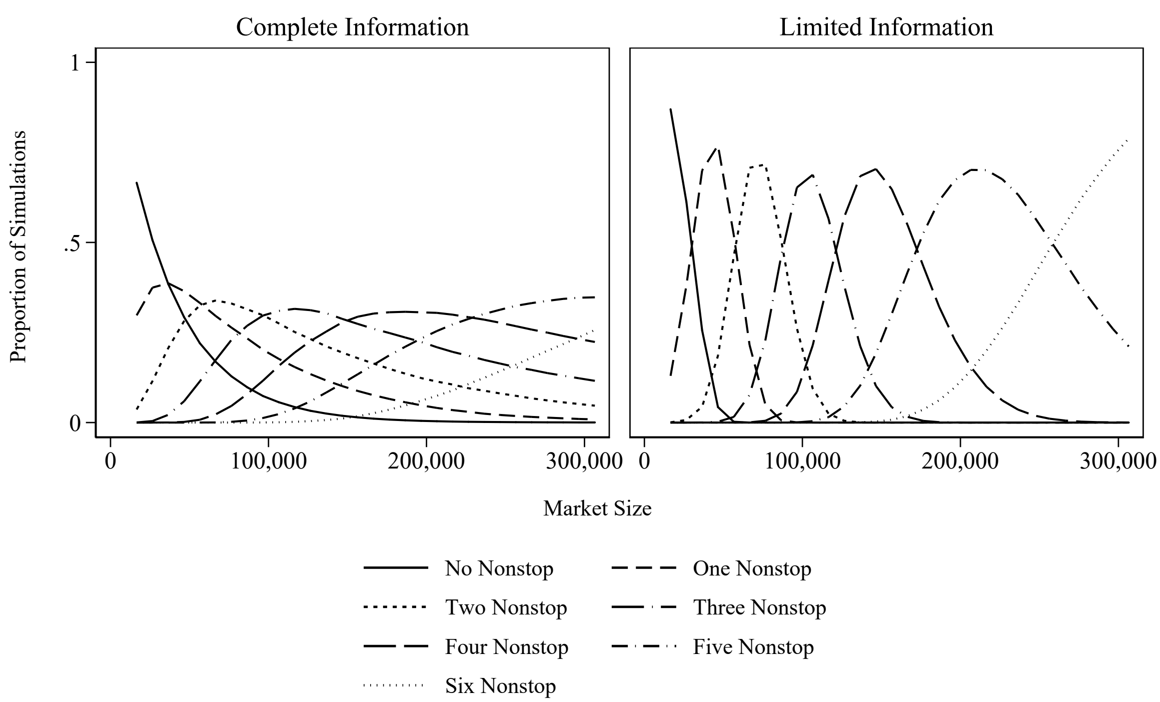

in a response to a merger will be even more limited in the counterfactuals that we now consider as observed characteristics will often place connecting carriers at a greater disadvantage than in the Austin-LAX market.

7 Merger Counterfactuals

We now perform a number of merger counterfactuals, focusing on the effects of endogenous service choices and different types of selection.

7.1 Effects of Mergers Holding Service Types Fixed

In Table 9 we predict the price effects of four different carrier mergers on four types of route when service types, qualities and marginal costs are held fixed, as is standard in a merger simulation. The four mergers include the three large legacy carrier mergers that happened after our sample period and a merger between United and US Airways which was proposed in May 2000 but abandoned in July 2001, and the types of route reflect whether the merging carriers are providing nonstop or connecting service in our Q2 2006 data. The reported pre-merger price is the average price paid by consumers traveling on the merging parties in our data. To calculate post-merger prices we back out carrier qualities and costs using our point estimates, observed prices and market shares, and then re-solve for new prices when we eliminate the merging carriers, replacing them with a "Newco" carrier which has the quality and marginal costs of the merging party with the highest average presence on the route.[40]

The results show that the predicted price effect varies according to the types of service offered by the merging carriers, which supports our distinction between nonstop and connecting service. While there is heterogeneity in the predicted effects across routes and mergers, we predict that when carriers are nonstop duopolists post-merger prices will increase by an average of 12.4%, with

[40] We could obviously consider alternative assumptions, including allowing for the merger to generate synergies, which might make the merger appear more beneficial or which could have the effect of discouraging rivals from providing nonstop service. We choose to make a plain vanilla assumption in order to avoid considering too many cases, given that we already want to consider how our results depend on the degree of selection that is assumed. We fix the nesting and price coefficients at their expected values for the market, although this is not a significant simplification given that the estimated variances for these coefficients are very small. When we invert back from prices and market shares we get a predicted marginal cost in each direction, whereas our model assumes that the marginal cost is the same in both directions. We therefore use the average of the two directional marginal costs, which are usually very close to each other, when performing counterfactuals.

Table 9: Price Effects of Mergers with Service Type Choices Held Fixed

	Delta/Northwest		United/Continental		American/US Airways		United/US Airways		Average	
Merging parties are	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Nonstop	\$566.39	\$593.20	\$503.75	\$556.17	\$459.13	\$521.15	\$479.32	\$549.49	\$481.40	\$541.25
duopolists	2 routes	+4.7%	4 routes	+10.4%	11 routes	+13.5%	7 routes	+14.6%	24 routes	+12.4%
Nonstop with	\$351.26	\$382.04	\$438.08	\$464.98	\$363.11	\$404.84	\$350.02	\$378.15	\$368.70	\$402.08
nonstop rivals	2 routes	+8.8%	4 routes	+6.1%	10 routes	+11.5%	10 routes	+8.0%	26 routes	+9.1%
Nonstop and	\$472.99	\$524.67	\$502.60	\$513.29	\$447.95	\$478.95	\$443.30	\$462.53	\$458.02	\$486.40
connecting	91 routes	+10.9%	59 routes	+2.1%	158 routes	+6.9%	163 routes	+4.3%	471 routes	+6.2%
Both connecting	\$433.26	\$444.63	\$487.04	\$486.86	\$464.20	\$457.77	\$484.25	\$479.62	\$466.00	\$465.97
	479 routes	+2.6%	334 routes	0.0%	471 routes	-1.4%	521 routes	-1.0%	1,805 routes	+0.0%

smaller price increases when there is at least one additional nonstop rival or one or both parties is offering connecting service.[41] When they are both offering connecting service we predict no price change on average, which reflects the fact that, even though there is a loss of competition, the Newco carrier can have a lower price than the pre-merger average if it has a lower marginal cost than the eliminated carrier (for example, because of a shorter connecting distance). We can also calculate changes in consumer surplus. For non-stop duopolies consumer surplus is expected to decline by \$35.09 per pre-merger traveler, whereas on routes where both carriers are connecting it is only expected to fall by \$4.91.

7.2 The United-US Airways Merger and a Proposed Remedy

In May 2000 United and US Airways announced their intention to merge. The merger was abandoned in July 2001 when it became clear that the Department of Justice would oppose it despite a number of remedies proposed by the parties, including a commitment by a third carrier, American, to initiate nonstop service on five routes where the parties were nonstop duopolists. The Department of Justice was not convinced that American would be an effective competitor on these routes even if American might find it profitable to make the commitment because of the East Coast assets that they would purchase from United if the merger was completed.[42]

Figure 5 shows routes where both US Airways and United were nonstop competitors in Q2 2006 (the routes were similar in 2000), which are mainly connections between United and US Airways hubs. Seven of these routes had no other nonstop carriers and are in our sample, while ten routes had additional nonstop competitors. The heavier line routes are the five nonstop

[41] The predicted patterns are different for the Delta/Northwest merger, which reflects different patterns in pre-merger market shares. For the other mergers, when one merging carrier is connecting and the other is nonstop, these carriers have a combined market share of 9.5% and other carriers have a combined market share of 12%. In a market where either Delta or Northwest is connecting and the other is nonstop, these carriers have a combined market share of 17% and other carriers have a combined market share of 12%, so that the merged Delta/Northwest will typically have greater market power than other merging carriers in these types of markets.

[42] "International Aviation Alliances: Market Turmoil and the Future of Airline Competition", speech by R. Hewitt Pate, Deputy Assistant Attorney General, November 7, 2001, available at: https://www.justice.gov/atr/department-justice-10 (accessed June 29, 2017): "And this summer, we announced our intent to challenge the United/US Airways merger, the second- and sixth-largest airlines, after concluding that the merger would reduce competition, raise fares, and harm consumers on airline routes throughout the United States and on a number of international routes, including giving United a monopoly or duopoly on non-stop service on over 30 routes. We concluded that United's proposal to divest assets at Reagan National Airport and American Airlines' promise to fly five routes on a nonstop basis were inadequate to replace the competitive pressure that a carrier like US Airways brings to the marketplace, and would have substituted regulation for competition on key routes. After our announcement, the parties abandoned their merger plans."

Figure 5: Routes Where United and US Airways Both Provide Nonstop Service

Table 10: Predicted Effects of United/US Airways Merger in Four Nonstop Duopoly Markets Affected By the Proposed Remedy

Selection/Service Change Considered	Mean Pre-Merger United/US Airways Price	Expected # of Rivals Initiating Nonstop Service: American	Expected # of Rivals Initiating Nonstop Service: Other Rivals	Expected Post-Merger "New United" Price
1. No Service Change	\$531.97	-	-	\$577.72 (+8.6%)
Allow Service Changes				
Selection on:				
2. Obs. and Unobserved Qualities and Costs	\$531.97	0.035	0.063	\$573.37 (+7.8%)
3. Only Obs. Qualities and Costs	\$531.97	0.148	0.298	\$563.73 (+6.0%)
4. No Selection	\$531.97	0.645	1.938	\$531.77 (-0.0%)
Remedy: AA Nonstop				
5. Obs. and Unobserved Qualities and Costs	\$531.97	1	0.030	\$566.34 (+6.5%)
6. Only Obs. Qualities and Costs	\$531.97	1	0.253	\$556.18 (+4.6%)
7. No Selection	\$531.97	1	1.820	\$529.73 (-0.4%)

duopoly routes where American agreed to provide nonstop service after the merger, and four of these are in our sample (Pittsburgh to Reagan National is too short to be included) and American was a connecting carrier in 2006 on each of these routes.[43]

In Table 10 we present the results of several counterfactuals for the four sample markets affected by the proposed remedy. The first row shows average prices on the merging parties and the post-merger Newco when we hold service types fixed, and we predict that they would have increased by \$46 or 8.6% after the merger. All of the subsequent rows allow for rival carriers to change their service types after the merger, under different assumptions about selection.

To explain the different assumptions, it is useful to begin with the second row where we allow for selection on both "observed and unobserved qualities and costs". We approximate, using simulation, the posterior distribution of the draws for the market random effect, carrier qualities and marginal costs for the type of service that is not offered (observed prices and market shares imply values for the type of service that is offered conditional on the random effect) and fixed costs of nonstop service given observed prices, market shares *and service choices* i.e., we use draws that can rationalize what we observe before the merger.[44] We perform a merger simulation using 100 sets of draws, for each market, from the posterior distributions for rival carriers, making the same assumptions about the merged "new United"'s quality and marginal costs that we made when service type choices were fixed. In all of the counterfactuals we assume that the new United will provide nonstop service, because we do not observe merging parties ceasing nonstop service on these types of routes, although this is also what we would predict in the majority of simulations. For the four routes we predict that the prices of the merged firm would increase by \$41 (7.8%) dollars on average, which is only slightly smaller than the predicted price increase when service types are held fixed. This reflects the fact that the expected number of rival carriers initiating nonstop service is only 0.1 (the probability that American initiates nonstop service is 0.035).

In the third row of Table 10 we perform a counterfactual where we use the same random effect draws that we used in the second row (this makes comparisons easier) and we infer the same carrier qualities and marginal costs for the service types that are actually chosen. For the

[43]The proposed remedy required American to provide nonstop service from Philadelphia to San Francisco or San Jose. For our counterfactuals we use Philadelphia-San Francisco as the affected market, as neither US Airways nor United provided nonstop service on the San Jose route in 2006.

[44]To reduce the computational burden we hold the price and nesting coefficients equal to their expected values for the route which is a minor simplification given their small estimated variances.

service types that are not chosen we take a random draw of incremental nonstop quality and fixed cost for each carrier given its observed characteristics but we do not select only those draws that rationalize the service choices we observed before the merger. In particular, this means that initiating nonstop service will tend to be more attractive for each connecting carrier (holding the choices of other carriers fixed) than it was in the second row. As a result we expect more carriers to initiate nonstop service (expected number is 0.45 compared to 0.1), but because the additional nonstop carriers tend to be those that find nonstop service marginally profitable, we still predict a significant post-merger price increase ($33 or 6.0%).

In the fourth row, we ignore selection in the sense that we assume that the carriers that provide connecting service in our data would be able to provide a nonstop service similar to that of the merging parties, by assuming that they would have the mean nonstop quality and mean nonstop marginal costs of the merging carriers and draw their fixed costs from a distribution with a mean equal to the average of the merging carriers.[45] If they provide connecting service they have the quality and marginal costs that are implied by the data. With these assumptions we predict that an average of 2.6 carriers would initiate nonstop service and that average prices on the merged carrier would remain almost the same as before the merger.[46]

Rows 5-7 consider the effects of the proposed remedy where American committed to initiate nonstop service using the alternative assumptions about selection, and we use the same draws for American's qualities and costs that we used in the upper part of the table. In each case, rival carriers know that both the new United and American will provide nonstop service when they make their service choices. The results demonstrate that the fact that American initiates nonstop service, which guarantees that the merger will not reduce the number of non-stop competitors, does not necessarily mean that there is a much more effective constraint on the market power of the merging firm: in rows 5 and 6 the predicted price increase is only 1.3-1.4 percentage points lower than in rows 2 and 3. The fact that American is not an effective competitor in the cases when it only enters because of the remedy is also illustrated by the fact that American's certain entry causes only a small drop in the expected number of other rival carriers initiating nonstop

[45] We view this counterfactual as approximating what might be done if one could not adequately control for how carrier characteristics affect the quality and costs of nonstop service.

[46] An alternative assumption is that other carriers would have the same nonstop quality, marginal cost and fixed cost distribution as the lower average presence merging carrier i.e., the carrier that we treat as being eliminated after the merger. In this case, we predict that 1.5 carriers would initiate nonstop service and that the merged firm's price would increase by $16 or just under 3%.

service, especially when we account for selection.

7.3 The Predicted Effects of Later Mergers Observed in the Data

We can repeat the analysis in the upper part of Table 10 for the mergers that took place after our sample period. We analyze these separately without trying to quantify their cumulative effects, as our main purpose is to understand whether the relationships between price changes, service changes and selection are robust across markets.

Table 11 presents predictions of service and price changes in markets where the merging parties are nonstop duopolists. In each case we report the average predicted post-merger price for the merged carrier and the average number of new carriers that are predicted to initiate nonstop service. While magnitudes vary, especially when we ignore selection, the basic pattern is the same as in Table 10: when selection is ignored the prices of the merging parties are not expected to increase significantly and significant new nonstop service is predicted; whereas, when we account for selection, we predict less initiation of nonstop service by rivals and significant price increases. Interestingly, as noted in Section 2, nonstop service was initiated by rival carriers in four out of seventeen nonstop duopoly markets within two years of the mergers that took place after our sample period, which is very similar to the average entry rate predicted by our model when we account for selection.

In Table 12 we perform a similar analysis for routes where the merging parties are nonstop and there was at least one nonstop rival prior to the merger (for all but one route there is exactly one nonstop rival). In these markets, we assume that, after the merger, the merged Newco is nonstop but, when we endogenize service choices, we allow the other nonstop carrier to potentially downgrade its service to connecting service. We include United/US Airways in this analysis as we did not consider these markets in our previous analysis.

When we do not allow for selection we predict significant churn in the set of carriers that are predicted to be nonstop: the probability that the observed nonstop rival ceases nonstop service is 0.33 and the overall growth in the total number of nonstop carriers comes from more connecting rivals switching to nonstop. On the other hand, when we fully account for selection the existing nonstop rival always maintains nonstop service, and, as in the duopoly case, we predict a small probability of new nonstop service. When we allow for selection only on observables we now predict a decline in the total number of nonstop rivals to the merged firm and, on average, the

Table 11: Predicted Price and Service Changes in Mergers Following the Sample Period on Routes where Merging Parties are Nonstop Duopolists

	Delta/Northwest		United/Continental		American/US Airways		Average for Completed Mergers	
	Price	Expected Numb. New Nonstop	Price	Expected Numb. New Nonstop	Price	Expected Numb. New Nonstop	Price	Expected Numb. New Nonstop
Pre-merger	\$566.39	-	\$503.75	-	\$459.13	-	\$482.25	-
Post-Merger								
Service Types Fixed	\$593.20	-	\$556.17	-	\$521.15	-	\$537.86	-
	+4.7%		+10.4%		+13.5%		+11.5%	
Allow Service Changes								
Selection on Obs. and Unobs. Qualities and Costs	\$590.34	0.07	\$547.65	0.14	\$511.33	0.21	\$529.17	0.18
	+4.2%		+8.7%		+11.4%		+9.7%	
Selection on Obs. Qualities and Costs	\$584.20	0.19	\$534.08	0.35	\$488.45	0.73	\$510.45	0.57
	+3.1%		+6.0%		+6.4%		+5.8%	
No Selection	\$573.83	0.93	\$454.36	2.62	\$460.25	2.10	\$472.23	2.08
	+1.3%		-9.8%		+0.2%		-2.1%	
Number of Routes	2		4		11		17	

Table 12: Price Effects of Mergers Where Both Merging Parties and at Least One Rival Carrier Provide Nonstop Service

	Delta/Northwest		United/Continental		American/US Airways		United/US Airways		Average	
	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals
Pre-merger	\$351.26	-	\$438.08	-	\$363.11	-	\$350.02	-	\$377.51	-
Post-Merger										
Service Types Fixed	\$382.04	-	\$464.98	-	\$404.84	-	\$378.15	-	\$412.27	-
	+8.1%		+6.1%		+11.5%		+8.0%		+9.1%	
Service Types Endogenous										
Selection on Obs.	\$378.90	0.16	\$464.86	0.01	\$404.41	0.03	\$377.24	0.05	\$411.07	0.06
and Unobs. Qualities	+7.9%		+6.1%		+11.4%		+7.8%		+8.9%	
and Costs										
Selection on Obs.	\$386.40	-0.51	\$466.18	-0.03	\$403.55	-0.27	\$375.17	-0.11	\$413.33	-0.28
Qualities	+10.0%		+6.4%		+11.1%		+7.2%		+9.5%	
and Costs										
No Selection	\$374.37	0.66	\$455.64	0.61	\$398.85	-0.03	\$367.68	0.48	\$404.95	0.34
	+6.6%		+4.0%		+9.8%		+5.0%		+7.3%	
Number of Routes	2		4		10		10		26	

largest price increases, showing that the duopoly result that predictions with observed selection lie between the no selection and full selection predictions depends on the market structure being considered. In all three cases we predict significant price increases in these markets that are broadly similar to those that we predict when service types are held fixed despite some non-trivial predicted changes in the set of non-stop competitors.

8 Conclusions

We have estimated a model of endogenous service choices and price competition in airline markets. Our model allows for carriers to have complete information about all demand and marginal cost shocks when choosing whether to provide nonstop service. As a result, no carrier will regret its choice ex-post in (pure strategy) equilibrium and, from a researcher's perspective, the set of carriers that provide nonstop service will be a selected subset of the carriers within the market. Selection can matter for counterfactuals, including the analysis of mergers, which is our focus, because selection will tend to imply that carriers will be more likely to maintain their pre-merger service choices. We find that, on average, many of the mergers that we consider would appear quite benign (in the sense that prices are not expected to rise significantly) if we allow for endogenous service choices but do not account for selection, whereas we can explain the limited service changes observed in the data when selection is accounted for. We also show that the type of remedy suggested by the parties in the case of the United/US Airways merger would likely have been ineffective in constraining market power on routes where the merging parties were nonstop duopolists, suggesting that the Department of Justice's view that the remedy was insufficient, even for these routes, was likely correct.

An important feature of our approach is that the computational burden of our estimator is not too large, especially for the purposes of academic research. This allows us to include quite rich specifications of observable controls, which turn out to be able to explain much of the variation in the data. We implement the model assuming a particular model of sequential entry, which generates a unique equilibrium prediction. However, we show that the point estimates are very similar when we do not impose this equilibrium selection rule, and instead base estimation on moment inequalities.

One could extend this research in many directions or apply the methodology to other industries. A natural extension in the airline industry would be to explicitly include carriers' capacity decisions, which are important strategic choices, but which will currently show up in our estimates as affecting quality and/or marginal costs. A model of capacity choices is also necessary to understand the effectiveness of merger remedies that involve divestitures of slots or gates at congested airports, and to analyze which carriers should be able to purchase these assets and to predict where purchasing carriers are likely to use them.

References

Ackerberg, D. (2009): "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*, 7(4), 343–376.

Aguirregabiria, V., and C.-Y. Ho (2010): "A Dynamic Game of Airline Network Competition: Hub-and-Spoke Networks and Entry Deterrence," *International Journal of Industrial Organization*, 28(4), 377–382.

——— (2012): "A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments," *Journal of Econometrics*, 168(1), 156–173.

Andrews, D. W., and P. J. Barwick (2012): "Inference for Parameters Defined by Moment Inequalities: A Recommended Moment Selection Procedure," *Econometrica*, 80(6), 2805–2826.

Andrews, D. W., and X. Shi (2013): "Inference Based on Conditional Moment Inequalities," *Econometrica*, 81(2), 609–666.

Andrews, D. W., and G. Soares (2010): "Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection," *Econometrica*, 78(1), 119–157.

Ashenfelter, O., D. Hosken, and M. Weinberg (2014): "Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers," *The Journal of Law and Economics*, 57(S3), S67–S100.

Azar, J., M. C. Schmalz, and I. Tecu (forthcoming): "Anti-Competitive Effects of Common Ownership," *Journal of Finance*.

Benkard, L., A. Bodoh-Creed, and J. Lazarev (2010): "Simulating the Dynamic Effects of Horizontal Mergers: U.S. Airlines," Discussion paper, Stanford University.

Berry, S. (1992): "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 60(4), 889–917.

Berry, S., and P. Jia (2010): "Tracing the Woes: An Empirical Analysis of the Airline Industry," *American Economic Journal: Microeconomics*, 2(3), 1–43.

Berry, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2), 242–262.

Bhattacharya, V., J. W. Roberts, and A. Sweeting (2014): "Regulating Bidder Participation in Auctions," *RAND Journal of Economics*, 45(4), 675–704.

Borenstein, S. (1990): "Airline Mergers, Airport Dominance, and Market Power," *American Economic Review*, 80(2), 400–404.

——— (2010): "An Index of Inter-City Business Travel for Use in Domestic Airline Competition Analysis," Discussion paper, UC Berkeley.

Bresnahan, T. F., and P. C. Reiss (1990): "Entry in Monopoly Markets," *Review of Economic Studies*, 57(4), 531–553.

——— (1991): "Entry and competition in Concentrated Markets," *Journal of Political Economy*, 99(5), 977–1009.

Chernozhukov, V., D. Chetverikov, and K. Kato (2016): "Testing Many Moment Inequalities," Discussion paper, University of California Los Angeles.

Chernozhukov, V., H. Hong, and E. Tamer (2007): "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, 75(5), 1243–1284.

Ciliberto, F., C. Murry, and E. T. Tamer (2016): "Market Structure and Competition in Airline Markets," Discussion paper, University of Virginia.

Ciliberto, F., and E. Tamer (2009): "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 77(6), 1791–1828.

Ciliberto, F., and J. W. Williams (2014): "Does Multimarket Contact Facilitate Tacit Collusion? Inference on Conduct Parameters in the Airline Industry," *RAND Journal of Economics*, 45(4), 764–791.

Draganska, M., M. Mazzeo, and K. Seim (2009): "Beyond Plain Vanilla: Modeling Joint Product Assortment and Pricing Decisions," *Quantitative Marketing and Economics*, 7(2), 105–146.

Dunn, A. (2008): "Do Low Quality Products Affect High Quality entry? Multiproduct Firms and Nonstop Entry in Airline Markets," *International Journal of Industrial Organization*, 26(5), 1074–1089.

Eizenberg, A. (2014): "Upstream Innovation and Product Variety in the US Home PC Market," *Review of Economic Studies*, 81(3), 1003–1045.

Fan, Y. (2013): "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market," *American Economic Review*, 103(5), 1598–1628.

Fan, Y., and C. Yang (2016): "Competition, Product Proliferation and Welfare: A Study of the US Smartphone Market," Discussion paper, University of Michigan.

Geweke, J. (1989): "Bayesian Inference in Econometric Models using Monte Carlo Integration," *Econometrica*, 57(6), 1317–1339.

Gillen, D., W. Morrison, and C. Stewart (2003): "Air Travel Demand Elasticities: Concept, Issues and Measurement," Discussion paper, Canadian Department of Finance.

Hüschelrath, K., and K. Müller (2014): "Airline Networks, Mergers, and Consumer Welfare," *Journal of Transport Economics and Policy*, 48(3), 385–407.

——— (2015): "Market Power, Efficiencies, and Entry Evidence from an Airline Merger," *Managerial and Decision Economics*, 36(4), 239–255.

Israel, M., B. Keating, D. L. Rubinfeld, and R. Willig (2013): "Airline Network Effects and Consumer Welfare," *Review of Network Economics*, 12(3), 287–322.

Jeziorski, P. (2013): "Empirical Model of Dynamic Merger Enforcement–Choosing Ownership Caps in US Radio," Discussion paper, UC Berkeley.

Koopman, S. J., N. Shephard, and D. Creal (2009): "Testing the Assumptions Behind Importance Sampling," *Journal of Econometrics*, 149(1), 2 – 11.

Li, S., J. Mazur, J. Roberts, and A. Sweeting (2015): "Airline Mergers and the Potential Entry Defense," Discussion paper, University of Maryland.

Mazur, J. (2016): "Can Stricter Bankruptcy Laws Discipline Capital Investment? Evidence from the US Airline Industry," Discussion paper, Purdue University.

Mazzeo, M. J. (2002): "Product Choice and Oligopoly Market Structure," *RAND Journal of Economics*, 33(2), 221–242.

Mizuno, T. (2003): "On the Existence of a Unique Price Equilibrium for Models of Product Differentiation," *International Journal of Industrial Organization*, 21(6), 761–793.

Morrison, S. A. (1996): "Airline Mergers: A Longer View," *Journal of Transport Economics and Policy*, 30(3), 237–250.

Nevo, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, pp. 395–421.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2015): "Moment Inequalities and their Application," *Econometrica*, 83(1), 315–334.

Peters, C. (2009): "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry," *Journal of Law and Economics*, 49(2), 627–649.

Reiss, P. C., and P. T. Spiller (1989): "Competition and Entry in Small Airline Markets," *Journal of Law and Economics*, 32(2, Part 2), S179–S202.

Roberts, J. W., and A. Sweeting (2013): "When Should Sellers Use Auctions?," *American Economic Review*, 103(5), 1830–1861.

Rosen, A. M. (2008): "Confidence Sets for Partially Identified Parameters that Satisfy a Finite Number of Moment Inequalities," *Journal of Econometrics*, 146(1), 107–117.

Seim, K. (2006): "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," *RAND Journal of Economics*, 37(3), 619–640.

Silva, J. S., and S. Tenreyro (2006): "The Log of Gravity," *Review of Economics and Statistics*, 88(4), 641–658.

Sweeting, A. (2009): "The Strategic Timing Incentives of Commercial Radio Stations: An Empirical Analysis Using Multiple Equilibria," *RAND Journal of Economics*, 40(4), 710–742.

——— (2013): "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 81(5), 1763–1803.

Sweeting, A., J. W. Roberts, and C. Gedge (2017): "A Model of Dynamic Limit Pricing with an Application to the Airline Industry," Discussion paper, University of Maryland.

Werden, G. J., A. S. Joskow, and R. L. Johnson (1991): "The Effects of Mergers on Price and Output: Two Case Studies From the Airline Industry," *Managerial and Decision Economics*, 12(5), 341–352.

Wollmann, T. (2016): "Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles," Discussion paper, University of Chicago.

APPENDICES

A Data Construction

This appendix complements the description of the data in Section 2 of the text.

Selection of markets. We use 2,028 airport-pair markets linking the 79 U.S. airports (excluding Alaska and Hawaii) with the most enplanements in Q2 2006. The markets that are excluded meet one or more of the following criteria:

- airport-pairs that are less than 350 miles apart as ground transportation may be very competitive on these routes;
- airport-pairs involving Dallas Love Field, which was subject to Wright Amendment restrictions that severely limited nonstop flights;
- airport-pairs involving New York LaGuardia or Reagan National that would violate the so-called perimeter restrictions that were in effect from these airports[47];
- airport-pairs where more than one carrier that is included in our composite "Other Legacy" or "Other LCC" (low-cost) carriers are nonstop, have more than 20% of non-directional traffic or have more than 25% presence (defined in the text) at either of the endpoint airports. Our rationale is that our assumption that the composite carrier will act as a single player may be especially problematic in these situations[48]; and,
- airport-pairs where, based on our market size definition (explained below), the combined market shares of the carriers are more than 85% or less than 4%.

Definition of players, nonstop and connecting service. We are focused on the decision of carriers to provide nonstop service on a route. Before defining any players or outcomes, we drop all passenger itineraries from DB1 that involve prices of less than \$25 or more than \$2000 dollars[49], open-jaw journeys or journeys involving more than one connection in either direction. Our next

[47] To be precise, we exclude routes involving LaGuardia that are more than 1,500 miles (except Denver) and routes involving Reagan National that are more than 1,250 miles.

[48] An example of the type of route that is excluded is Atlanta-Denver where Airtran and Frontier, which are included in our "Other LCC" category had hubs at the endpoints and both carriers served the route nonstop.

[49] These fare thresholds are halved for one-way trips.

Figure A.2: Proportion of DB1 Passengers Traveling with Connections, Based on the Type of Service

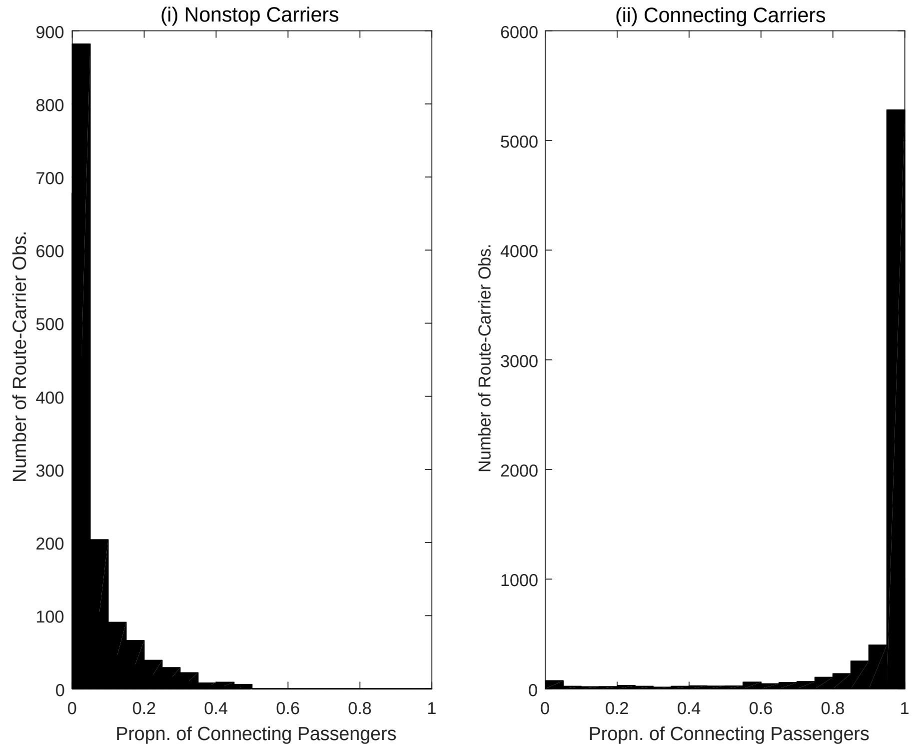

step is to aggregate smaller players into composite "Other Legacy" and "Other LCC" carriers, in addition to the "named" carriers (American, Continental, Delta, Northwest, Southwest, United and US Airways) that we focus on. Our classification of carriers as low-cost follows Berry and Jia (2010). Based on the number of passengers carried, the largest Other Legacy carrier is Alaska Airlines, and the largest Other LCC carriers are JetBlue and AirTran.

We define the set of players on a given route as those ticketing carriers who achieve at least a 1% share of total travelers (regardless of their originating endpoint) and, based on the assumption that DB1 is a 10% sample, carry at least 200 return passengers per quarter, with a one-way passenger counted as one-half of a return passenger. We define a carrier as providing nonstop service on a route if it, or its regional affiliates, are recorded in the T100 data as having at least 64 nonstop flights in each direction during the quarter and at least 50% of the DB1 passengers that it carries are recorded as not making connections (some of these passengers may be traveling on flights that make a stop but do not require a change of planes). Other players are defined as providing connecting service.

There is some arbitrariness in these thresholds. However, the 64 flight and 50% nonstop thresholds for nonstop service have little effect because almost all nonstop carriers far exceed these thresholds. For example, Figure A.2 shows that the carriers we define as nonstop typically carry only a small proportion of connecting passengers. For this reason, we feel able to ignore the fact that carriers may provide both nonstop and connecting service on the same route. On the other hand, our 1% share/200 passenger thresholds do affect the number of connecting carriers. For example, if we instead required players to carry 300 return passengers and have a 2% share, the average number of connecting carriers per market falls by almost one-third as marginal carriers are excluded.

Market Size. As in many settings where discrete choice demand models are estimated, the definition of market size is important but not straightforward. Ideally, variation in market shares across carriers and markets should reflect variation in prices, carrier characteristics and service types rather than variation in how many people consider flying on a particular route which is what the market size measure should be capturing.

A common approach is to use the geometric average of endpoint populations as the measure of market size (e.g., Berry and Jia (2010), Ciliberto and Williams (2014)).[50] However, as illustrated in the left-hand panel of Figure A.3(a), using this measure results in considerable heterogeneity in (the natural log of) total market shares (i.e., summing across all carriers) across routes. It also leads to significant variation in the proportion of the market traveling in each direction on many routes even though the services offered by the carriers are usually very similar in both directions (right-hand panel). This is a problem as we model competition on directional routes.

We address these issues in two ways. First, conditional on our market size measure, our demand model allows for a route-level random effect, unobserved to the econometrician but known to the carriers. This random effect is common to all carriers and all types of service, and it can explain why more people travel on some routes holding service, prices and observed variables constant. Second, we define market size using the regression-based gravity model of Silva and Tenreyro (2006) where the log of the number of passengers traveling on a directional route is projected onto a set of interactions between the total number of originating and destination passengers (i.e., aggregating across all carriers and routes) at the endpoint airports and the nonstop distance between the airports. We then multiply the predicted traveler number by 3.5 so that, on average, the combined market shares of carriers is just under 30%. Figure A.3(b)

[50] Reiss and Spiller (1989) use the minimum endpoint population as their market size measure.

Figure A.3: Market Size Measures and their Impact on Market Shares

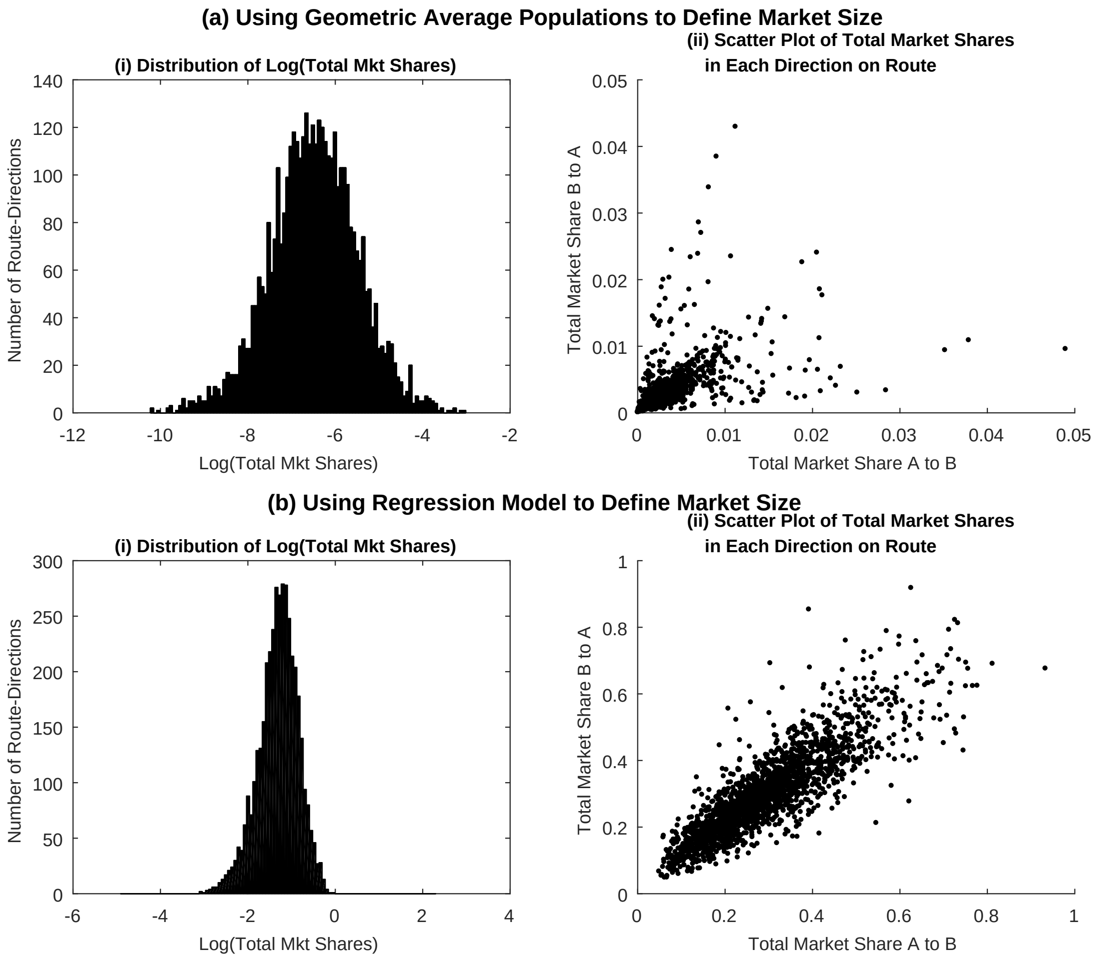

repeats the figures in Figure A.3(a) using this new definition, and the distribution of the log of total market shares and the relationship between total market shares in each direction display much more limited heterogeneity.

Prices and Market Shares. As is well-known, airlines use revenue management strategies that result in passengers on the same route paying quite different prices. Even if more detailed data (e.g., on when tickets are purchased) was available, it would likely not be feasible to model these type of strategies within the context of a combined service choice and pricing game. We therefore use the average price as our price measure, but allow for prices and market shares (defined as the number of originating passengers carried divided by market size) to be different in each direction, so that we can capture differences in passenger preferences (possibly reflecting frequent-flyer program membership) across different airports.[51]

Explanatory Variables Reflecting Airline Networks. The legacy carriers in our data operate hub-and-spoke networks. On many medium-sized routes nonstop service may be profitable only because it allows a large number of passengers who use the route as one segment of a longer trip to be served. While our structural model captures price competition for passengers traveling only the route itself, we allow for connecting traffic to reduce the effective fixed cost of providing non-stop service by including three carrier-specific variables in our specification of fixed costs. Two variables are indicators for the principal domestic and international hubs of the non-composite carriers. We define domestic hubs as airports where more than 10,000 of the carrier's ticketed passengers made domestic connections in DB1 in Q2 2005 (i.e., one year before our estimation sample). Note that some airports, such as New York's JFK airport for Delta, that are often classified as hubs do not meet our definition because the number of passengers using them for domestic connections is quite limited even though the carrier serves many destinations from the airport. International hubs are airports that carriers use to serve a significant number of non-Canadian/Mexican international destination nonstop. Table A.2 shows the airports counted as hubs for each named carrier.

We also include a continuous measure of the potential connecting traffic that will be served if nonstop service is provided on routes involving a domestic hub. The construction of this variable, as the prediction of a Heckman selection model, is detailed in Appendix A.1.

A.1 An Ancillary Model of Connecting Traffic

As explained in Section 2, we want to allow for the amount of connecting traffic that a carrier can carry when it serves a route nonstop to affect its decision to do so. Connecting traffic is especially important in explaining why a large number of nonstop flights can be supported at domestic hubs in smaller cities, such as Charlotte, NC (a US Airways hub), Memphis (Northwest) and Salt Lake City (Delta). While the development of a model where carriers choose their entire network structure is well beyond the scope of the paper, we use a reduced-form model of network

[51] Carriers may choose a similar set of ticket prices to use in each direction but revenue management techniques mean that average prices can be significantly different. Fares on contracts that carriers negotiate with the federal government and large employers, which may be significantly below list prices, may also play a role, but there is no data available on how many tickets are sold under these contracts.

Table A.2: Domestic and International Hubs for Each Named Carrier

Airline	Domestic Hub Airports	International Hub Airports
American	Chicago O'Hare, Dallas-Fort Worth, St. Louis	Chicago O'Hare, Dallas-Fort Worth, New York JFK, Miami, Los Angeles
Continental	Cleveland, Houston Intercontinental	Houston Intercontinental, Newark
Delta	Atlanta, Cincinnati, Salt Lake City	Atlanta, New York JFK
Northwest	Detroit, Memphis, Minneapolis	Detroit, Minneapolis
United	Chicago O'Hare, Denver, Washington Dulles	Chicago O'Hare, San Francisco, Washington Dulles
Southwest	Phoenix, Las Vegas, Chicago Midway, Baltimore	none
US Airways	Charlotte, Philadelphia, Pittsburgh	Charlotte, Philadelphia

flows that fits the data well[52] and which gives us a prediction of how much connecting traffic that a carrier can generate on a route where it does not currently provide nonstop service, taking the service that it provides on other routes as given. We include this prediction in our model of entry as a variable that can reduce the effective fixed or opportunity cost of providing nonstop service on the route.[53]

Model. We build our prediction of nonstop traffic on a particular segment up from a multinomial logit model of the share of the connecting passengers going from a particular origin to a particular destination (e.g., Raleigh (RDU) to San Francisco (SFO)) who will use a particular carrier-hub combination to make the connection. Specifically,

$$s_{c,i,od} = \frac{\exp(X_{c,i,od}\beta + \xi_{c,i,od})}{1 + \sum_l \sum_k \exp(X_{l,k,o,d}\beta + \xi_{l,k,od})} \tag{3}$$

where $X_{c,i,od}$ is a vector of observed characteristics for the connection (c)-carrier (i)-origin (o)-destination (d) combination and $\xi_{c,i,od}$ is an unobserved characteristic. The Xs are functions of variables that we are treating as exogenous such as airport presence, endpoint populations and geography. The outside good is traveling using connecting service via an airport that is not one of the domestic hubs that we identify.[54] Assuming that we have enough connecting passengers that the choice probabilities can be treated as equal to the observed market shares, we could potentially estimate the parameters using the standard estimating equation for aggregate data (Berry 1994):

$$\log(s_{c,i,od}) - \log(s_{0,od}) = X_{c,i,od}\beta + \xi_{c,i,od}. \tag{4}$$

However, estimating (4) would ignore the selection problem that arises from the fact that some connections may only be available because the carrier will attract a large share of connecting traffic. We therefore introduce an additional probit model, as part of a Heckman selection

[52] This is true even though we do not make use of additional information on connecting times at different domestic hubs which could potentially improve the within-sample fit of the model, as in Berry and Jia (2010). As well as not wanting to avoid excessive complexity, we would face the problem that we would not observe connection times for routes that do not currently have nonstop service on each segment, but which could for alternative service choices considered in our model.

[53] We also use the predicted value, not the actual value, on routes where we actually observe nonstop service.

[54] For example, the outside good for Raleigh to San Francisco could involve traveling via Nashville on any carrier (because Nashville is not a domestic hub) or on Delta via Dallas Fort Worth because, during our data, Dallas is not defined as a domestic hub for Delta even though it is for American.

model, to describe the probability that carrier i does serve the full *ocd* route,

$$\Pr(i \text{ serves route } ocd) = \Phi\left(W_{i,c,od}\gamma\right). \tag{5}$$

Sample, Included Variables and Exclusion Restrictions. We estimate our model using data from Q2 2005 (one year prior to the data used to estimate our main model) for the top 100 US airports. We use DB1B passengers who (i) travel from their origin to their destination making at least one stop in at least one direction (or their only direction if they go one way) and no more than one stop in either direction; and, (ii) have only one ticketing carrier for their entire trip. For each direction of the trip, a passenger counts as one-half of a passenger on an origin-connecting-destination pair route (so a passenger traveling RDU-ATL-SFO-CVG-RDU counts as $\frac{1}{2}$ on RDU-ATL-SFO and $\frac{1}{2}$ on RDU-CVG-SFO). Having joined the passenger data to the set of carrier-origin-destination-connecting airport combinations, we then exclude origin-destination routes with less than 25 connecting passengers (adding up across all connecting routes) or any origin-connection or connection-destination segment that is less than 100 miles long.[55] We also drop carrier-origin-destination-connecting airport observations where the carrier (or one of its regional affiliates) is not, based on T100, providing nonstop service on the segments involved in the connection. This gives us a sample of 5,765 origin-destination pairs and 142,506 carrier-origin-destination-hub connecting airport combinations, of which 47,996 are considered to be served in the data.

In $X_{c,i,od}$ (share equation), we include variables designed to measure the attractiveness of the carrier i and the particular *ocd* connecting route. Specifically, the included variables are carrier i's presence at the origin and its square, its presence at the destination and its square, the interaction between carrier i's origin and destination presence, the distance involved in flying route *ocd* divided by the nonstop distance between the origin and destination (we call this the 'relative distance' of the connecting route), an indicator for whether route *ocd* is the shortest route involving a hub, an indicator for whether *ocd* is the shortest route involving a hub for carrier i and the interaction between these two indicator variables and the relative distance.

The logic of our model allows us to define some identifying exclusion restrictions in the form

[55] Note while we will only use routes of more than 350 miles in the estimation of our main model, we use a shorter cut-off here because we do not want to lose too many passengers who travel more than 350 miles on one segment but less than 350 miles on a second segment.

of variables that appear in W but not in X. For example, the size of the populations in Raleigh, Atlanta and San Francisco will affect whether Delta offers service between RDU and ATL and ATL and SFO, but it should not be directly relevant for the choice of whether a traveler who is going from RDU to SFO connects via Atlanta (or a smaller city such as Charlotte), so these population terms can appear in the selection equation for whether nonstop service is offered but not the connecting share equation. In $W_{c,i,od}$ we include origin, destination and connecting airport presence for carrier i; the interactions of origin and connecting airport presence and of destination and connecting airport presence; origin, destination and connecting city populations; the interactions of origin and connecting city populations and of destination and connecting city populations, a count of the number of airports in the origin, destination and connecting cities[56]; indicators for whether either of the origin or destination airports is an airport with limitations on how far planes can fly (LaGuardia and Reagan National) and the interactions of these variables with the distance between the origin or destination (as appropriate) and the connecting airport; indicators for whether the origin or destination airport are slot-constrained. In both $X_{i,c,od}$ and $W_{i,c,od}$ we also include origin, destination and carrier-connecting airport dummies.

Results. We estimate the equations using a one-step Maximum Likelihood procedure where we allow for residuals that are assumed to be normally distributed in both (5) and (4) to be correlated, although our predictions are almost identical using a two-step procedure (correlation in predictions greater than 0.999). The coefficient estimates are in Table A.3, although the many interactions means that it is not straightforward to interpret the coefficients

To generate a prediction of the connecting traffic that a carrier will serve if it operates nonstop on particular segment we proceed as follows. First, holding service on other routes and by other carriers fixed, we use the estimates to calculate a predicted value for each carrier's share of traffic on a particular *ocd* route. Second, we multiply this share prediction by the number of connecting travelers on the *od* route to get a predicted number of passengers. Third, we add up across all *oc* and *cd* pairs involving a segment to get our prediction of the number of connecting passengers served if nonstop service is provided. There will obviously be error in this prediction resulting from our failure to account for how the total number of connecting passengers may be affected by service changes and the fact that network decisions will really be made simultaneously.

[56]For example, the number is 3 for the airports BWI, DCA and IAD in the Washington DC-Baltimore metro area.

Table A.3: Estimation Coefficients for Ancillary Model of Connecting Traffic

	Connecting Share	Serve Route	$\frac{1}{2}log\frac{1+\rho}{1-\rho}$	$log(std.\ deviation)$
Constant	4.200***	-8.712***	-0.109	0.308***
	(0.338)	(0.823)	(0.0860)	(0.0150)
Presence at Origin Airport	4.135***	6.052***		
	(0.396)	(1.136)		
Presence at Connecting Airport		11.90***		
		(0.721)		
Presence at Destination Airport	2.587***	6.094***		
	(0.396)	(1.126)		
Origin Presence * Connecting Presence		-5.536***		
		(1.311)		
Destin. Presence * Connecting Presence		-5.771***		
		(1.303)		
Population of Connecting Airport		-1.20e-07***		
		(3.16e-08)		
Origin Population * Origin Presence		-5.09e-08**		
		(2.23e-08)		
Destin. Population * Destination Presence		-4.46e-08*		
		(2.35e-08)		
Number of Airports Served from Origin		0.543***		
		(0.101)		
Number of Airports Served from Destination		0.529***		
		(0.0984)		
Origin is Restricted Perimeter Airport		0.0317		
		(0.321)		
Destination is Restricted Perimeter Airport		-0.0865		
		(0.305)		
Origin is Slot Controlled Airport		-1.098***		
		(0.321)		
Destination is Slot Controlled Airport		-1.055***		
		(0.331)		
Distance: Origin to Connection		-0.00146***		
		(0.000128)		
Distance: Connection to Destination		-0.00143***		
		(0.000125)		
Origin Restricted * Distance Origin - Connection		0.000569***		
		(0.000207)		
Destin. Restricted * Distance Connection - Destin		0.000602***		
		(0.000211)		
Relative Distance	-4.657***			
	(0.441)			
Most Convenient Own Hub	-0.357*			
	(0.192)			
Most Convenient Hub of Any Carrier	-0.574			
	(0.442)			
Origin Presence^2	-2.797***			
	(0.429)			
Destination Presence^2	-1.862***			
	(0.449)			
Relative Distance^2	0.745***			
	(0.129)			
Most Convenient Own Hub * Relative Distance^2	0.479***			
	(0.151)			
Most Convenient Hub of Any Carrier *	0.590			
Relative Distance	(0.434)			
Origin Presence * Destination Presence	-5.278***			
	(0.513)			
Observations	142,506	-	-	-

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

However we find that the estimated model does a pretty accurate job of predicting how many connecting travelers there are on the segments that airlines fly in 2005. For example, for the identified legacy carriers in our primary model, the correlation between the number of connecting passengers served on one of these segments and the number of passengers the model predicts is 0.96, and the model captures some natural geographic variation. For example, for many destinations a connection via Dallas is likely to be more attractive for a passenger originating in Raleigh-Durham (RDU) than a passenger originating in Boston (BOS), while the opposite may hold for Chicago. Our model predicts that American, with hubs in both Dallas (DFW) and Chicago (ORD), should serve 2,247 connecting DB1 passengers on RDU-DFW, 1213 on RDU-ORD and 376 on RDU-STL (St Louis), which compares with observed numbers of 2,533, 1,197 and 376. On the other hand, from Boston the model predicts that American will serve more connecting traffic via ORD (2265, observed 2765) than DFW (2040, observed 2364).

B Estimation

This Appendix provides additional information on the algorithm that we use to estimate our model. Appendix B.1 lays out the set of moments that are used in our preferred specification. Appendix B.2 explains how we estimate our model when we do not impose a known order of moves. Appendix B.3 provides Monte Carlo evidence that both estimators work well for a simplified model. Appendix B.4 provides some evidence that the algorithm works well when applied to our data.

B.1 Moments

When estimating our preferred specification, we minimize a standard simulated method of moments objective function in the second step

$$m(\Gamma)'Wm(\Gamma)$$

where W is a weighting matrix. $m(\Gamma)$ is a vector of moments where each element has the form $\frac{1}{2{,}028}\sum_{m=1}^{m=2{,}028}\left(y_m^{data} - \widehat{E_m(y|\Gamma)}\right)Z_m$, where subscript ms represent markets. y_m^{data} are observed outcomes and Z_m are exogenous observed variables.

We use a large number (1,384) of moments in estimation. To understand how we get to this number, Table B.2 presents a cross-tab describing the interactions that we use between outcomes and exogenous variables. There are two types of outcomes: market-specific and carrier-specific, and for each of these types, we are interested in prices, market shares and service choices. For example, market-specific outcomes include weighted average connecting and nonstop prices in each direction. Carrier-specific outcomes include the carrier's price in each direction, its market share in each direction and whether it provides nonstop service. The exogenous Z variables can be divided into three groups: market-level variables, variables that are specific to a single carrier, and variables that measure the characteristics of the other carriers that are in the market (e.g., Delta's presence at each of the endpoint airports when we are looking at an outcome that involves United's price or service choice).

Table B.2: Moments Used in Estimation

Exogenous Variables	Market Specific (y_M) Endogenous Outcomes 7 outcomes	Carrier Specific (y_C) Endogenous Outcomes 5 per carrier	Row Total
Market-Level Variables (Z_M) (7 per market)	49	315	364
Carrier-Specific Variables (Z_C) (up to 5 per carrier)	280	200	480
Other Carrier-Specific (Z_{-C}) (5 per "other carrier")	315	225	540
Column Total	644	740	1,384

Notes: $Z_M = \{$constant, market size, market (nonstop) distance, business index, number of low-cost carriers, tourist dummy, slot constrained dummy$\}$

$Z_C = \{$presence at each endpoint airport, our measure of the carrier's connecting traffic if the route is served nonstop, connecting distance, international hub dummy$\}$ for named legacy carriers and for Southwest (except the international hub dummy). For the Other Legacy and Other LCC Carrier we use $\{$presence at each endpoint airport, connecting distance$\}$ as we do not model their connecting traffic. Carrier-specific variables are interacted with all market-level outcomes and carrier-specific outcomes for the same carrier.

$Z_{-C} = \{$the average presence of other carriers at each endpoint airport, connecting passengers, connecting distance, and international hub dummy$\}$ for each other carrier (zero if that carrier is not present at all in the market).

$y_M = \{$market level nonstop price (both directions), connecting price (both directions), sum of squared market shares (both directions), and the square of number of nonstop carriers$\}$

$y_C = \{$nonstop dummy, price (both directions), and market shares (both directions)$\}$ for each carrier.

B.2 Estimation Using Moment Inequalities

Our baseline estimates assume that carriers play a sequential service choice game. However we also present estimated coefficients based on moment inequality estimation where we allow for the observed outcome to be associated with any pure strategy equilibrium in a simultaneous move game or a sequential move game with any order of moves. Estimation is based on moment inequalities of the form

$$\mathbb{E}(m(y,X,Z,\Gamma)) = \mathbb{E}\begin{bmatrix} y_m^{data} - \underline{\mathbb{E}(\widehat{y_m(X},\Gamma))} & \\ & \otimes Z_m \\ \overline{\mathbb{E}(\widehat{y_m(X},\Gamma))} - y_m^{data} & \end{bmatrix} \geq 0$$

where y_m^{data} are observed outcomes in the data and Z_m are non-negative instruments. $\underline{\mathbb{E}(\widehat{y_m(X},\Gamma))}$ and $\overline{\mathbb{E}(\widehat{y_m(X},\Gamma))}$ are minimum and maximum expected values for y_m given a set of parameters Γ, and these are calculated using importance sampling where, for each set of draws, we now calculate the minimum and maximum values of the outcome across different equilibria. For example, suppose that the outcome is whether firm A is nonstop. The lower bound (minimum) would be formed by assuming that whenever there are equilibrium outcomes where A is **not** nonstop, one of them will be realized, whereas the upper bound (maximum) would be formed by assuming that whenever there are equilibrium outcomes where A is nonstop, one of them is realized.[57] The instruments are the same as for the baseline estimation.

The objective function that is minimized is

$$Q(\Gamma) = \min_{t\geq 0}[m(\widehat{y,X,Z},\Gamma) - t]W[m(\widehat{y,X,Z},\Gamma) - t]$$

where t is a vector equal in length to the vector of moments, and it sets equal to zero those moment inequalities which hold so that they do not contribute to the objective function. W is a weighting matrix.

[57]Because only a subset of outcomes, or combinations of outcomes, are considered when forming moments, estimates based on these inequalities will not be sharp.

B.3 Monte Carlo

We present the results of several Monte Carlo exercises that examine the performance of our 'Simulated Method of Moments with Importance Sampling' estimator when applied to a model of airline entry. To make the Monte Carlo exercises computationally feasible, we use a slightly simpler model by reducing the number of covariates and using a binary choice of whether or not to enter a market rather than a service choice decision. However, compared with many Monte Carlos, the number of parameters that we estimate is still large, illustrating that we can accurately estimate many parameters using our approach.

B.3.1 Model

All of the Monte Carlo exercises are based on the same economic model.

Industry Participants. At the industry level, there are six carriers, A, B, C, D, E and F. A, B, C and D are 'legacy' carriers ($LEG_i = 1$) whereas E and F are low-cost carriers ($LCC_i = 1$). A carrier's legacy/low-cost status can affect both its demand and costs.

Potential Entrants. We create datasets with observations from either 500 or 1,000 independent local markets, which one can think of as airport-pairs. For each market, we first draw the number of potential entrants (2, 3 or 4 with equal probability), and then randomly choose which of the six carriers will be potential entrants.

Demand, Costs and Market and Carrier Characteristics. Each carrier has a demand quality and a marginal cost (which does not depend on quality if it enters). Carrier i's quality, $\beta^D_{i,m}$, is a draw from a truncated normal distribution

$$\beta^D_{i,m} \sim TRN(\underset{0.2}{\beta^{D,LEG}} LEG_i + \underset{0}{\beta^{D,LCC}} LCC_i + \underset{0.3}{\beta^D_1} X^D_{i,m} \text{ x } LEG_i, \underset{0.2}{\sigma^D}, -2, 10)$$

where the terms in parentheses are the mean, the standard deviation and the lower and upper truncation points respectively. The numbers beneath the Greek parameters are their true values. Carrier i's marginal cost, $c_{i,m}$, is also drawn from a truncated normal

$$c_{i,m} \sim TRN(\underset{0}{\gamma^{C,LEG}} LEG_i + \underset{-0.5}{\gamma^{C,LCC}} LCC_i + \underset{0.5}{\gamma^C_1} X^C_m, \underset{0.2}{\sigma^C}, 0, 6).$$

Each carrier also has a truncated normal fixed cost, $F_{i,m}$, that is paid if it enters the market

$$F_{i,m} \sim TRN(\underset{7500}{\theta_1^F} + \underset{1000}{\theta_2^F} X_{i,m}^F + \underset{5000}{\theta_3^F} X_m^F, \underset{2500}{\sigma^F}, 0, 30000)$$

As shown in these equations, demand and cost depend on a combination of observed market and carrier characteristics. Carrier characteristics include the carrier's type (legacy/LCC), the demand shifter $X_{i,m}^D$ (which we loosely interpret as the carrier's presence at the endpoints, and we assume that this only affects demand for legacy carriers, reflecting their greater use of frequent-flyer programs), and the carrier-specific fixed cost shifter $X_{i,m}^F$. $X_{i,m}^D$ and $X_{i,m}^F$ are drawn from independent $U[0,1]$ distributions. Market characteristics, X_m^C (which we interpret as distance) and X_m^F (which we interpret as a measure of airport congestion), affect marginal costs and entry costs. X_m^C is drawn from a $U[1,6]$ distribution. X_m^F is drawn from a $U[0,1]$ distribution.

We also allow for some additional unobserved market-level heterogeneity that affects demand. Specifically, a consumer j's indirect utility for traveling on carrier i is

$$u_{i,j,m} = \underbrace{\beta_{i,m}^D + \eta_m - \alpha_m p_{im}}_{\delta_{i,m}} + \zeta_{j,m} + (1 - \lambda_m)\varepsilon_{i,j,m}$$

where there is cross-market unobserved heterogeneity in the level of demand through a market random effect, η_m, the price sensitivity parameter, α_m, and the nesting parameter, λ_m. $\varepsilon_{i,j,m}$ is the standard Type I extreme value logit error. We make the following distributional assumptions:

$$\eta_m \sim TRN(0, \underset{0.5}{\sigma^\eta}, -2, 2)$$

$$\alpha_m \sim TRN(\underset{0.45}{\mu^\alpha}, \underset{0.1}{\sigma^\alpha}, 0.15, 0.75)$$

$$\lambda_m \sim TRN(\underset{0.7}{\mu^\lambda}, \underset{0.1}{\sigma^\lambda}, 0.5, 0.9)$$

where setting the mean of the random effect to zero is a normalization as we included separate mean quality coefficients for legacy and LCC carriers. Market size is assumed to be observed, and is drawn from a uniform distribution on the interval 10,000 to 100,000.

Order of Entry We study Monte Carlos under different assumptions on the equilibrium being played and what the researcher knows about equilibrium selection. In each case there is complete information and carriers set prices simultaneously once entry decisions have been made. We assume that the true model is that there is sequential entry. Legacy carriers are assumed to move first, ordered by $X^D_{i,m}$ (highest first), followed by low-cost carriers who are ordered randomly. The firms know the order. Firms enter when they expect their profits from entering to be greater than zero. Given the specification of the entry game, and the fact that there will be a unique equilibrium in any of the pricing games that follow entry[58], the game will have a unique subgame perfect Nash equilibrium.

B.3.2 Summary Statistics

We briefly summarize some of the patterns that emerge when we simulate outcomes for 2,000 markets given these parameters. 15.1% of the markets have no entrants, while 51.8%, 28.0%, 4.6% and 0.5% of markets have one, two, three and four entrants respectively. In 11.7% of markets, all of the potential entrants enter. 48.8% and 26.0% of legacy and LCC potential entrants enter respectively, which partly reflects the demand advantage of legacy carriers, but also their first mover advantage in the entry game. Variation in market size and the demand parameters α (price coefficient), λ (nesting coefficient) and η (market demand random effect) have sensible effects on entry. Moving from the lowest to the highest tercile of market size increases the average number of entering firms from 0.7 to 1.7. Similarly, going from the lowest to the highest tercile of $-\alpha$ (demand become less price sensitive), λ (carriers become closer substitutes) and η (market demand increases) changes the average number of entrants from 1.4 to 2.0, from 2.0 to 1.4 and from 1.5 to 1.9 respectively. There are both direct and indirect (via entry) effects on prices. For example, going from the lowest to the highest tercile of $-\alpha$ increases average prices, from 3.2 to 3.8, consistent with demand becoming less elastic, but it also increases the standard deviation of prices, from 1.0 to 1.4, because prices will tend to fall if more entry occurs. We also observe the standard deviation of prices increasing with λ (1.1 to 1.5). This reflects the fact that, because entering carriers will be closer substitutes when the nesting parameter is large, there will be a greater spread between monopoly and duopoly prices. Observed market marginal cost and fixed cost shifters also affect both price and entry outcomes. For example, going from the lowest to the highest tercile of the marginal cost shifter (X^C_m) increases average prices from

[58]This follows from Mizuno (2003) due to the assumptions that demand has a nested logit structure, each firm produces a single product and marginal costs are non-decreasing with quantity.

2.7 to 4.6, while reducing the number of entrants from 1.9 to 1.6. For the market fixed cost shifter (X_m^F) moving from the lowest to the highest tercile reduces expected entry from 1.9 to 1.6 carriers, and because of the reduced entry, average prices increase from 3.4 to 3.7.

B.3.3 Monte Carlo Exercises

There are 17 parameters, $\Gamma = \{\beta^{D,LEG}, \beta^{D,LCC}, \beta_1^D, \sigma^D, \gamma^{C,LEG}, \gamma^{C,LCC}, \gamma^C, \sigma^C, \theta_1^F, \theta_2^F, \theta_3^F, \sigma^F, \sigma^\eta,$ μ^α, σ^α, μ^λ, $\sigma^\lambda\}$, to be estimated. Label the true parameters Γ_0. We present results for three Monte Carlo exercises below.

Monte Carlo Exercise 1: Estimation When the True Distributions Are Used To Form the Importance Sampling Density & Known Order of Entry. Recall that an importance sampling estimate of the expected value for a particular outcome h_m in market m, $\widehat{E(h_m)}$, is calculated as

$$\frac{1}{S}\sum_{s=1}^{S} y(X_m, \theta_{ms}) \frac{f(\theta_{ms}|x_{m,}\Gamma')}{g(\theta_{ms}|X_m)}$$

where, in our setting, θ_{ms} is a vector of draws for the market-level parameters and demand and cost draws for all of the potential entrant carriers, f is the density of these draws given parameters Γ', g is the importance density from which θ_{ms} is drawn, and $y(X_m, \theta_{ms})$ is the value of the outcome of interest given observed market characteristics and θ_{ms} (e.g., a dummy for whether firm A enters, or the combined market share of entrants).

In the first exercise, we use the true distribution as the importance density, i.e., $g(\theta_{ms}|X_m) \equiv f(\theta_{ms}|x_{m,}\Gamma_0)$. While this estimator is generally infeasible, it is the efficient estimator in the sense that the variance of the importance sampling estimate of each expected outcome is minimized. It therefore provides a benchmark against which we can compare other results.

To perform this exercise, we first create one hundred datasets, each with 1,000 markets. We perform the estimation using 1,000 importance sampling draws per market.[59] We use the following observed outcomes in estimation: the entry decision (represented by a 0/1 dummy), the price and the market share of each of the firms (A-F)[60], and three market outcomes: the

[59] We first create the data and 2,000 draws for 2,000 different markets. Given that the importance density is the true density of the parameters, this effectively involves doing 2,001 sets of draws, and arbitrarily calling the first set 'data'. We then create the one hundred datasets. For each dataset, we draw 1,000 markets from the sample of 2,000 without replacement and, for each of the drawn markets, taking a sample of one thousand draws, without replacement, from the 2,000 that were created for that market.

[60] Obviously, if a carrier is not a potential entrant in a particular market these outcomes will be zero.

average transaction price (i.e., the average price of the entrants weighted by their market shares), the sum of squared market shares for the entering carriers[61] and the square of the number of entrants.

These outcome measures are then interacted with several observed variables to create moments for estimation. Market-level variables include a constant, market size, X_m^C, X_m^F and the number of LCC potential entrants. Carrier-level variables are $X_{i,m}^D, X_{i,m}^F$ and the average of these variables for *other* potential entrants, although we do not use $X_{i,m}^D$ for the LCC carriers as, by assumption, it does not affect their demand or their entry order. We then create moments by interacting market outcomes with the market-level variables and the carrier variables for each of the six carriers, and the carrier outcomes with the market level variables and the carrier variables for that firm. This gives us a total of 237 moments for estimation. We weight these moments by the inverse of their variances (evaluated at the true parameters, which, recall, we are using to form the importance densities) in forming the objective function.[62]

Column (1) of Table B.3 reports the mean and standard deviation of the parameters estimated for the one hundred repetitions. For all of the parameters, the mean estimated value is close to the true value, indicating that there is no systematic bias, and the standard deviations are small enough that, if they were interpreted as standard errors, all of the parameters whose true values are not equal to zero, would be statistically different from zero at the 5% level, with the exception of θ_2^F.

Another way of assessing the accuracy of the Monte Carlo estimates is by looking at how accurately we are able to predict how market outcomes would change in response to a change in the market environment. As an illustration we consider an increase in mean fixed costs of all legacy carriers by 10,000 (taking their mean fixed cost from 10,500 to 20,500). The fixed costs of LCC carriers are not affected. The first column of Table B.4 reports the expected changes in entry, the cumulative market share of entering carriers and average prices under the true parameters.[63] As expected, fewer legacy carriers enter, while there is some increased entry by LCCs. The reduction in entry causes weighted average prices to rise and the number of travelers

[61]For this calculation, market shares are defined allowing for some consumers to purchase the outside good so this is not the same as the HHI.

[62]We found that in practice the estimator performed more reliably from a wider range of starting values when we used a diagonal weighting matrix rather than the usual inverse covariance matrix of the moments.

[63]We use the outcomes for the 2,000 markets in our "data", and then re-compute outcomes increasing the fixed costs of legacy carriers but leaving the other draws unchanged.

Table B.3: Monte Carlo Results with Known Order of Entry

			(1)	(2)	(3)
		IS Density:	Same As True Distribution	50% Increase in Std. Devs.	Same As True Distribution
		# of Mkts.:	1000	1000	500
		# of IS Draws:	1000	1000	1000
		True Values			
Market Demand Random Effect	Std. Dev. (σ^{η})	0.5	0.494	0.473	0.511
			(0.073)	(0.151)	(0.089)
Mkt Demand Slope	Mean (μ^{α})	-0.45	-0.421	-0.429	-0.425
			(0.024)	(0.026)	(0.022)
	Std. Dev. (σ^{α})	0.1	0.038	0.075	0.051
			(0.039)	(0.056)	(0.044)
Nesting Parameter	Mean (μ^{λ})	0.7	0.694	0.689	0.701
			(0.033)	(0.035)	(0.054)
	Std. Dev. (σ^{θ})	0.1	0.051	0.062	0.089
			(0.039)	(0.033)	(0.072)
Carrier Quality	Legacy ($\beta^{D,LEG}$)	0.2	0.189	0.190	0.188
			(0.064)	(0.103)	(0.076)
	LCC ($\beta^{D,LCC}$)	0	0.000	-0.031	0.003
			(0.064)	(0.087)	(0.069)
	$X^D_{i,m} * LEG_i$ (β^D_1)	0.3	0.295	0.295	0.293
			(0.067)	(0.142)	(0.097)
	Std. Dev. (σ^D)	0.2	0.176	0.209	0.170
			(0.043)	(0.064)	(0.050)
Carrier Marginal Cost	Legacy Constant ($\gamma^{C,LEG}$)	0	0.031	0.040	0.054
			(0.111)	(0.133)	(0.130)
	LCC Constant ($\gamma^{C,LCC}$)	-0.5	-0.507	-0.470	-0.483
			(0.135)	(0.141)	(0.158)
	X^C_m (γ^C)	0.5	0.500	0.479	0.489
			(0.034)	(0.047)	(0.042)
	Std. Dev. (σ^C)	0.2	0.216	0.169	0.205
			(0.069)	(0.081)	(0.072)
Carrier Fixed Cost	Constant ($\theta^F_1/10,000$)	0.75	0.738	0.725	0.743
			(0.096)	(0.131)	(0.101)
	$X^F_{i,m}$ ($\theta^F_2/10,000$)	0.1	0.110	0.118	0.121
			(0.081)	(0.166)	(0.130)
	X^F_m ($\theta^F_3/10,000$)	0.5	0.556	0.599	0.548
			(0.126)	(0.163)	(0.142)
	Std. Dev. ($\sigma^F/10,000$)	0.25	0.210	0.246	0.209
			(0.065)	(0.084)	(0.086)

Notes: Reported numbers are the mean estimates of each parameter across 100 repetitions, with the standard deviations reported in parentheses.

Table B.4: Illustrative Counterfactual: The Effects of Increasing the Fixed Entry Costs of Legacy Carriers Using Parameters Estimated Using IS Distributions that are the Same as True Distribution of the Parameters and 1,000 Markets

Change in ...	Using True Parameters	Mean (Std. Dev.) Prediction Across MC Repetitions
Total Number of Entrants	-0.335	-0.332
		(0.0254)
Number of Legacy Entrants	-0.493	-0.478
		(0.0332)
Number of LCC Entrants	+0.158	+0.145
		(0.0241)
Total Market Share	-0.054	-0.053
		(0.003)
Average Price (conditional on	+0.228	+0.219
at least one firm entering)		(0.056)

to fall.[64] The second column reports the mean changes and standard deviations (in parentheses) across the 100 Monte Carlo repetitions.[65] We can see that the Monte Carlo counterfactuals predict the true effects accurately, with small standard deviations.[66]

Column (3) of Table B.3 shows the results when there are only 500 markets, rather than 1,000, in each of the datasets (we continue to use 1,000 importance draws for each market). In this case, the standard deviation of the parameter estimates increase, but only by a relatively small amount, while the means remain very close to the true values of the parameters. We also note that with either 500 or 1,000 markets, estimation is quite quick: each optimization takes less than four hours even when we rely on numerical derivatives. We also get similar Monte Carlo results when starting each optimization at parameters that are significantly perturbed from their true values.[67]

[64] Average prices are only calculated for markets where entry occurs, so average prices are calculated for the subset of markets where entry occurs before the increase in fixed costs.

[65] To isolate the effects of using different parameters, we use the same percentile for each parameter draw as in our "data" for each market, before calculating predicted outcomes with and without the increase in legacy carrier fixed costs. So, for example, suppose that in market 17 (out of 2,000), α_m was drawn from the 43$^{\text{rd}}$ percentile of the true distribution that has (untruncated) mean -0.45 and standard deviation 0.1. When we are considering a Monte Carlo repetition where the estimates of the mean and standard deviation of α are -0.6 and 0.2, we would use the 43$^{\text{rd}}$ percentile draw from this distribution.

[66] The standard deviation for the predicted change in prices is larger simply because differences in predictions of entry, either with or without the change in fixed costs, can have a large effect on prices. However, the mean prediction is close to the true value.

[67] This comment comes with the caveat that in a small number of cases when we start with perturbed parameters, a parameter drifted to some very extreme value (e.g. an estimated mean of the untruncated distribution of the nesting parameter λ of -9.96, whereas only values of λ between 0 and 1 can be rationalized if consumers maximize their utility) in which case we rejected the repetition and added a new repetition. We only drop estimates that

Monte Carlo Exercise 2: Estimation When Wider Distributions Are Used To Form the Importance Sampling Density & Known Order of Entry Our second exercise considers the case where we use an importance distribution that is more dispersed than the true parameters. This reflects the fact that in practice we do not know what the true parameters are and that, when estimating unknown parameters, it makes sense to use an importance distribution that will contain some draws that will still have reasonable density when the parameters are changed. As an illustration, we therefore repeat the first exercise, but the importance distributions are formed by increasing all of the standard deviation parameters by 50%. The mean parameters remain unchanged. Column (2) of Table B.3 reports the results when each dataset contains 1,000 markets. The mean estimates continue to be very close to the true parameter values. The standard deviations increase for most parameters, as one might expect, but the magnitude of the increases is fairly small.

Monte Carlo Exercise 3: Estimation When the Econometrician Only Knows that a Pure Strategy Nash Equilibrium is Played Our third exercise considers estimation when we relax the assumption that entry decisions are made in a known sequential order. Instead, we follow the strand of the literature (most notably, Ciliberto and Tamer (2009)) that has based estimation on moment inequalities formed under the assumption that firms play some pure strategy Nash equilibrium in a simultaneous move game.[68] The idea is that, as long as the set of equilibrium outcomes (i.e., entry decisions, prices and market shares) can be enumerated, one can use the set to calculate lower and upper bound predictions for moments of the data, and then, in estimation, search for the parameters that make inequalities based on these lower and upper bounds hold.

We keep the same assumptions on the set of potential entrants, demand and costs as in the previous exercises. The change is that now we assume that the potential entrants make entry decisions simultaneously and that they play a complete information, pure strategy Nash equilibrium (as competition always reduces profits, at least one pure strategy Nash equilibrium will exist). With at most six potential entrants it is straightforward to find all of the pure strategy Nash equilibria for a given draw of all of the cost and demand shocks. When creating

are truly extreme as in this example. We also observed examples where μ^α drifted to extreme values.

[68] In our application we also allow for the observed outcome to be the equilibrium outcome in a sequential move game with any order.

our data, we choose an equilibrium randomly if more than one equilibrium exists for a given set of draws. Given the assumed parameters, there are multiple equilibrium outcomes in 24.2% of the 2,000 sample data markets. In most cases, the equilibria differ only in the identity of entrants rather than the number of firms that enter.

The details of estimation are explained in Appendix B.2 and we follow Exercise 1 in using the true distributions of the parameters when taking our importance sample draws. The one difference to what we do in the text is that we restrict ourselves to examining pure strategy equilibria in a simultaneous move game, rather than also allowing for sequential move games with any order.

There are now many papers that propose approaches for inference for moment inequality models (for example, Chernozhukov, Hong, and Tamer (2007) Rosen (2008), Andrews and Soares (2010), Andrews and Barwick (2012), Andrews and Shi (2013), Pakes, Porter, Ho, and Ishii (2015)), and these methods often involve a significant amount of simulation making them somewhat impractical for a Monte Carlo where the procedure would have to be repeated multiple times. For our example, we therefore restrict ourselves to minimizing the objective function and reporting the mean and standard deviations (across Monte Carlo runs) of the objective function-minimizing parameters. While asymptotically the objective function should be equal to zero at the true parameters (all of the inequalities satisfied), in practice we always found that the objective function was minimized slightly above zero by a unique set of parameters (the mean minimized value is 0.0026, with a standard deviation of 0.001 across our Monte Carlo runs).[69] Table B.5 reports the Monte Carlo results, using 1,000 markets and 1,000 IS draws for each market in each Monte Carlo run.[70]

Comparing the results to those from column (1) of Table B.3 (which used the same number of observation and the same distribution to generate the importance sample draws), we see that the estimator performs almost as well, with all of the mean parameters close to their true values with the exception of the standard deviation of the carrier quality which is underestimated. The standard deviations of the estimated parameters also remain similar. Of course, it is possible that estimates would become less accurate if we assumed parameters that generated multiple

[69] As before we use the inverse of the variance of the moments, evaluated at the true parameters, as the weighting matrix.

[70] As in Exercise 1 we initially create a sample of 2,000 markets and 2,000 IS draws for each market, and then randomly sample from these sets when creating datasets for each Monte Carlo run.

Table B.5: Monte Carlo Results with Unknown Equilibrium Selection in a Simultaneous Move Game

Parameters		True Value	Estimated Value Mean (Std. Dev.)
Market Demand Random Effect	Std. Dev. (σ^{η})	0.5	0.472 (0.078)
Market Demand Slope	Mean (μ^{α})	-0.45	-0.422 (0.020)
	Std. Dev. (σ^{α})	0.1	0.072 (0.037)
Nesting Parameter	Mean (μ^{λ})	0.7	0.744 (0.057)
	Std. Dev. (σ^{λ})	0.1	0.113 (0.085)
Carrier Quality	Legacy constant ($\beta^{D,LEG}$)	0.2	0.191 (0.081)
	LCC constant ($\beta^{D,LCC}$)	0	0.004 (0.066)
	$X_{i,m}^{D} * LEG_i$ (β_1^D)	0.3	0.281 (0.121)
	Std. Dev. (σ^{D})	0.2	0.091 (0.036)
Carrier Marginal Cost	Legacy constant ($\gamma^{C,LEG}$)	0	-0.020 (0.127)
	LCC constant ($\gamma^{C,LCC}$)	-0.5	-0.562 (0.126)
	X_m^C	0.5	0.488 (0.032)
	Std. Dev. (σ^{C})	0.2	0.189 (0.059)
Carrier Fixed Cost	Constant ($\theta_1^F/10,000$)	0.75	0.696 (0.104)
	$X_{i,m}^F$ ($\theta_2^F/10,000$)	0.1	0.213 (0.147)
	X_m^F ($\theta_3^F/10,000$)	0.5	0.586 (0.109)
	Std. Dev. ($\sigma^F/10,000$)	0.25	0.204 (0.060)

Notes: Reported numbers are the mean estimates of each parameter across 100 repetitions, with the standard deviations reported in parentheses.

equilibria in a higher proportion of markets.

One of the advantages of using importance sampling, with or without equilibrium selection, is that the objective function is smooth, so that we can use derivatives to find the minimum. In Figure B.2 we examine the the shape of the objective function using moment inequalities based on the first Monte Carlo run when we change each of the parameters in turn. The black dot on each horizontal axis marks the true value of the parameter. On the other hand, for three parameters (γ^C, $\beta^{D,LEG}$, $\beta^{D,LCC}$) it is also clear that there are multiple local minima even when we are only changing a single parameter at a time. The fact that objective function can have multiple local minima makes the a second feature of the importance sampling approach, the ability to calculate the value of the objective function quickly, without having to re-solve a large number of games, particularly valuable.

B.4 Performance of the Estimation Algorithm Using the Actual Data

In this section we examine two features of the estimator in the context of our application for the case where we assume a known, sequential order of entry (i.e., the estimates in column (1) of Table 3). Figure B.3 shows the shape of the continuous objective function when we vary the parameters one-at-a-time around their estimated values. While these pictures do not show the shape of the objective function is well-behaved in multiple dimensions, there is at least some grounds for optimism that a global minimum has been found.

We also address the question of whether our importance sampling estimator satisfies the condition that the variance of $y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$ must be finite, identified by Geweke (1989). One informal way to assess this property in an application (Koopman, Shephard, and Creal (2009)) is to plot how an estimate of the sample variance changes with S, and, in particular, to see how 'jumpy' the variance plot is as S increases. The intuition is that if the true variance is infinite, the estimated sample variance will continue to jump wildly as S rises. Figure B.4 shows these recursive estimates of the sample variance for the moments associated with the three market-level outcomes, namely the weighted nonstop fare, the weighted connecting fare and the quantity-based sum of squared market shares for the carriers in the market, for the estimated parameters. The log of the number of simulations is on the x-axis and the variance of $\frac{1}{M}\sum y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms})}$ across simulations $s = 1, .., S$ is on the y-axis. Relative to examples in Koopman, Shephard, and Creal (2009), the jumps in the estimated sample variance are quite

Figure B.2: Shape of the Objective Function Based on Inequalities Around the Estimated Parameters for the First Monte Carlo Run (black dot marks the true value of the parameter)

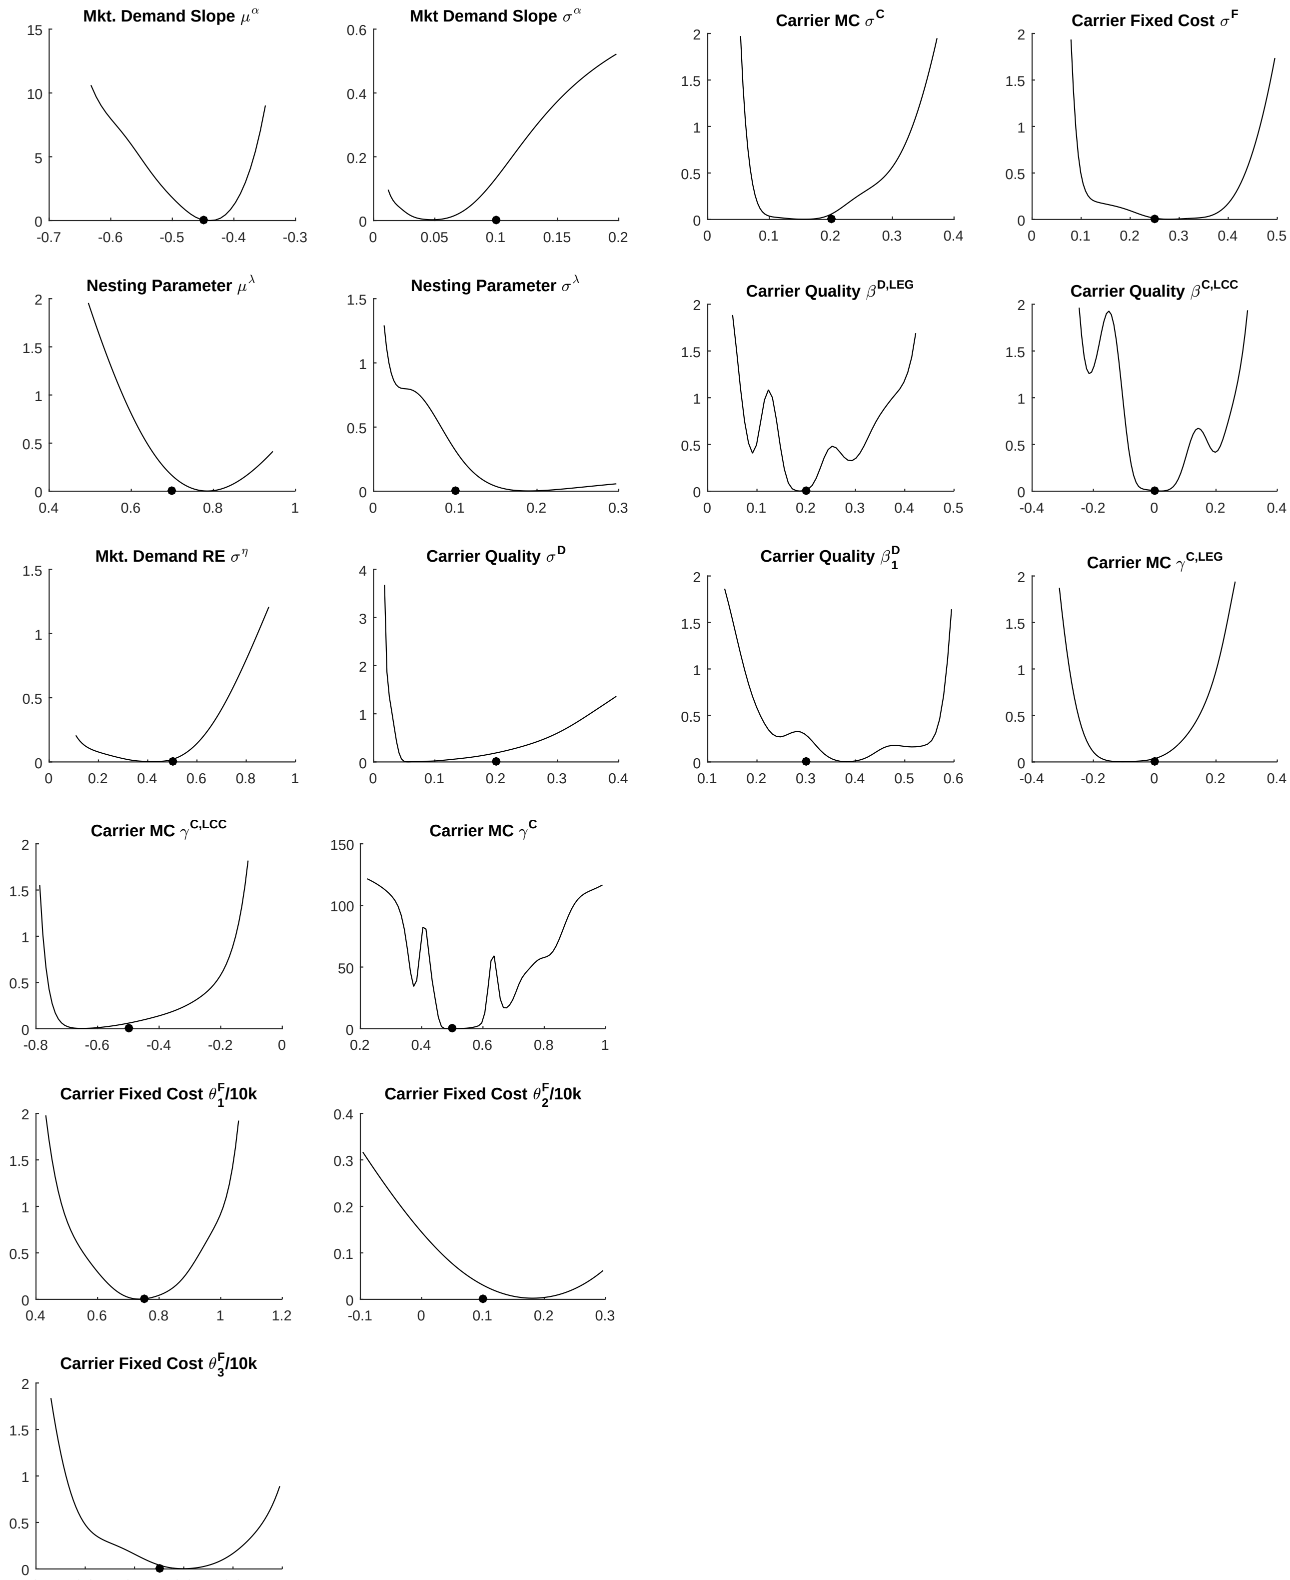

Figure B.3: Shape of the Objective Function Around the Estimated Parameters For the Parameter Estimates in Column (1) of Table 3 (black dot marks the estimated value)

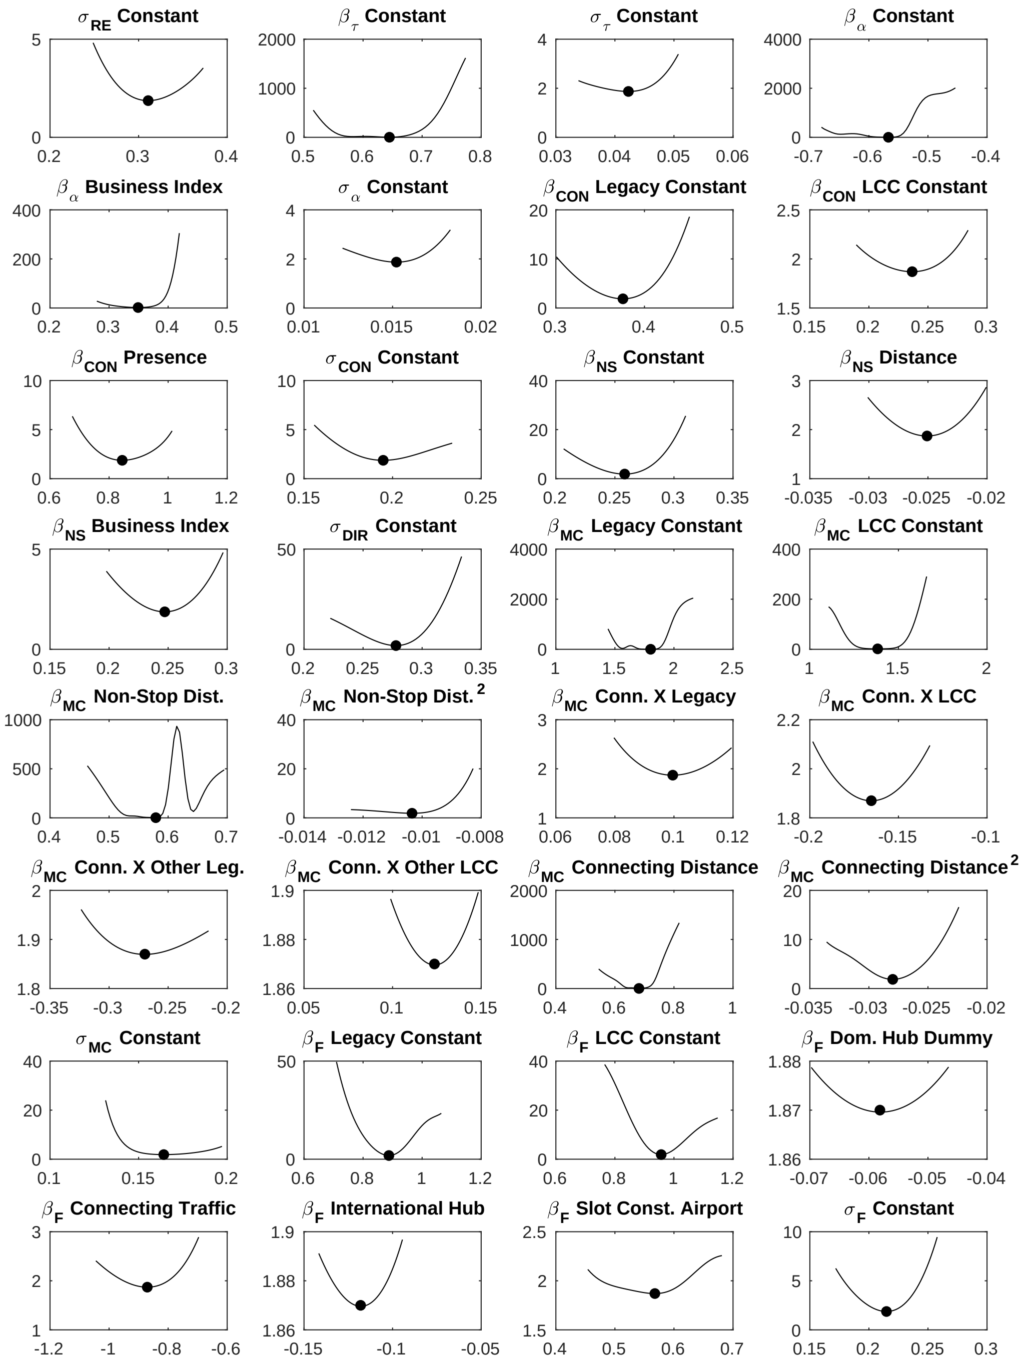

Figure B.4: Sample Variance of Three Moments as the Number of Simulation Draws is Increased (logarithm of the number of draws on the x-axis)

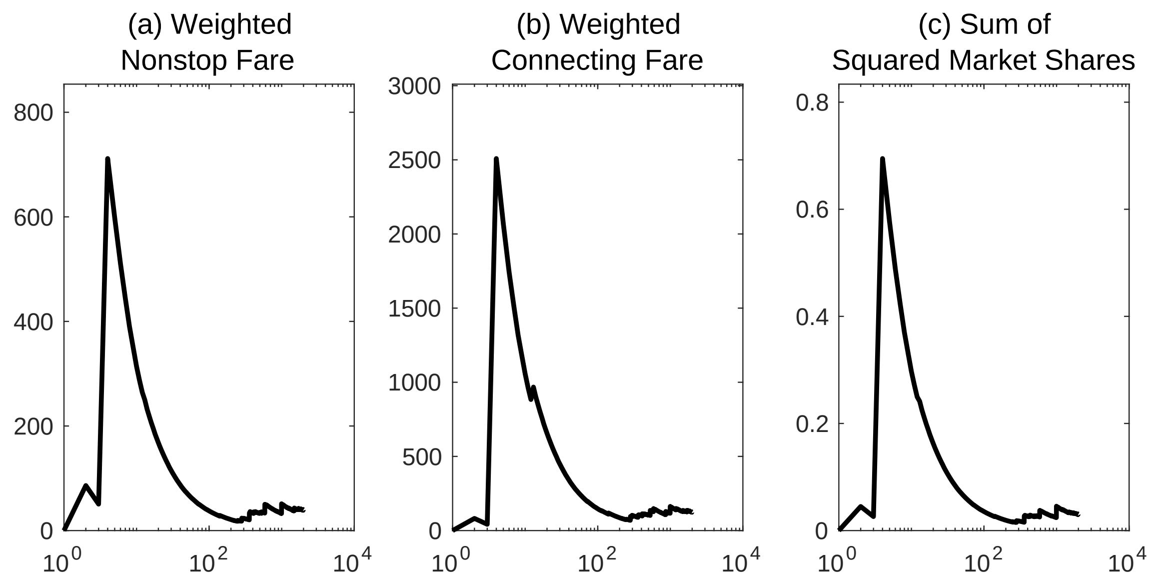

small for $S > 500$. In our application we are using $S = 1,000$.