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

November 2017

Abstract

We develop and estimate a model of service choice and price competition in airline markets. Our model allows for the carriers that provide non-stop service to be a selected subset of the carriers competing on the route. The model can be estimated without an excessive computational burden, so that we can allow for up to nine competitors, directional demand and pricing, and for carrier demand and costs to depend on many observable variables. These specifications allow us to fit service patterns in the data accurately and contribute to making our estimates robust to alternative equilibrium selection assumptions. We use the model to show how selection affects equilibrium market structure and to investigate the effects of mergers as a function of the types of service offered on the route and allowing for rivals to change their service in response to the merger. We show that the predictions, and the analysis of possible merger remedies, depends on what is assumed about selection.

1 Introduction

When mergers are proposed in concentrated 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 formal equilibrium models that assume 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 often based on more informal analyses

*Corresponding author: sweeting@econ.umd.edu. Some of this research was conducted 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. 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".

of the history of entry or repositioning in the industry and perceptions of the firms' likely business plans. While the lack of formal modeling may seem surprising given the large "discrete choice/entry games" literature 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 would allow the likelihood of repositioning or its sufficiency in constraining market power to be easily or convincingly quantified.

In this paper, we develop and estimate an integrated model of positioning and price competition and use it to analyze post-merger service-type choices in the airline industry. Our service-type choices involve carriers deciding whether to offer non-stop or connecting service on a particular route. We make the conventional assumption that carriers first choose the type of service that they offer, and then, in a second stage, choose equilibrium prices given these service choices. The distinction between non-stop and connecting service has usually been ignored in the airline entry literature, even though it plays an important role in the analysis of actual airline mergers (Dunn (2008)).[1]

A key feature of our model is that we allow for the carriers that provide non-stop service in our data to be a selected subset of the carriers that compete in the market, where selection may reflect both observable and unobservable factors that affect carrier quality, carrier marginal costs or carrier fixed costs. Selection will mean that carriers will be seen offering non-stop service when they can offer a more competitive non-stop product (i.e., one that has higher quality or lower marginal cost), and that carriers that currently only provide connecting service may have limited ability to compete effectively at providing non-stop service if, for some reason, they had to do so. Selection follows from our assumption that, for the carriers, there is complete information about all carrier characteristics throughout the game, so that, in the first-stage, carriers take service decisions knowing both their own and their rivals' qualities and costs from offering different types of service, whether or not these qualities or costs depend on variables that the researcher can observe.

We show that selection on both observables and unobservables will matter for equilibrium market structure, and, as a consequence, for what will happen after mergers. For example, if carriers had to decide whether to provide non-stop service with information only on their fixed costs we find that there would be significantly more non-stop carriers in larger markets, but that, for a given market structure, consumer surplus will tend to be significantly lower (because it may be fairly low quality or high cost carriers that provide non-stop service) compared with our model. Without selection, market size is a much more important determinant of how many carriers provide non-stop service, implying that a merger that eliminates one non-stop carrier will be quite likely to induce another carrier to add this type of service. Consistent with this, we find that when we perform merger counterfactuals but ignore selection, we predict that mergers would typically cause other carriers to initiate non-stop service, and that this would be sufficient to prevent prices from rising. A merger may even cause consumer surplus to increase, as in the theoretical model of Cabral (2003). When we account for selection, initiation of non-stop service is less likely, and, if it happens, it is also less likely to be effective at constraining

[1] See also the Department of Justice's 2013 Competitive Impact statement on the American/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/Continental merger, http://www.gao.gov/new.items/d10778t.pdf (accessed June 26, 2017).

market power. We illustrate these points in the context of mergers that were consummated after the period of our data and by considering a merger between United and US Airways, that was proposed prior to our data, in July 2001, but abandoned. For these mergers we find that, on average, our model, which allows for endogeneity of service choices and accounts for selection implies that the average prices paid by travelers on the merging carriers will increase by an amount (over 8% on average) as in a model where service choices are held fixed, on routes where the merging parties are non-stop duopolists before the merger. This is true even though these mergers do lead to new non-stop service being provided by rivals on around 30% of routes, because selection implies that most of these new services are only marginally profitable and they are relatively ineffective at constraining market power.

We also examine a remedy which was suggested by the parties in the United/US Airways case. Specifically, a third carrier, American Airlines, offered to commit to providing non-stop service for ten years on a set of routes, such as Denver-Philadelphia, where only the merging parties provided non-stop service. We show that when we account for selection this remedy appears ineffective at constraining market power, even though it guarantees that the number of non-stop carriers does not fall. To be specific, even though it increases the probability that American will enter after the merger by a factor of eight, the predicted post-merger price increase is reduced by only a couple of percentage points relative to the prediction without the remedy. This suggests that the Department of Justice was right to reject the proposed remedy as inadequate.

As we discuss below, the objective of building an estimable model for analyzing post-merger repositioning is shared by several papers in the literature. The main contribution of our paper is that we allow for selection, which we show can have important effects on counterfactuals, while maintaining a computational burden that is small enough that we can use relatively rich specifications. The rich specifications lead us to find that much of the selection into non-stop service is based on observable carrier or market characteristics, which makes our results and counterfactuals easier to interpret.

Estimating models that combine discrete service choices and continuous price competition is challenging because, first, it is not possible to separate the estimation of the demand and marginal cost functions from the fixed cost parameters in the discrete choice game; second, it is computationally expensive to re-solve equilibrium pricing games many times during estimation, with the burden increasing when there are more players; third, the game may have multiple equilibria with different firms entering even if, given a particular configuration of entrants, the price equilibrium is unique; and, fourth, the discreteness of entry decisions means that, even if a unique equilibrium is selected, the objective functions produced by a standard Nested Fixed Point estimation routine will be discontinuous. As a result, it is necessary to use a very large number of iterations, implying a a very great computational burden, in order to be confident that a global minimum has been found.

We address these practical problems by estimating the game using a method-of-moments estimator where we approximate the 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. The resulting objective function is smooth, which allows the use of standard minimization routines, and during estimation, we only need to re-evaluate and multiply a set of probability density functions to calculate the objective function rather than re-solving for equilibria in many different pricing games. While we focus most of our results on a model where service type 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.[2]

Before discussing the related literature, we identify several limitations of our analysis that the reader should be aware of up front. 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 the 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.[3] While these responses could have economically significant 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.[4]

Second, we do not model choices of route-level capacity or schedules. This is a significant limitation and it means that we must attribute some differences in carrier market shares to differences in unobserved quality or costs, which we assume that carriers cannot influence, rather than to strategic capacity or scheduling choices.[5] We hope to extend our model to allow for capacity or flight frequency choices, in order to analyze merger remedies that involve capacity or slot divestitures. A computationally-light approach to estimation is likely to be even more important for this type of extension.

Third, the approximation method that our estimation approach relies on will only be valid under some assumptions (see Geweke (1989) for discussion of when these assumptions fail). We can only test some of these assumptions, and even if all of the assumptions are valid, the use of approximation involves some sacrifice of econometric efficiency. The benefit is computationally feasibility.

Fourth, we examine a game where carriers that provide some service on a route decide whether to provide non-stop service, ignoring decisions about whether to provide any type of service, and we assume that carriers that were merely potential entrants can have no effects on prices. Sweeting, Roberts, and Gedge (2017) show that the presence of Southwest as a potential entrant caused incumbent carriers to reduce their prices. However, we documented these effects on a subset of routes with an existing dominant incumbent, which, given our focus on mergers, are not the routes that we focus on here. In an earlier version of this paper, we estimated a model where potential entrants (defined as carriers operating

[2]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 either be 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.

[3]For example, after the United-Continental merger the merged firm eliminated Continental's hub in Cleveland, which was followed by Frontier substantially expanding its Cleveland operations and Spirit and Allegiant Airlines entering the airport.

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

[5]For example, a carrier will be able to influence the appeal of its connecting service on a route by how it schedules incoming and outgoing flights at the connecting airport.

on any routes out of the endpoints) choose between providing non-stop, connecting or no service, i.e., a trinomial entry and service choice. This richer model had a greater computational burden, and the decision to provide connecting service turned out to appear to be fairly random. This finding may be explained by the fact that, when setting up the data, the definition of connecting and non-entrant carriers can depend on an arbitrary threshold for how many passengers a carrier has to have to be included as a competitor (see later discussion). As a result the estimates and the counterfactuals were harder to interpret than the binary non-stop service decision that we consider here.

The rest of the Introduction discusses the related literature. Section 2 describes the data that we use, emphasizing how we define non-stop service, how we define market size, which is novel to the literature, and how we allow for connecting traffic to affect the profitability of non-stop service. Section 3 describes our model, while Section 4 describes estimation and discusses identification, and, in particular, the role of exclusion restrictions in the various demand and cost equations of the model. Section 5 describes the parameter estimates both with and without a known order of entry, and 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.

Related Literature

Our paper is closely related to several distinct parts of the literature. Several papers estimate the price effects of consummated airline mergers, exploiting the abundance of publicly available data and the large number of mergers that have occurred since deregulation in 1978. Ashenfelter, Hosken, and Weinberg (2014) summarize this literature up to 2013. Prior to 1989, the Department of Transportation had primary jurisdiction over airline mergers and gave more credence than the Department of Justice to claims that the threat of new entry or service changes by existing competitors would prevent significant post-merger prices increases (Werden, Joskow, and Johnson (1991)). As a result, the Department of Transportation approved all proposed mergers. Most papers examining the pre-1989 period have identified significant post-merger price increases, although magnitudes vary depending on the choices of the time-window and the control group.[6] Kwoka and Shumilkina (2010) also identify, in the context of the USAir/Piedmont merger, price increases when one of the merging carriers serves the route and the other is a potential entrant. In a model with selection, reducing the number of potential entrants can also raise expected prices even if the number of actual competitors remains the same because it may mean that, on average, the set of entrants has higher marginal costs. Analyses of more recent Department of Justice-approved mergers have provided more mixed results. Hüschelrath and Müller (2014) and Hüschelrath and Müller (2015) identify short-run price increases of as much

[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, 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.

as 10% after recent mergers, with smaller long-run price increases, suggesting that efficiencies are realized or competition re-asserts itself. Israel, Keating, Rubinfeld, and Willig (2013) and Carlton, Israel, MacSwain, and Orlov (2016) argue that some of these recent mergers, such as Delta/Northwest and United/Continental, have benefited consumers through improved network benefits. Ciliberto and Williams (2014) and Azar, Schmalz, and Tecu (2016) argue that industry pricing behavior became more collusive in the late 2000s. One implication of this allegation is that comparisons of price changes between merging and non-merging carriers may not be so compelling in this era. We estimate our model using data from 2006 partly to minimize concerns that carriers are colluding, as our model assumes non-cooperative behavior.

A second related literature is work on entry games that developed following the seminal contributions of Bresnahan and Reiss (*inter alia* Bresnahan and Reiss (1991), Bresnahan and Reiss (1990)) and Berry (1992b), who considered airline markets, with subsequent important contributions by Mazzeo (2002), Seim (2006) and Ciliberto and Tamer (2009). These papers model firms' payoffs from entry or service-type decisions using a reduced-form profit function where competition effects enter in an additively separable way. A limitation for policy analysis is that there is no direct link between the estimated profit function and prices or consumer surplus.

A subsequent strand of the literature has tried to integrate models of entry/product positioning and models of post-entry competition, which requires facing the computational challenges laid out above. One approach, adopted, for example, by Draganska, Mazzeo, and Seim (2009), Eizenberg (2014), Wollmann (2016) and Fan and Yang (2016), assumes that any unobserved demand or marginal cost shocks are not known to firms when they take their entry/product-choice decisions.[7] This eliminates selection on unobservables and allows marginal cost and demand functions to be estimated separately from the entry game, but it means that some firms may regret their choices ex-post which is unsatisfactory if the data is to be interpreted as reflecting an industry in steady-state equilibrium. This type of assumption does not seem reasonable in the airline industry where carriers are likely to know more than a researcher about their competitiveness in the market. If so, selection will be present. 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.[8]

The most closely related papers are Reiss and Spiller (1989) and the working paper by Ciliberto, Murry, and Tamer (2016) (CMT, hereafter). Reiss and Spiller estimate a model of service choice and

[7] Related work, most notably Fan (2013), has examined how mergers may affect continuous measures of quality, as well as price. The 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 assumption that unobservables will remain the same when qualities change may be more reasonable.

[8] In the empirical auctions literature use has also been made of incomplete information entry models where either potential bidders themselves have no information about their values before deciding to enter (e.g., Athey, Levin, and Seira (2011)) or bidders know their values exactly before deciding to enter, even though they have no information on their competitors (Li and Zhang (2015)). Li and Zheng (2012) compare variants of these models, and Gentry and Li (2014) demonstrate non-parametric identification of selective entry auction models under a variety of assumptions.

subsequent price competition in airline markets, and they distinguish between non-stop 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 non-stop 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 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 overcome by considering at most six players and by using a simulated annealing minimization algorithm on a supercomputer. Our approach has a significantly lower computational burden, so that researchers should find it easier to apply, and to experiment with different specifications, in a range of different contexts. That said, our method will have lower econometric efficiency. Substantively, CMT, following Ciliberto and Tamer (2009) and Berry (1992b), focus on the decision of carriers to enter a market, without making a distinction between non-stop and connecting service, whereas we focus on the decision to offer non-stop service. We choose our focus for two reasons. First, the analysis of airline mergers usually makes a clear distinction between non-stop and connecting service, on the basis that non-stop service is a distinct, higher quality product for many consumers (Dunn (2008), Berry and Jia (2010)). Given our interest in merger counterfactuals, allowing for this difference is important. Second, if potential entrants are defined as carriers serving both endpoints, the fixed costs involved in providing connecting service may be quite small, and the data, where many connecting carriers have very small market shares, is consistent with this hypothesis. It is therefore more natural to construct the discrete choice game around the decision to provide non-stop service on a route, as non-stop service will require a costly commitment of aircraft and gate resources.

2 Data and Empirical Setting

We estimate our model using publicly available data from the second quarter of 2006. While more recent data is available, we use older data because we believe that our static model should be viewed as a representation of an industry that is roughly in steady-state with firms behaving non-cooperatively. Subsequent years were associated with, first, the significant, if transitory, effects of the financial crisis, which lowered demand for business air travel, and then, as the economy recovered, a number of significant mergers, and allegations of significant cooperative behavior among carriers resulting from either multi-market contact or common ownership.[10]

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

[10] Of course, the industry was being affected by some significant changes in 2006, such as the America West-US Airways merger, which was financially completed in 2005, but the carriers only obtained joint operating certificates in 2007, and

Our data is drawn from the Department of Transportation's DB1 10% ticket sample and its T100 Origin and Destination database. While these data are widely used, there is some useful information, such as the exact flights chosen and when tickets are purchased relative to the departure date, that are not available. This data would be required if choices of flight frequencies or revenue management practices were to be incorporated into the model.

Selection of markets. Our estimation sample consists of 2,028 airport-pair markets.[11] We create this sample by first identifying the largest 79 airports in the lower 48 states based on enplanements, and considering the routes between these airports. From this set of routes we exclude:

- airport-pairs that are less than 350 miles apart as ground transportation may be very competitive on these routes, which drops a large number of routes on the east coast, and some routes in California and Texas;
- airport-pairs involving Dallas Love Field, which was subject to Wright Amendment restrictions that severely limited which routes carriers could fly non-stop;
- airport-pairs involving New York LaGuardia or Washington Reagan that would violate the so-called perimeter restrictions that were in effect from these airports[12];
- airport-pairs where more than one carrier that is included in our composite "Other Legacy" or "Other LCC" (low-cost) carriers (defined below) has a large market share, as our assumption that the composite player will act as a single player will be especially problematic.[13]
- airport-pairs where, based on our market size definition (explained below), the combined market shares of the carriers is more than 85% or less than 4%.

Definition of players, non-stop and connecting service. As described in the Introduction, we focus on the decision of carriers to provide non-stop 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[14], open-jaw journeys or journeys involving more than one connection in either direction. Our next 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).

the closure of Independence Air which took place in Q1 2006. In our analysis we treat US Airways and America West as a single carrier.

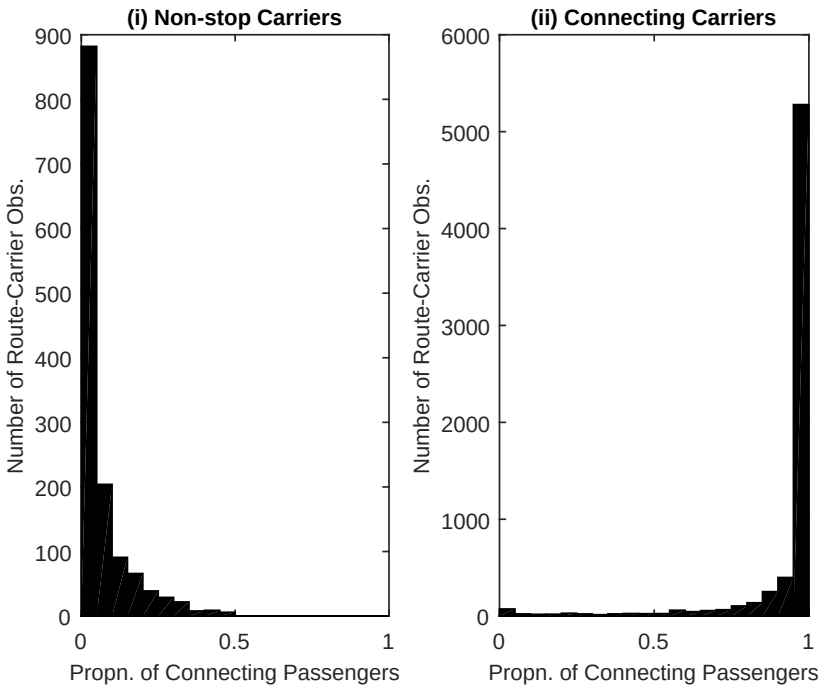
[11] It is often debated whether airport-pair or city-pair markets are more appropriate, and the treatment in merger cases has depended on exactly which airports and cities are being considered.

[12] 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.

[13] We exclude routes where two of these carriers are non-stop or have more than 20% shares of non-directional traffic; or where two of these carriers have presence of more than 25% at either of the endpoints (see our definition of shares, non-stop entry and presence below). 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 non-stop.

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

Figure 1: Proportion of DB1 Passengers Traveling with Connections, Based on the Classification of the Carrier



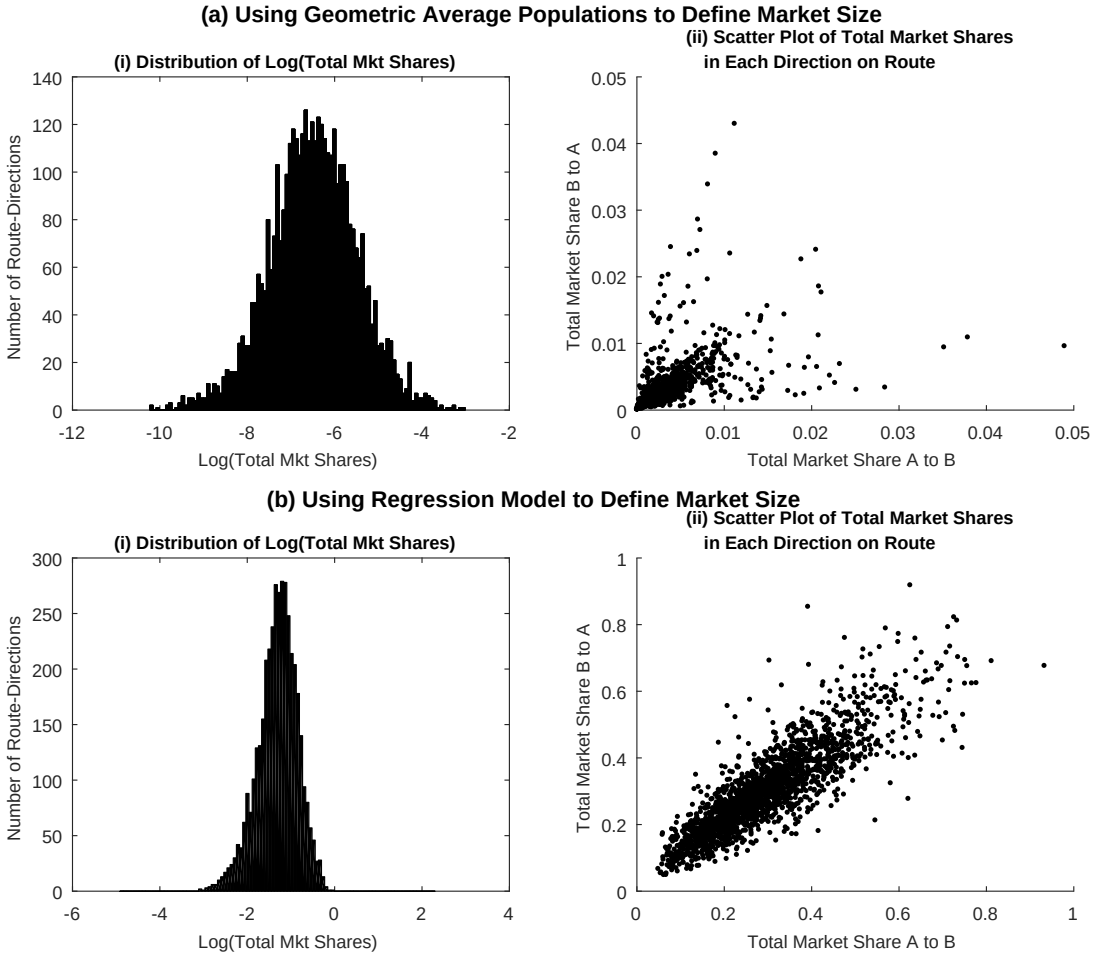
Based on the number of passengers carried, the largest Other Legacy carrier is Alaska Airlines, and the largest Other LCC carriers are JetBlue and AirTran.

We define the set of players on a given route as those carriers who achieve at least a 1% share (i.e., 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 non-stop service on a route if it is recorded in the T100 data as having at least 64 non-stop 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% non-stop thresholds for non-stop service have little effect because almost all non-stop carriers far exceed these thresholds. For example, Figure 1 shows that the carriers we define as non-stop typically carry only a small proportion of connecting passengers. For this reason, we feel able to ignore the fact that carriers may provide both non-stop and connecting service on the same route. On the other hand, our share threshold for which carriers should be counted as players does affect how many players are in each market. 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 from 2.9 to 2.4 as marginal carriers are excluded.

Our model assumes that there is no potential competition effect on market outcomes from carriers

Figure 2: Market Size Measures and their Impact on Market Shares



that we do not count as serving the route at all, even if they are active at the endpoint airports. Some indirect support for this assumption comes from the fact that the set of carriers providing non-stop service is relatively stable: of the 1,172 carriers that we define as providing non-stop service in Q2 2006, 1,027 of them were providing non-stop service on the same route in Q2 2005, 81 were providing connecting service and only 26 of them were present at the endpoints but not serving the market at all (given our definitions). By dropping these carriers we reduce the computational burden, while still capturing most of the carriers that are most likely to begin providing non-stop service if a merger takes place.

Market Size. As in many settings, the definition of market size is not straightforward. Ideally, we would like variation in market shares across markets to reflect variation in carrier characteristics and the type of service that carriers offer 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)).[15] However, as illustrated in the left-hand panel of Figure 2(a), using this measure results in considerable heterogeneity in 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, because these numbers vary across directions when the geometric average population cannot, even though the services offered by the carriers are usually very similar (right-hand panel). This is a problem as we will be modeling 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 unobserved random effect in demand, common to all carriers and types of service, that 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 non-stop 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%. As a result, once we employ the market selection rules laid out above, the distribution of total market shares and the correlation between total market shares in each direction display more limited heterogeneity, as illustrated in Figure 2(b).

Prices and Market Shares. To model competition between carriers, we define carrier prices and carrier 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 these prices and market shares (defined as the number of originating passengers carried divided by market size) to be different in each direction.[16]

Explanatory Variables. We construct a number of variables that we include in the demand and/or cost equations of our model. We define American, Continental, Delta, Northwest, United, US Airways and the composite Other Legacy carrier as legacy carriers. Carrier presence at an airport is defined by the number of routes that a carrier (or its regional affiliates) serve non-stop out of an airport divided by the total number of routes served by all carriers out of the airport, and we distinguish between presence at the origin and the destination of a directional route. Non-stop distance is defined as the great circle (haversine) distance, in miles, of a return trip. We define Reagan National, LaGuardia and JFK as slot-constrained airports. As business travelers are likely to have different price sensitivities, we use a measure of the proportion of business travelers that flying the route based on data provided to us by

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

[16] 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 significantly different. This may also explain some differences between average non-stop and connecting fares. Fares on contracts that carriers negotiate with the US government and large employers, which may be significantly below list prices, may also play a role, but there is no data available on these prices. 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.

Severin Borenstein (Borenstein (2010)). This data is based on surveys that measure the proportion of travelers flying for business at the state-level, so it is likely to be an imperfect proxy at the route-level. 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. On many medium-sized routes non-stop 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, non-stop, a significant number of non-Canadian/Mexican international destinations. Table 1 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 non-stop 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 non-stop service only when the carrier can serve unusually large amounts of connecting traffic. The model, and the exclusion restrictions, are detailed in Appendix A. We acknowledge that this model is not completely consistent with the structural model that we are estimating, as it does not account for how a carrier's network is composed of routes that are served as part of the outcome of a multi-carrier game, but we have found that including this variable can help to explain patterns of entry in the data. We view it as an attempt to approximate the types of models that carriers may use to predict flows of connecting passengers.

Summary Statistics. Table 2 contains route-level and route-carrier-level summary statistics for the primary variables in our data. On average, there are four carrier players in each market. There are more carriers, on average, in longer markets where a greater number of connecting routes may be attractive. For example, Seattle to Baltimore and Seattle to Orlando, which have one non-stop carrier, both have the maximum nine players. While most routes have no non-stop service, larger markets or markets that include the hub of multiple carriers have as many as four carriers providing non-stop service. As an illustration of how market structure varies, Figure 3 shows the number of carriers serving markets non-stop between ten airports serving cities of different sizes and with varying hub status. Most non-stop service involves a hub airport, illustrated by the fact that Salt Lake City, which is a hub for Delta, has more non-stop 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 non-stop service to hubs. Markets between larger cities and shorter routes are also more likely to have non-stop service.

Table 1: Domestic and International Hubs for Each Carrier

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

Figure 3: Number of Carriers Offering Non-Stop Service Between Selected Airports

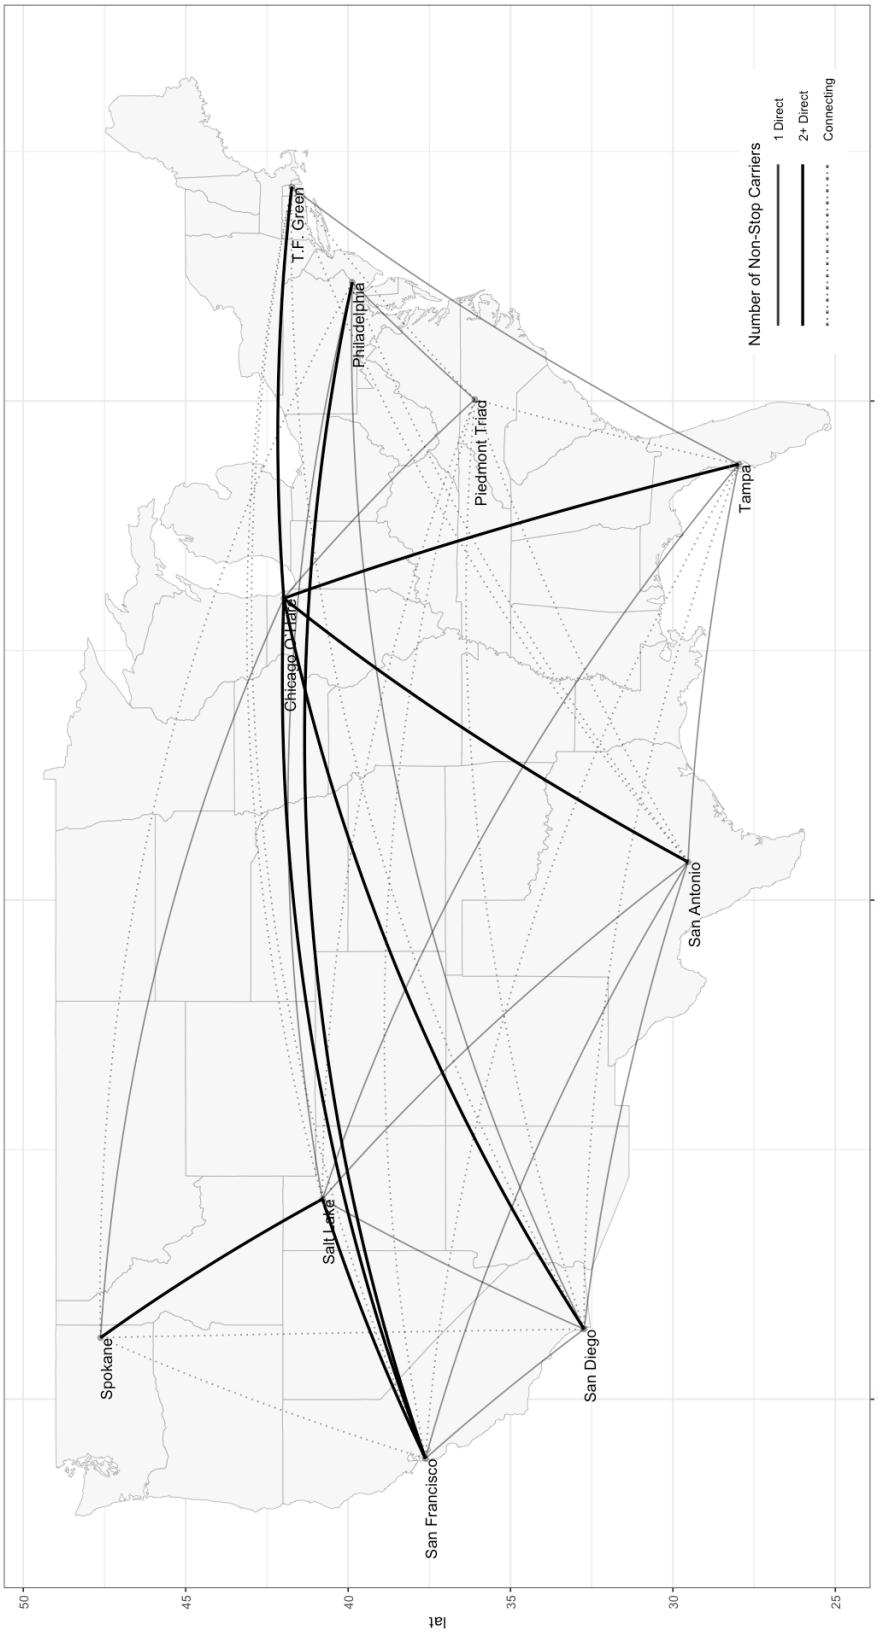


Table 2: Summary Statistics for the Estimation Sample

	Obs.	Mean	Std. Dev.	10th percentile	90th percentile
Market Level					
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. non-stop	2,028	0.668	0.827	0	2
Total Passengers (directional)	4,056	6970.90	10830.06	625	17,545
Non-stop 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 Level					
Non-stop	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
At Least One Endpoint is A Domestic Hub	8,065	0.13	0.33	0	1
At Least One Endpoint is an International Hub	8,065	0.10	0.30	0	1
Connecting Distance	7,270	3,161	1,370	1,486	4,996
Log(Predicted Connecting Traffic)	1,036	6.44	0.81	5.31	7.47

Our model allows for carriers to set different prices and have different market shares in both directions, reflecting the fact that market shares vary systematically with a carrier's presence at the originating endpoint. A simple regression shows that a one standard deviation increase in the difference in presence at the endpoint airports increases the expected difference in market shares by 1.3 percentage points, which is large given that average market shares are 7.1%.[17] Fares are also significantly higher for passengers originating from the higher presence endpoint, although the magnitude of the difference is quite small ($5). Fares vary systematically with distance (an increase in the return distance of 1,000 miles increases fares on average by $30), whether service is non-stop (non-stop service fares are $43 higher than connecting fares), whether the carrier is low-cost (low-cost carriers fares are $70 lower than legacy fares) and the degree of competition, and especially the number of non-stop carriers (controlling for route distance and the identity of the carrier, the presence of one non-stop carrier is associated with connecting fares falling by $10, while a second non-stop carrier is associated with a $40 reduction in non-stop fares and a $30 reduction in connecting fares. This pattern motivates our focus on what determines the number of carriers providing non-stop service in equilibrium.[18]

2.1 Changes in Service Choices After Actual Mergers

Given our focus on service choices, it is natural to ask what types of service choice changes are observed after actual mergers. To do this, we have examined the set of routes where in the quarter that a merger closed financially, the merging parties were both providing non-stop service and no other carriers were doing so, as these are the routes where both intuition and our estimates suggest that there may be the largest anti-competitive effects. For this exercise we define non-stop service using only T100 and treat a route as being served by a carrier non-stop when either the carrier itself or one of its regional affiliates flies at least 130 flights (in either direction) in each quarter. We examine the Delta/Northwest (5 routes), United/Continental (7), Airtran/Southwest (3) and American/US Airways (5) mergers.

For all of these routes, the merged firm continued providing non-stop service for the four years after the merger was financially completed. After one, two, three and four years the number of routes where at least one other carrier had initiated non-stop service were three, five, seven and eight respectively, out of twenty routes, so 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. This is consistent with a model where the carriers not currently providing non-stop service are a selected set of carriers who do not find this type of service profitable, and we will see below that we can only match this rate of non-stop initiation in our counterfactuals when we allow for selection on both observables and unobservables.[19]

[17] This estimate is taken from a route-fixed effects regression where the difference in market shares in each direction is regressed on the difference in presence. The standard deviation of the difference in endpoint presence is 0.33.

[18] The distance, non-stop service and competition estimates come from regressions of a carrier's weighted (across directions) average fare on a route on non-stop distance, carrier dummies, a dummy for whether the carrier provides non-stop service and interactions between whether a carrier provides non-stop service and the number of non-stop 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.

[19] Two of the routes where entry did occur within two years included routes from Newark or Washington National where entry was limited by slot constraints that were partially eased at a condition of the merger being approved. For example,

3 Model

Consistent with the majority of the airline literature we focus on carriers' strategic decisions at the route level (see Mazur (2016), who analyzes carriers' total capacity choices for an exception). Consider a particular route-market, m, connecting two airports A and B. There are a well-defined set of competing carriers, $i = 1, ..., I_m$, who already serve both endpoint airports. The carriers play a two-stage, complete information game. In the first stage they decide whether to enter with non-stop or with only 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. Non-stop service implies paying a fixed cost, F_{im}, whereas we assume that there is no fixed cost associated with providing connecting service given the fact that the carrier is already providing services to both endpoints. Our model does not allow for the possibility that a carrier provides both non-stop and connecting service on the same route, motivated by the fact that when non-stop service is offered almost all passengers travel non-stop (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.

3.1 Second Stage: Post-Entry Price Competition

We assume that, having chosen their service types, the carriers in a market play two standard, 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 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

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

where

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

and

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

Southwest was able to enter Newark, and the Denver-Newark route that had been a United and Continental non-stop duopoly, through an approved purchase of slots from United/Continental. Southwest also began service from Washington National to Nashville after the American/US Airways merger where slot divestitures at National were required. 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.

so that quality can depend on observed characteristics, such as the type of carrier (legacy vs. LCC) and route characteristics. TRN denotes a truncated normal and the lower truncation of β_{im}^{NS} at zero implies that the perceived quality of non-stop 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 vary across markets in a parametric way

$$\alpha_m \sim N(X^\alpha \beta_\alpha, \sigma_\alpha^2)$$

where X^α will include the business index for the route, and

$$\tau_m \sim N(\beta_\tau, \sigma_\tau^2)$$

although we assume that α_m and τ_m are the same across directions for the same route.[20]

ν_m is a market-specific random effect that is designed to capture the fact that in some markets all carriers tend to have more travelers in both directions relative to our chosen definition of market size than could be rationalized by independent heterogeneity across carriers. We assume that it 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 providing service. We abstract away from capacity constraints and capacity choices, although we admit that these choices may significantly affect competition in practice. Specifically we assume that

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

where the cost can depend on the type of carrier (e.g., low-cost or legacy), the type of service and the distance traveled through the parameters β_{MC}. For non-stop service the distance is simply the non-stop 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.[21] As we assume that travelers are making return trips we treat the marginal cost as non-directional.

This specification is restrictive in two senses. First, we are assuming that the random component of marginal costs does not vary with the service type 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 that our model cannot necessarily explain every realization of market shares and prices in the data. We have adopted these restrictions based on the fact that the models we have estimated with independent cost shocks across either directions or service type choices have fit the data less well (for example, implying more variability in prices and market shares across directions than is

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

[21] For the composite Other Legacy and Other Low Cost carriers it is not straightforward to assign connecting routes. Therefore we use the non-stop distance for these carriers, but include additional dummies in the connecting marginal cost specification for these carriers.

actually observed).[22]

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 a vector of variable profits, $\pi_m^{A \to B}(s)$, for each carrier as a function of service-type choices (or market structure), 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.

3.2 First Stage: Service Type Selection

In the first stage of the game carriers choose whether to commit to the fixed cost associated with non-stop 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 this gives us an extensive form game where the payoff of a carrier i at the endpoint of the game tree where the vector of service types is s is

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

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

$$F_{im} \sim TRN(X_{im}^F \beta^F, \sigma^F, 0, \infty).$$

We assume that all of the market and carrier-level demand and cost draws are known, by all carriers, when choices about whether to offer non-stop service 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. Berry (1992a) has previously estimated a model of sequential entry for airline markets, assuming that profitability and incumbency affect the order.

We also consider the robustness of our results to more agnostic assumptions about the order or timing of service choices, as we recognize that the baseline assumption of sequential choices is strong. Specifically, we allow for the equilibrium that is played to be any of the pure strategy Nash equilibria in the simultaneous service choice game[23] *or* a subgame perfect Nash equilibrium in a sequential move game with any order of moves. We show that the assumed order has little effects on the estimates.

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 the first-stage of the model is by backwards induction where one solves for equilibrium variable profits in all possible market structures (adding up across the two directions on the route) in a first-step and then solves the extensive form game

[22] One natural solution would be to allow for the costs shocks to be correlated and to estimate the degree of correlation by introducing additional parameters. In our experience allowing for additional covariance parameters can lead to the objective function have more local minima.

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

backwards from the end of the game tree in the second-step. However, this approach is computationally inefficient because of the time spent solving for variable profits in market structures that may be strictly dominated for one or more firms. Therefore, when we are assuming a sequential structure, we instead use a procedure where we selectively grow the game tree *forward*.

To be precise, we first calculate whether it would be profitable for the first firm (i.e., the one with the highest presence) to operate as a non-stop carrier if it is the only carrier in the market, given its F.[24] If not, then we do not even need to consider any of the branches where it provides non-stop 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 firm, and ask, for each of the branches still under consideration (i.e., holding fixed the choice of the first carrier) whether the firm would find it profitable to provide non-stop service if no subsequent firms were in the market. If it is not, then we only need to consider future growth where the second firm is providing connecting service. Otherwise we can continue forward with both non-stop and connecting branches for the second firm. This process continues until we have considered all of the firms, and we can then solve the game backwards, as usual, on the reduced game tree that results.

In the game that we consider the gains from this type of selective growth are useful but not absolutely necessary because our estimation routine is itself fairly efficient, and, indeed, we need to use a more standard approach when we consider the simultaneous move game. However, if we consider larger games (e.g., three options per firm and nine firms) then this type of approach appears absolutely necessary for estimation to be feasible.

4 Estimation and Identification

One approach to estimation is to use a nested fixed point procedure where, every time one of the parameters is changed, the model is re-solved to find equilibrium outcomes and to calculate an objective function that reflects the match between the observed outcomes and the data. We view this approach as not being feasible without access to massive computational resources when there are up to nine competitors per market, directional demand and pricing, and the demand and cost specifications which contain more than a few parameters.

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 the researchers. In the second stage 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 reweighting the outcomes from the games solved in the first step.

The main benefit of this approach is that we avoid re-solving a large number of games each time we change the value of one or more of the parameters; instead, we just need to calculate and to multiply the probability density functions. The first step can easily spread across a number of machines as each

[24] To be clear, here we are testing whether the profits from providing non-stop 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.

game is solved independently. An additional benefit is that we can estimate variants of our model with different linear combinations of variables (e.g., those affecting the mean of marginal costs), without having to repeat the first step.[25] In this section we detail 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 a Monte Carlo experiment that examines how well the procedure works and shows how it can be extended to models with multiple equilibria (e.g., a simultaneous entry model), or where the order of entry is unknown to the researcher.

4.1 Importance Sampling

As outlined by Ackerberg (2009), who proposed its use for a range of models used in empirical Industrial Organization, the potential advantage of importance sampling is that it provides a way to quickly approximate the value of an objective function when it is relatively costly to solve the economic model.

In our setting, suppose that we want to calculate the expected value, $E_m(y)$ of a particular outcome, y, such as the probability that American provides non-stop service in a particular market, m, such as Boston-Raleigh.

Denote a given realization of the quality and cost draws for each carrier as θ_m, and the parameters that describe the distribution of these draws, and which we want to estimate, as Γ. Assuming that the density of these draws is characterized by $f(\theta_m|X_m, \Gamma)$

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

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, and it is too time consuming to use a standard simulation approach where $y(\theta_m, X_m)$ is recalculated for different θs every time we adjust Γ. However, 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.[26] For a given set

[25] For example, suppose in the simplest specification we assume that the mean of carrier fixed costs associated with non-stop only depends on whether the carrier is a legacy or low-cost carrier. To estimate this model we have to solve games where we take a large number of draws of fixed costs for each carrier. We can use these same draws to estimate a model where we also allow the mean of fixed costs to depend on our measure of potential connecting traffic and whether the carrier has an international hub, which makes specification testing particularly straightforward.

[26] 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.

of S draws from g we can then approximate $E_m(y)$ using

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

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.[27] A significant advantage of this re-weighting procedure is that $\widehat{E_m(y|\Gamma)}$ will be a smooth function of Γ even when the outcome of interest is discrete (such as an entry or service type decision).

A disadvantage of this procedure is that estimates of the moments will be inaccurate if the density g is too dissimilar to f, intuitively because expected outcomes will have been calculated using the average of some outcomes that are not really relevant. Connected to this issue, Geweke (1989) shows that an importance sampling estimate of an expected value will only be consistent if the variance of $y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$ is finite. We provide evidence that the variances are finite given that the number of draws that we use in both our application (Section 5) and the Monte Carlo (Appendix B).

4.2 Method of Moments

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

$$m(\Gamma)'W^{-1}m(\Gamma)$$

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 3 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-type choices. For example, market-specific outcomes include weighted average connecting and non-stop 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 non-stop 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).

We construct a diagonal weighting matrix so that there is equal weight on the three groups of moments, and, within each group, the value of the weight is proportional to the reciprocal of the variance of the moment based on some initial estimation described below.[28] While this approach should still give consistent results, it is a little different to the usual asymptotically optimal GMM procedure where a consistent estimator of the inverse of the full covariance matrix of the moments is used. There are

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

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

Table 3: 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 (non-stop) 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 non-stop, 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 non-stop price (both directions), connecting price (both directions), HHI (both directions), and the square of number of non-stop carriers}
y_C = {non-stop dummy, price (both directions), and market shares (both directions)} for each carrier

Table 4: Description of g For the Final Round of Estimation

Market Draw	Symbol	Support	g
Market Random Effect	v_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	α_m	[-0.75,-0.15]	$N(X_m^{\alpha}\beta_{\alpha}, 0.022^2)$
Carrier Draw			
Carrier Connecting Quality	$\beta_{im}^{CON,A\to B}$	[-2,10]	$N(X_{im}^{CON}\beta_{CON}, 0.219^2)$
Carrier Incremental Non-stop Quality	β_{im}^{DIR}	[0,5]	$N(X_{im}^{DIR}\beta_{DIR}, 0.257^2)$
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_{DIR}.constant= 0.374 (0.500), β_{MC}.legacy= 1.802 (1.600), β_{MC}.LCC= 1.408 (1.400), β_{MC}.non-stop_distance= 0.533 (0.600), β_{MC}.non-stop_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, non-stop quality premium 0.5, marginal cost 0.15, fixed cost 0.25.

two reasons for our choice. First, with a large number of moments relative to the number of markets, we cannot claim to be able to estimate the full variance-covariance matrix consistently. Second, we found that using the full covariance matrix gave worse numerical performance and less stable coefficient estimates, which may be explained by the lack of consistency. Our approach ensures that equal weight is given to pricing, market share and entry outcomes in estimation.

4.3 Supports and Starting Values

Our method requires us to specify the support of each of the θ draws (these supports cannot depend on the parameters) and to choose the importance density g. To generate the results reported in this paper, we use the supports and densities listed in Table 4, where the densities are truncated at the ends of the supports. The supports were chosen to be broad in the sense that they contained all of the values that were likely to be relevant, with the exception of the nesting parameter whose support was restricted because we found that, during estimation, there could be local minima which implied implausibly high or low values of τ. The assumed range of τ is consistent with many 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 accounted for). Draws from the gs are taken independently for each market, carrier (except for the market-level random demand effect) and type of draw.

To get to the parameters that we use to form the gs, 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. While the estimator can be consistent for any set of gs where the expectations have finite moments, Ackerberg (2009) recommends using a multi-round estimation procedure to improve efficiency.[29]

We take 2,000 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.

4.4 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. Focusing on the demand equation (e.g., the aggregate data estimating equation for a nested logit model taken from Berry (1994)) in their model, where each carrier makes a binary choice over whether to enter or not enter, the problem is that the unobserved product qualities for carriers that are observed in the market will be correlated with observed characteristics and, more importantly, will not have unconditional mean zero. This implies that the use of instruments for the observed characteristics will not be sufficient for consistent estimation. In our model, we are looking at a service choice within active carriers, but we still face the same problem because we are allowing for an incremental quality of non-stop service, so that the value of the quality unobservable for the type of service that is actually offered will not have mean zero. This problem can be corrected only by understanding the form of selection that emerges from the first-stage game.

This having been said, the intuition for how identification works will come from the fact that in our three groups of structural equations (demand, marginal costs and fixed costs of non-stop service) we have made a number of exclusion restrictions with respect to the included exogenous variables. For example, we have assumed that carrier endpoint presence affects the preferences of travelers originating at that endpoint only, without having any direct effect on marginal costs or fixed costs, although we do, for example, allow the hub status of an airport to affect fixed costs. Route distance, which can vary across routes and across carriers depending on the location of their domestic hubs, is assumed to directly affect marginal costs, but not demand (given that we controlled for its effect on market size using our gravity model) or fixed costs. Domestic and international connecting traffic (i.e., network structure) and slot constraints are assumed to affect the effective fixed cost of serving a route, but not

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

the demand or marginal costs associated with passengers who only want to fly the segment. Of course, these characteristics for other carriers will also create variation in a carrier's entry decisions through the structure of the service choice game. Our continuous measure of the amount of connecting traffic generated by serving a route may be especially valuable because when this variable is large non-stop service may be essentially certain, so that the incremental draw of non-stop quality for this carrier-route should be effectively uncensored. In our baseline model, the assumption of a particular order of moves will also provide identifying power, although, as we show below, our estimates are robust to using more agnostic assumptions.

Of course, our model also has many parametric restrictions on the distribution of the residuals, including the fact that we are restricting them to be independent of each other (although the market random effect allows for correlation in demand across carriers and across service types in a given market). In principle one could allow correlation across the different types of residuals for a given carrier (e.g., allowing higher quality to be associated with higher marginal costs), which is something that CMT do allow for in their model, although without the market-level random effects which are a feature of our model. However, we have chosen to not go down this route for two reasons. First, our included covariates are able to explain most of the variation in which carriers provide non-stop service, which limits the role that we need unobservables to play. Second, and more practically, we have found that the objective function tends to have more local minima when we allow for correlations, which reduces the confidence we would have in the results.

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 5. 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 sets of 2,000 draws for each market) for each of the drawn markets. We discuss the estimates themselves before describing the fit of the model and the performance of the estimation algorithm.

Demand Parameters. The first finding is that, as we suggested in the data section, there is unobserved heterogeneity in demand for air travel across markets, reflected in the estimated standard deviation for the market random effect. In contrast, there is very limited unobserved heterogeneity in the nesting and demand slope parameters across markets. All else equal, demand on business routes is less elastic (the business index parameter for the slope). This is consistent with the results from the richer demand models estimated by Berry and Jia (2010) and Ciliberto and Williams (2014), where business travelers have been identified as having less elastic demand. The expected slope coefficient for

Table 5: Coefficient Estimates (bootstrapped standard errors in parentheses)

				(1) Assumed Order of Entry		(2) No Eqm. Selection
Demand: Market Parameters						
Market 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
Market 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 Non-stop Service			Distance	-0.025	(0.034)	-0.041
			Business Index	0.247	(0.494)	0.227
	Std. Dev.	σ_{DIR}	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 Legacy"	-0.270	(0.680)	0.197
			Conn. X "Other LCC"	0.124	(0.156)	0.164
			Non-Stop Distance	0.579	(0.117)	0.589
			Non-Stop Distance^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 Constrained Airport	0.568	(0.094)	0.602
	Std. Dev.	σ_F	Constant	0.215	(0.035)	0.198
			Function value	1.870		
			Run Time	29 CPU-hours		47 CPU-hours

the market with the highest business index (Dayton to Dallas-Fort Worth) is -0.34 compared to the cross-market average of -0.57.

Given our point estimates, the average (absolute value) own-price elasticity of demand is 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 number of people choosing to fly on any carrier when the prices of all carriers are increased) is 1.29. This number is consistent with the existing literature: for example, Gillen, Morrison, and Stewart (2003) report a median elasticity of 1.33 aggregating 85 airline demand studies, and Berry and Jia (2010) estimate a value of 1.67 when they also use data from 2006 but apply a more disaggregated model where carriers sell tickets at a large number of different prices.

The remaining market 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 non-stop 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 slightly 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. Our experience also suggests that the marginal cost parameters are the least sensitive parameters to changes in the specification. The coefficients indicate that legacy carriers have higher marginal costs for both non-stop and connecting service, and that marginal costs of both non-stop and connecting service increase with distance in a reasonably similar way. For a legacy carrier, the average marginal cost of providing non-stop service on a roughly 3,000 mile round-trip route (Miami to Minneapolis) is $345, compared to $367 for connecting service. For Southwest the marginal costs of non-stop and connecting service are lower and, for this route, almost identical ($303 and $298 respectively). Naturally our estimates of marginal costs are somewhat higher than in studies that have estimated carrier demand to be less elastic, but how costs vary with distance is quite similar to previous studies, even when different types of service have not been distinguished. The degree of unobserved heterogeneity in marginal costs is relatively small (estimated standard deviation of the marginal cost draw is $16), which leads to our finding that there is only very limited selection on unobserved marginal costs.

Fixed Cost Parameters. The estimated average fixed cost for providing non-stop service between two airports that are neither hubs not slot-constrained is around $900,000 for both types of carrier, although the cost paid by carriers that choose non-stop service is only $480,000. In our static model we are obviously unable to distinguish between costs that are incurred every period that the carrier provides non-stop service and costs that are only incurred when non-stop service is initiated. We estimate higher costs for routes out of slot-constrained airports, reflecting the opportunity costs involved in using a slot on a specific route. The remaining parameters also have the expected sign and allow us to explain why carriers serve most routes non-stop from their domestic and international hubs. The connecting traffic variables indicate that, as expected, a carrier is more likely to serve a route when doing so will deliver more connecting traffic for other destinations. For some routes to and from domestic hubs this

effect will make the expected fixed cost very small for the hub carrier, making non-stop service almost inevitable.[30] Estimated unobserved heterogeneity in fixed costs is relatively small, reflecting the fact that observables are able to explain which carriers provide non-stop 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, so it is important to address this issue. For example, Table VII of Berry (1992a) reports quite different estimates when changing the assumed order of entry from being based only on profitability to one based on previous incumbency, and Ciliberto and Tamer (2009) find that there are multiple equilibria in over 90% of market simulations given their estimates. Column (2) of Table 5 reports the point estimates when we minimize an objective function based on moment inequalities, allowing for the observed outcome to be 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 is implemented by collecting together all of the possible equilibrium outcomes of the game, for a given set of draws θ_{ms}, calculating the maximum and minimum values of each predicted outcome. We can then use importance sampling, as before, to calculate the expected values of the maximum and minimum outcomes for a given set of parameters Γ. We can then generate two sets of moments. For the upper bound moments we replace $\left(y_m^{data} - \widehat{E_m(y|\Gamma)}\right)$ from our baseline specification with $\max\left(0, \left(y_m^{data} - \widehat{E_m(y^{max}}|\Gamma)\right)\right)$, and a lower bound moment $\max\left(0, \left(\widehat{E_m(y^{min}}|\Gamma) - y_m^{data}\right)\right)$.

Our analysis indicates that objective function has a unique minimum. The value of the objective function for this estimator is 77.59. We can also evaluate the objective function of this new 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 an inequality estimator.[31] This is not a straightforward task given the fact we have a large number of moments and the need to search a 32-dimensional space hundreds of times in order to construct critical values using a method such as the one proposed by Chernozhukov, Hong, and Tamer (2007), which, in any case, assumes a small number of moments compared to the number of market observations. This assumption does not hold here. We note however, that when we apply the tests proposed 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 estimates in column

[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 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] 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 identified from outcomes that are predicted uniquely, so the baseline order may simply serve as an overidentifying restriction that can be rejected.

Table 6: 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	Number of Markets	Average Number of Equilibria
1	141	1
2	304	1.005
3	416	1.014
4	413	1.019
5	342	1.025
6	228	1.022
7	136	1.028
8	46	1.037
9	2	1.042
Total	2,028	1.017

(1) than in column (2).[32]

More importantly, however, 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), where an order of entry was imposed. The similarity partly reflects the fact that, given the estimates in column (1), multiple equilibria are rare even with weaker assumptions. Table 6 shows the average number of distinct equilibrium outcomes when we solve for all of the pure strategy Nash equilibria and all of the subgame perfect Nash equilibria in sequential move games 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 it is less than 1.05 even for the markets with the most carriers. 98.4% of market 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 illustrate this claim.

First, note that in Table 5, the reported time taken to estimate the parameters in the second step of the routine is under 30 hours.[33] Optimization is performed on a single processor without using parallelization or even analytic derivatives, although for specifications where we have provided

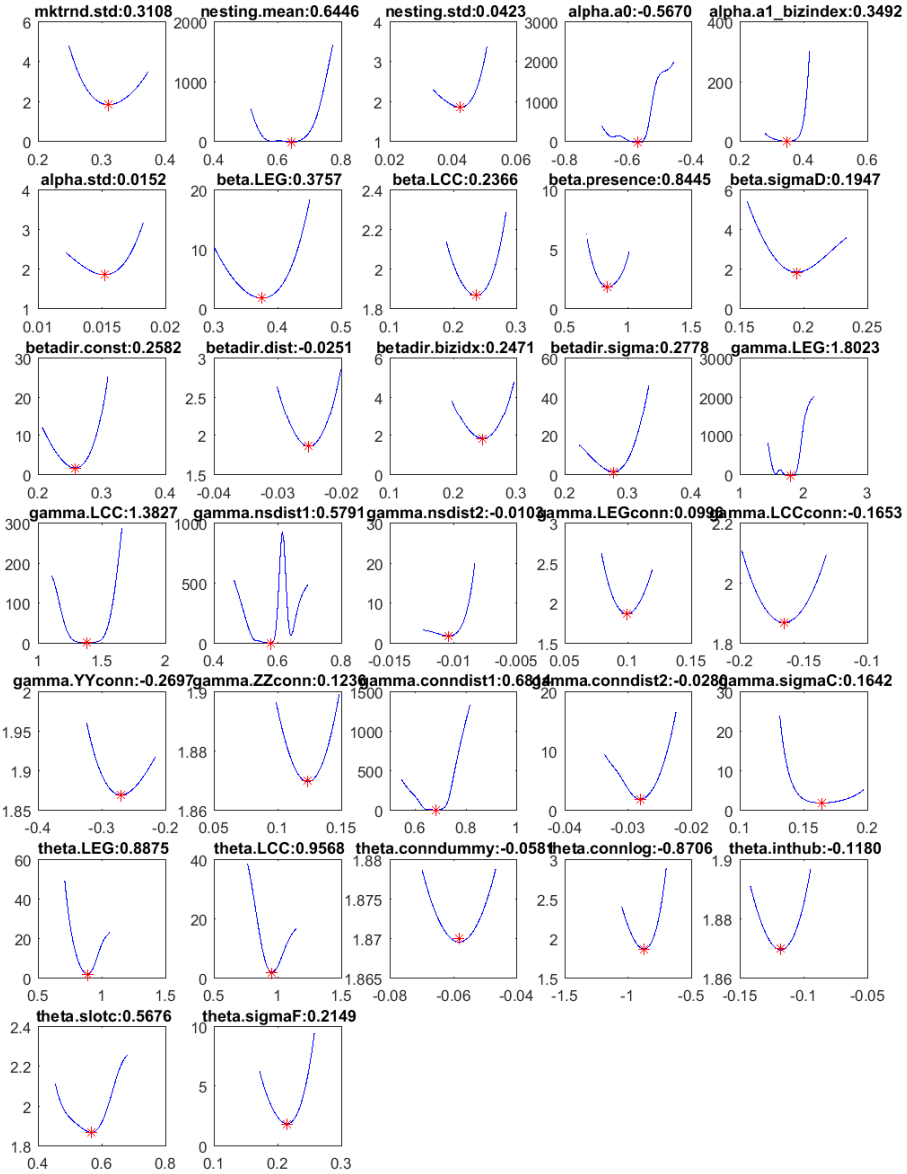
[32]The CCK approach is potentially very useful in our application because it is valid when the number of moments is large and the calculation of the critical values does not require the minimization of the objective function many times. However, the approach cannot easily be adjusted to account for the fact that some of the moments will appear to be 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. 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.

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

derivatives the estimation time is significantly lower. Therefore the resources required to perform the estimation, and to investigate alternative specifications, should be available to a wide range of academics and practitioners.

Second, Figure 4 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. Interestingly, the shape of the objective function using the real data appears to be even more regular than in our Monte Carlo experiments (see Appendix B).

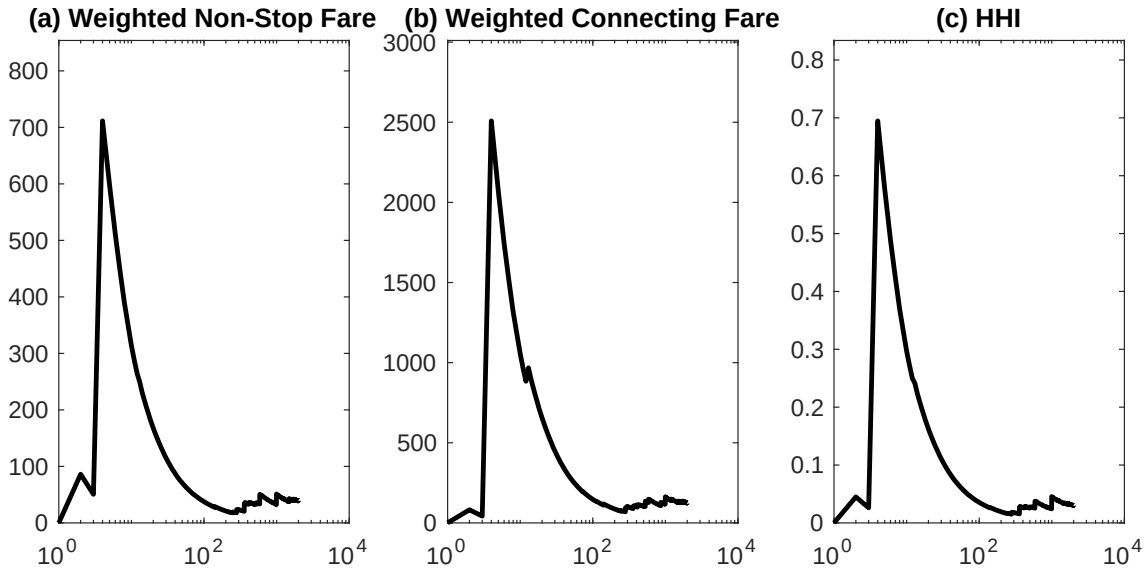
Figure 4: Shape of the Objective Function Around the Estimated Parameters



Third, we address the question of whether our importance sampling estimator satisfies the condition

that the variance of $y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$ should 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 quite wildly as S rises. Figure 5 shows these recursive estimates of the sample variance for the moments associated with the three market-level outcomes, namely the weighted non-stop fare, the weighted connecting fare and the quantity-based HHI 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 seem quite small for $S > 500$. In our application we are using $S = 1,000$.

Figure 5: Sample Variance of Three Moments as the Number of Simulation Draws is Increased (logarithm of the number of draws on the x-axis)



5.4 Model Fit

We now discuss the fit of the model, both on average and for some specific airports. Table 7 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 compute the predicted fit by taking the model estimates and simulating 20 new sets of parameter draws for each market in the data from their implied distributions (i.e., we do not use the same draws that we use in estimation, which would tend to give us a better fit than the new draws that we use). The standard errors in parentheses reflects the standard deviation in the reported means when we repeat the calculations using the estimates from each of our bootstrap samples.

On average, we match the differences in market shares and prices across different types of service

Table 7: Model Fit: Average Market Shares and Prices (bootstrapped standard errors in parentheses)

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

accurately, although the predicted level of prices and market shares are both slightly higher than in the data. These cross-carrier averages hide some patterns at the carrier-level. For example, in the data United's average non-stop and connecting fares are \$479 and \$436 directly, while our model predicts \$472 and \$445, so the match is close. On the other hand, Delta's average non-stop fare (\$498) is greater than its average connecting fare (\$448), whereas, our model predicts a connecting fare of \$466 and a non-stop fare of only \$453.

The estimated model also predicts service choices accurately. Using our 20 simulated outcomes per market, our success rate at predicting whether a carrier actually chooses to enter non-stop is 87.5% (1.1%). This involves correctly predicting 91.7% (1.0%) of decisions to provide connecting service and 67.1% (2.8%) of decisions to provide non-stop service. However, if, for a given carrier-route, 11 or more of our 20 simulations predict non-stop service, this is what we observe in the data for 82.6% (2.2%) of market-carrier observations. For the counterfactuals, the relationship between airport presence and carrier entry decisions is important. In the data, 44.9% of carriers with the highest mean presence on a route enter direct, compared with a prediction of 36.5% (2.0%). Table 8 shows the performance of our model at predicting direct entry at a number of domestic hubs/focus cities for the hub carrier. While the performance is a little uneven (e.g., we predict less non-stop service by Delta in Salt Lake City and Continental at Newark than they actually provide), the fit is generally impressive.

This is also true for airports that are not hubs. Take, for example, Raleigh-Durham International in North Carolina, which in 2006 was not a hub for any carrier (it had been a small, unprofitable American hub in the 1990s). Table 9 shows the percentage of routes out of Raleigh on which each carrier

Table 8: Model Fit: Proportion of Routes Served by Carriers from a Selection of Domestic Hubs

		Number of Routes	Propn. Mkts Entered Non-Stop Data	Prediction
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%)
Midway	Southwest	44	72.7%	64.5% (6.0%)

Table 9: Model Fit: Predictions of Service Decisions at Raleigh-Durham

	Number of Routes	Mean Presence Endpoints	% Non-Stop Data	% Direct Sim.
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%)

provided non-stop service both in the data and in simulations from our estimated model. The number of routes is different for each carrier depending on the other airports that they serve: for example, Delta (US Airways) serves many passengers and many routes from Raleigh by providing connecting service through Atlanta (Charlotte). The non-stop choices of the carriers providing the most non-stop service at Raleigh, American and Southwest, are predicted very accurately. The largest difference between the prediction and the data is for United, where most simulations predict that the carrier would provide non-stop service to its hubs in Denver and San Francisco. While United did not serve these routes non-stop in 2006, it has subsequently added them. Delta, whose service is also underpredicted, has also significantly increased the number of routes that it serves non-stop from Raleigh.

6 Evidence of Selection and Its Effects on Outcomes

We now use our estimated model to understand how much selection is implied by the parameters and the effects of selection on market outcomes. We will examine how selection affects merger counterfactuals in the next section.

Our first exercise examines which types of observed and unobserved carrier- and market-level factors explain which carriers provide direct service. We do this by solving for equilibrium outcomes for 20 sets of draws for each of the markets in our data, based on our point estimates. In Table 10, we report, at the carrier-level, the average values of each factor or draw for all observations and for those market-carrier observations where the carrier provides non-stop service in equilibrium. A difference between these mean values is evidence of selection on the observed or unobserved variable of interest. We do this for all markets, and for three groups of markets based on market size, as we might expect to observe stronger selection in smaller markets where more favorable draws may be required to make non-stop service profitable.

Consistent with non-stop service requiring significant fixed costs, carriers providing non-stop service tend to be found in larger markets, with a higher demand random effect and less elastic demand (i.e., a smaller, in absolute value, price coefficient). Non-stop carriers tend to be those with higher observed and unobserved (to the researcher) quality for both types of service (recall, that quality of non-stop service equals the quality of connecting service plus an incremental draw, with factors such as presence at the origin affecting both connecting and non-stop quality). Similarly, non-stop carriers tend to have lower marginal costs and lower fixed costs, although the difference in the unobserved component of marginal cost (assumed to be the same across connecting and non-stop service in out model) is small. Carriers providing non-stop service tend to move earlier in the assumed sequential order, although it is unclear whether this reflects first-mover advantage or the fact that we have assumed that carriers with higher presence, which also affects quality, move first.

To further understand which factors are most important in determining which carriers provide non-stop service, we estimate a number of simple linear regressions using the same factors, where the dependent variable is whether a market-carrier provides non-stop service in equilibrium. The continuous explanatory variables are normalized so that they have mean zero and standard deviation 1, so it is easier to make comparisons between the relative effects of variables and draws that are measured in

Table 10: Mean Value of Observed and Unobserved Factors for All Carriers and for Carriers Choosing Non-Stop Service

	Market Size Groups							
	1st Tercile		2nd Tercile		3rd Tercile		All Markets	
	All	Non-Stop	All	Non-Stop	All	Non-Stop	All	Non-Stop
Market Demand Parameters								
Market Size	5,032	6,324	14,294	15,630	56,596	70,343	25,284	56,331
Market Random Effect	0.004	0.080	0.005	0.059	0.000	0.030	0.003	0.038
Nesting Parameter	0.646	0.647	0.645	0.646	0.644	0.643	0.645	0.643
Price Coefficient	-0.426	-0.410	-0.426	-0.416	-0.429	-0.425	-0.427	-0.423
Components of Carrier Quality								
Observed Connecting Quality	0.512	0.827	0.520	0.771	0.551	0.669	0.528	0.697
Unobserved Connecting Quality	-0.001	0.061	-0.001	0.062	-0.002	0.042	-0.001	0.047
Observed Non-Stop Quality	0.804	1.148	0.811	1.088	0.846	0.977	0.820	1.007
Unobserved Non-Stop Quality	0.075	0.323	0.075	0.285	0.073	0.212	0.074	0.231
Components of Carrier Marginal Cost								
Observed Connecting MC ($)	325.1	293.7	328.4	290.0	314.1	292.4	322.5	291.9
Observed Non-Stop MC ($)	320.2	286.1	323.2	284.2	309.2	287.5	317.6	286.7
Unobserved MC ($)	-0.2	-2.9	-0.3	-0.2	-0.2	-2.4	-0.2	-2.6
Components of Fixed Cost								
Observed Fixed Cost ($m)	0.882	0.480	0.844	0.524	0.797	0.607	0.841	0.585
Unobserved Fixed Cost ($m)	0.002	-0.151	0.007	-0.037	0.028	0.033	0.012	0.012
Structure of Game								
Number of Carrier Players	4.106	3.466	5.121	3.990	4.996	4.296	4.741	4.204
Carrier's Position in Order	2.553	1.246	3.061	1.576	2.998	2.116	2.871	1.972

different units. Obviously our model implies a non-linear relationship between the factors and entry decisions, but we view the estimated coefficients in this linear form as still being informative about the relative importance of different parts of the model (the conclusions are essentially identical using probit or logit specifications). We are especially interested in the relative importance of observed and unobserved factors, as in most entry models in the literature, it is assumed that unobserved demand and quality shocks could only be observed after entry/product-type choices have been made.

The results of seven specifications are reported in Table 11. In the first specification, we only include market-level variables. As expected, higher demand and less elastic demand (due to a lower absolute price coefficient or nesting parameter) make non-stop service more likely. Of these factors, market size is clearly the most important, but all coefficients are significant at the 1% level.[34] In the second specification, we include the components of the carrier's own qualities and costs that are based on observable factors. These five variables increase the adjusted R^2 from 0.23 to 0.42, and they indicate that higher quality, lower non-stop marginal costs and, especially, lower fixed costs, are associated with a higher probability of non-stop service. Observed non-stop quality has a negative sign, which is unexpected, but may reflect how the same variables that affect incremental non-stop quality can also affect other variables in a non-linear way, as well as the fact that the coefficients on these variables in the structural estimates have quite large standard errors. The third specification adds variables measuring unobserved quality and cost components, which raises the adjusted R^2 by a further 0.1. Based on the magnitude of the coefficient, unobserved quality seems to have similar importance to observed quality, while the two unobserved cost components are less important than their observed equivalents. For example, a one-standard deviation increase in (un)observed fixed costs reduces the probability of non-stop service by 0.134 (0.048).

The fourth specification adds dummies for the carrier's position in the move order and the number of carriers that we are using as players. While carriers later in the move order are less likely to enter direct than the first mover, it is striking that these controls have almost no effect on the other coefficients and that the adjusted R^2 increases by less than 1 percentage point. This is consistent with our finding that, at the estimated parameters, there is almost always a single equilibrium outcome whatever move order is considered.

The remaining specifications include either market, market-simulation (recall we are doing twenty simulations per market) or market-carrier fixed effects. In these cases, the coefficients are identified from cross-carrier or cross-simulation variation within markets, so that coefficients that only vary at the market level (including potential passengers' valuation of non-stop over connecting service) are not identified. The small estimated standard deviation of the price coefficient leads to a very imprecise estimate of the price coefficient in these regressions when we include fixed effects but the other coefficients are almost unchanged. The robust patterns that emerge from these results are that observable variation in market demand, which reflects our definition of market size, marginal costs and fixed costs explain most of the variation in which carriers provide non-stop service, and that the type of unobservable that

[34]Note that here we are not correcting standard errors to allow for the fact that parameters are estimated. However, given the small size of the standard errors on the predictions made above there is no reason to believe that taking into account this effect would change the statistical significance of the coefficients.

Table 11: Determinants of Non-Stop 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 Effect	0.021***	0.021***	0.021***	0.021***	0.022***		0.022***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)		(0.001)
Nesting Parameter	-0.003***	-0.002***	-0.003***	-0.003***	-0.003***		-0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)		(0.001)
Price Coefficient	0.033***	0.023**	0.025***	0.042***	-645.709		-639.643
	(0.003)	(0.009)	(0.009)	(0.009)	(1,271.633)		(1,293.485)
Components of Carrier Quality							
Observed Connecting Quality		0.080***	0.078***	0.091***	0.019***	0.019***	
		(0.020)	(0.019)	(0.019)	(0.003)	(0.004)	
Unobserved Connecting Quality			-0.049***	-0.049***	-0.049***	-0.050***	-0.049***
			(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Observed Non-Stop Quality		-0.021	-0.018	-0.052***			
		(0.019)	(0.018)	(0.018)			
Unobserved Non-Stop Quality			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							
Observed Connecting MC		0.011	0.010	0.149***	0.135***	0.133***	
		(0.030)	(0.029)	(0.031)	(0.029)	(0.033)	
Observed Non-Stop MC		-0.086**	-0.089**	-0.238***	-0.220***	-0.218***	
		(0.042)	(0.041)	(0.042)	(0.040)	(0.045)	
Unobserved 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 Fixed Cost		-0.124***	-0.134***	-0.126***	-0.158***	-0.158***	
		(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	
Unobserved Fixed Cost			-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 Carrier	Market	Market-Simulation	Market-Carriers
Observations	161,300	161,300	161,300	161,300	161,300	161,300	161,300
Adjusted R-squared	0.230	0.427	0.521	0.528	0.588	0.550	0.632

matters most is route-carrier-level variation in the quality of non-stop service.

Our second exercise shows how our assumption of complete information affects the distribution of equilibrium outcomes relative to the assumption of incomplete information given the estimated parameters. We base our analysis on a single market, Austin-Los Angeles LAX, which has six carriers in our data (Southwest, American, United, Other Legacy, US and Continental with this order of entry) with two of them, Southwest and American, providing non-stop service. None of these carriers has overwhelming presence at either endpoint (the highest values are for United, which serves 39% of destinations out of LAX, and Southwest which serves 37% of destinations out of Austin), which allows for the possibility we can see a variety of service choices for different draws and as we vary market size, which we will do in our experiments. On this route the market sizes in both directions are similar (63,231 and 69,891), and for simplicity we first fix them both at the average (65,561) and then, in our experiments, vary market size upwards and downwards in steps of 5,000, between 16,561 and 306,561.[35]

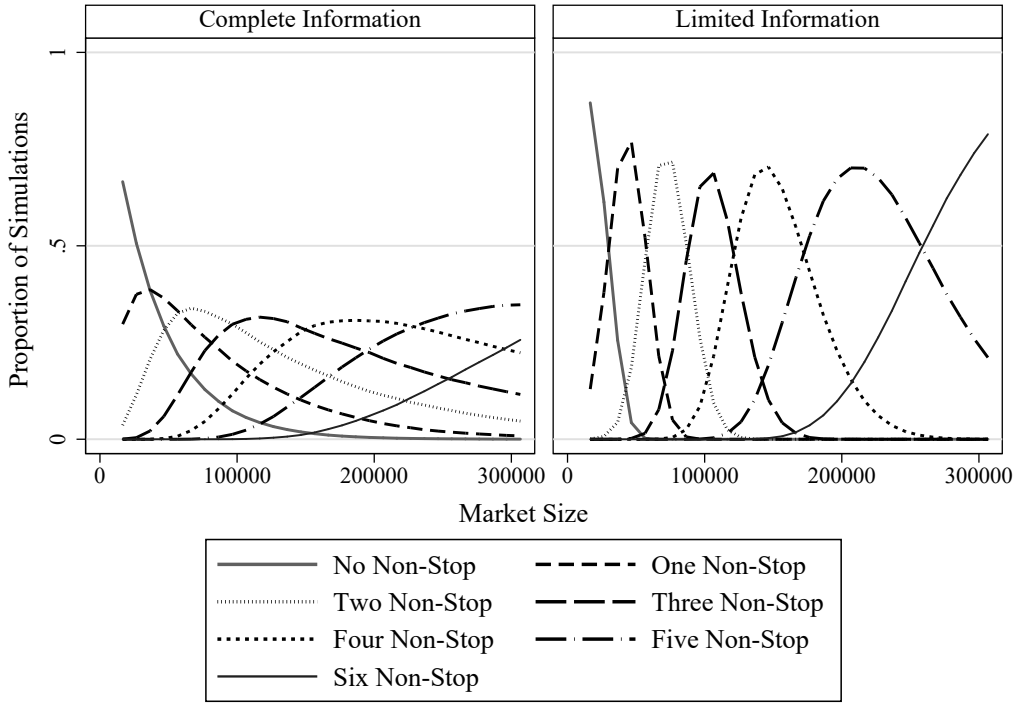
We solve for equilibrium outcomes under two different models, which we will call complete information and limited information, using 50,000 sets of the quality and cost draws for each market size in each case. The complete information model is the model with selection and a known order of entry that we have estimated. The limited information model assumes that carriers do not know the value of marginal cost or quality shocks when they take service decisions but that they do know all carriers' fixed costs of providing non-stop service. We continue to assume an order of entry to avoid multiplicity of equilibria. For each market configuration we approximate the expected profits of each carrier in every possible market configuration by taking 1,000 draws of marginal costs and qualities. We then solve the sequential, limited information entry game for each of the 50,000 draws of fixed costs before simulating realizations of the marginal cost and quality draws to compute expected market shares and consumer surplus.

The most striking differences in the predictions of the two models concern how the number and the identity of non-stop carriers varies with market size. This is partially illustrated in Figure 6 which shows the proportion of simulations resulting in different equilibrium numbers of non-stop carriers for the two models. The limited information model has a lower average number of non-stop carriers for the smallest market sizes, but a larger average number for market sizes above 46,561.[36] This result is intuitive as, when the market size is small, it may only be profitable to provide non-stop service when either the fixed cost is very low or the quality associated with non-stop service is very high, which may be unlikely even for a high presence carrier ex-ante. Therefore knowledge of the quality draw can make entry more likely to occur (i.e., profitable for higher fixed cost draws). On the other hand, when market sizes are large, non-stop service may only be unprofitable for high presence carriers when their quality draws are unfavorable, so that information about carrier quality can make non-stop service less likely. For example, for a market size of 106,561, where the modal number of non-stop carriers is 3 for both models, the probability that Southwest, the highest presence/first mover carrier, provides non-stop

[35] We 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.

[36] For example, for market sizes of 16,561, 66,561 and 106,561, the limited (complete) information models yield 0.13 (0.37), 1.15 (1.11) and 3.13 (2.41) non-stop carriers respectively.

Figure 6: The Relationship Between Market Size and Equilibrium Market Structure Under Different Informational Assumptions



service is 0.85 under limited information and only 0.51 under complete information.

The patterns shown in Figure 6 are also relevant for thinking about the effects of mergers on service type decisions. Under limited information, the number of non-stop carriers is much more clearly dependent on the size of the market, meaning that if one non-stop carrier was to be eliminated and we allowed carriers to re-optimize their service choices then it will be likely that at least one carrier will initiate non-stop service. On the other hand, with complete information, the quality and marginal cost draws of the connecting carriers will make it less likely that one of them will want to initiate non-stop service given that it was not profitable to do so before the merger. This finding will also show through in our merger counterfactuals below (where observable features of connecting carriers will tend to make it even less likely that they will want initiate non-stop service, compared with the Austin-LAX market where carriers are fairly symmetric).

We compute the consumer surplus for each simulation. Figure 7 shows how expected consumer surplus varies with market size for the two different models. Consumer surplus will tend to increase when there are more non-stop carriers (these are better products, and carriers cannot capture the full incremental value of non-stop service when they set prices), and especially when the non-stop carriers have higher quality and lower marginal costs. Even though the expected number of non-stop carriers is lower in the complete information model for market sizes above 46,561, consumer surplus is higher until market size exceeds 126,561 as the greater selection that occurs in the complete information model tends to increase consumer surplus.

Figure 7: The Relationship Between Market Size and Expected Consumer Surplus

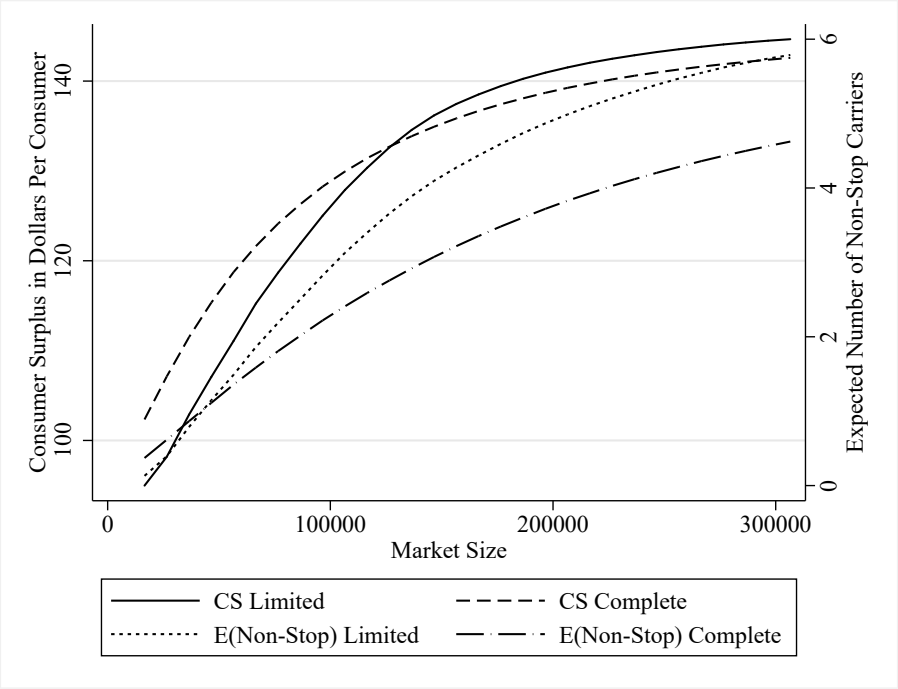
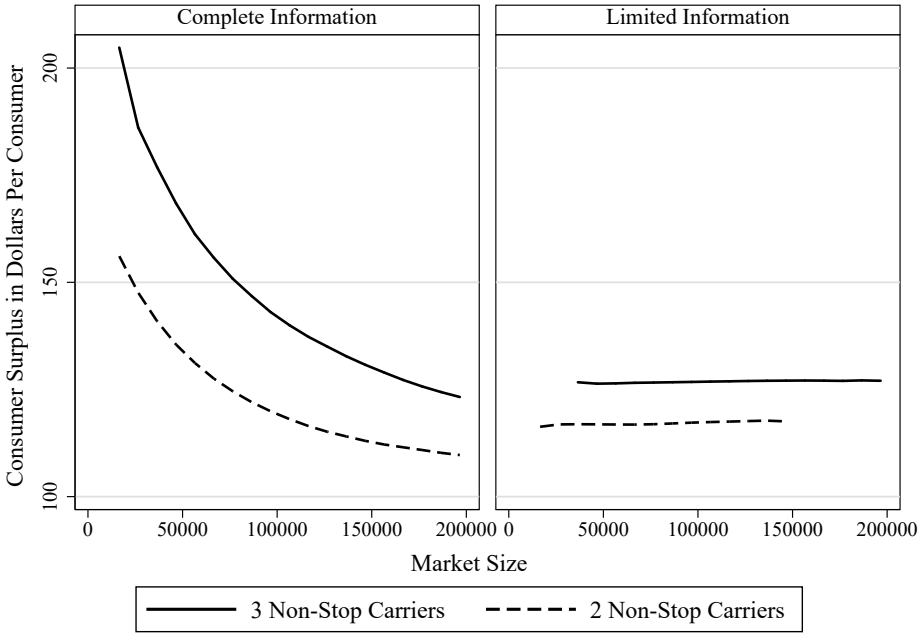


Figure 8: The Relationship Between Market Size and Expected Consumer Surplus Conditional on a Number of Non-Stop Carriers



This can also be seen in Figure 8 which shows expected consumer surplus conditional on a given number of non-stop carriers being an equilibrium outcome of the model (we only show market sizes up to 200,000 as outcomes with two or three non-stop carriers do not occur in our simulations for the limited information model when market sizes are large). Unless the configuration is only an equilibrium in the complete information model because carrier quality and marginal cost draws are likely to be very unfavorable (i.e., at large market sizes), selection tends to mean that consumer surplus is significantly higher in the complete information model because the carriers with the best non-stop qualities and marginal costs will be selected into non-stop service.

7 Merger Counterfactuals

We now examine a number of merger counterfactuals. We are interested in the predictions of the model, their policy implications, and the relationship between the predictions and different types of selection.

7.1 The United-US Airways Merger and Proposed Remedy

In May 2000 United and US Airways announced their intention to merge. The proposed merger was abandoned in July 2001 when it became clear that the Department of Justice would oppose it. While these events occurred several years before our data, this is an interesting hypothetical merger to look at because the parties proposed a remedy in which a third carrier, American, would commit to provide non-stop service on several routes where the merger would have created a non-stop monopoly based on pre-merger service choices.[37] In particular, there was a concern that while American might be willing to serve these routes because of other benefits that it might get as part of the transaction (in particular, it was buying many United assets on the East Coast and would be a co-owner of a new carrier, DC Air, that would operate out of Washington DC), it might not be an effective competitor in the sense of constraining the merged firm's pricing on the routes that it pledged to serve.

Our starting point is to examine what we predict would happen to prices and market shares on different types of routes, if we hold carriers' service type choices fixed as would be done in a typical merger simulation. Of course, we are doing this using 2006 as our base year, not 2000, which was when the merger was proposed, but as we are only interested in understanding the effects of this type of merger and this type of remedy in general, this is not a significant issue.

Figure 9 shows the routes where, in Q2 2006, both US Airways and United were non-stop competitors. Most of these routes connect a United (domestic or international) hub and a US Airways hub. There are seven of these routes that are in our sample and where the merging parties were the

[37] "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 non-stop 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."

Table 12: Predicted Effects of United/US Airways Merger with Service Types Held Fixed

Pre-Merger Market Structure	No. of Markets	Mean Pre-Merger UA/US Price	Mean Post-Merger "New UA" Price
UA & US Non-stop Duopoly	7	\$479	\$549 (+14.6%)
UA & US Non-stop, With Other Non-stop Rivals	10	\$350	\$378 (+8.0%)
One Party Non-stop, Other Party Connecting	163	\$443	\$463 (+4.5%)
Both Parties Connecting	521	\$484	\$480 (-0.9%)

only non-stop carriers. There are an additional ten routes where the merging parties were non-stop competitors but another carrier provided non-stop service as well. The red lines mark the five non-stop duopoly routes where American agreed to provide non-stop service as part of the proposed remedy. One of these markets (Pittsburgh to Washington National) is too short to be in our sample, but the other four markets were non-stop United and US Airways duopolies during our data and American provided connecting service.[38] Figure 10 shows a randomly selected sample of 16 of the 163 routes in our sample where one of the merging parties provided non-stop service and the other party provided connecting service. Many of these routes involve a link between two (domestic or international) hubs owned by the same party, which provides non-stop service while the other party does not, such as Chicago O'Hare-Denver (both United hubs), or a link between a hub and a non-hub city (e.g., San Francisco-Seattle). We will be examining the effects of the merger on these routes later in the paper.

To perform our analysis, we back out carrier qualities and costs using our point estimates, observed prices and market shares, and then simulate new prices when we eliminate the merging carriers, replacing them with a "new United" carrier which has the quality and marginal costs of the merging party with the highest presence on the route.[39] We assume that it has the service type of the higher market share party pre-merger, which in practice means that the new United is non-stop if one of the merging parties was non-stop.

Table 12 reports the pre-merger weighted average price of the merging firms and the new United's predicted post-merger price based on the markets in our data for four groups of markets where both carriers provided service in Q2 2006, which vary in the type of service that they offered and in the case of the routes which both merging parties served non-stop, whether another carrier also provided non-stop service. The results are intuitive. When the merging firms were non-stop duopolists, our average prediction, is that the new United's prices would rise significantly, by around 15% on average,

[38]The agreement required American to provide non-stop 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 non-stop service on the San Jose route prior to the merger.

[39]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 non-stop 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 are holding service types fixed, we do not need to distinguish between carrier-service type quality and the random market-level demand shock.

Figure 9: Routes Where United and US Airways Both Provide Non-Stop Service

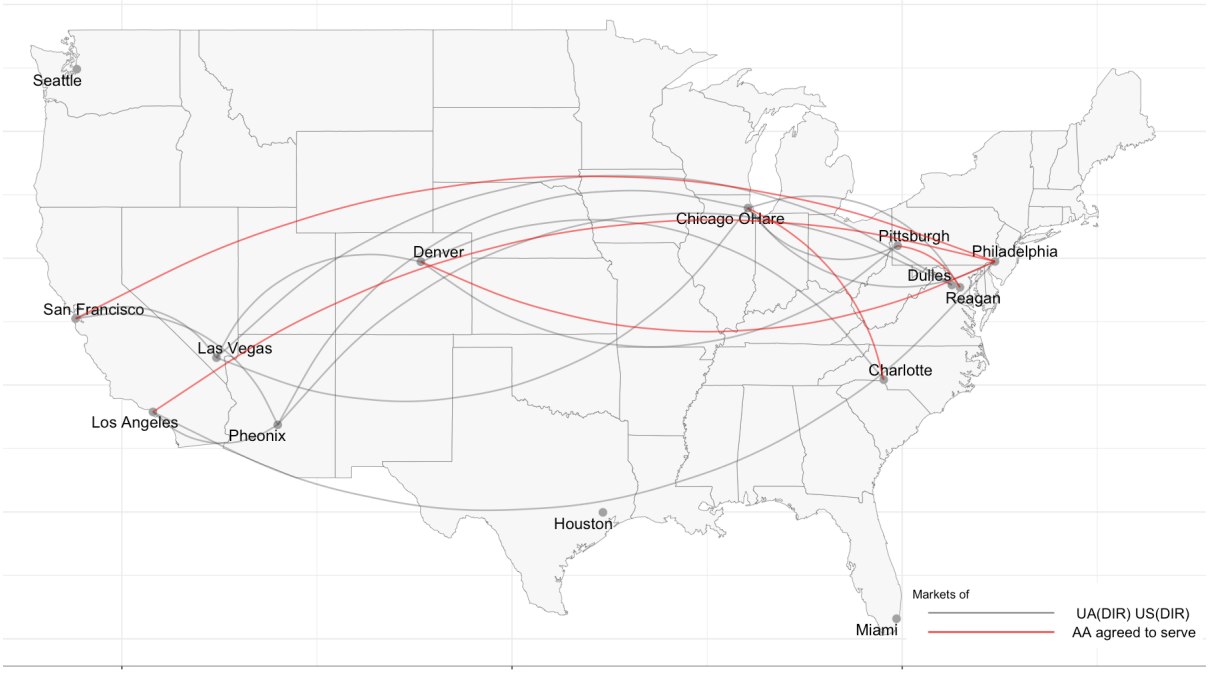
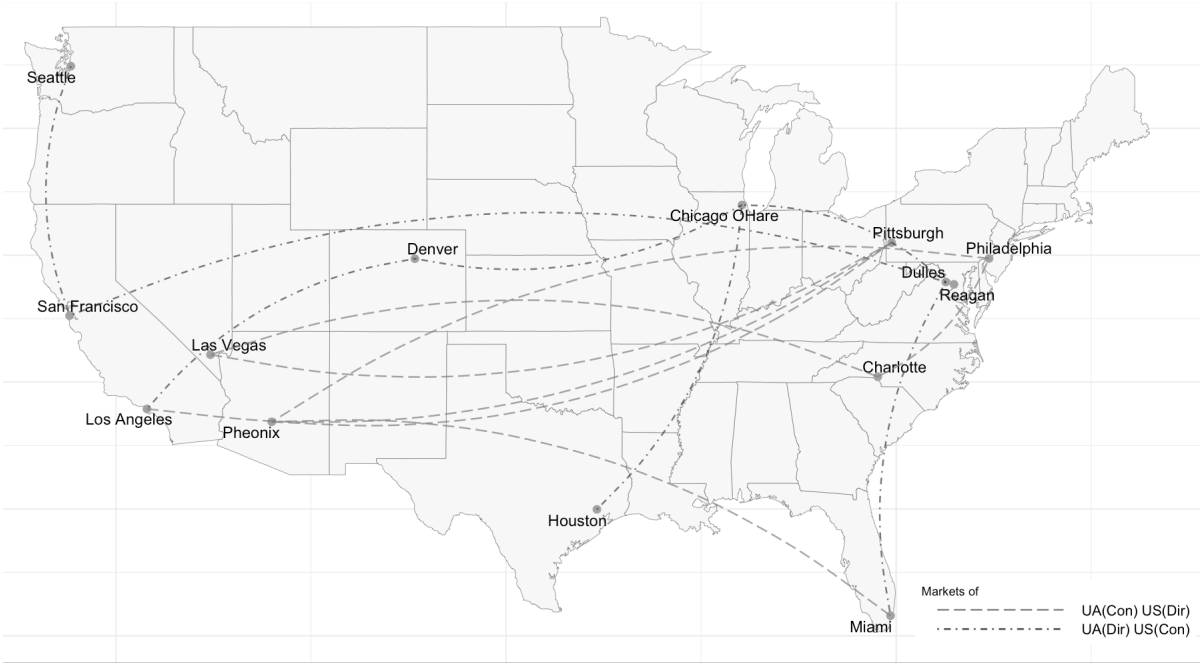


Figure 10: A Random Sample of the Routes Where One of United and US Airways Provide Non-Stop Service and the Other Provides Connecting Service



despite the fact that there are an average of 2.84 connecting carriers in each market.[40] In contrast, when another carrier also provides non-stop service, the average predicted increase is smaller at around 8%. The difference illustrates the potential importance of distinguishing between non-stop and connecting competitors, as well as whether the merging parties themselves provide non-stop service, in merger analysis.

The predicted price increase is smaller (4%) when one of the parties is connecting and the other is non-stop (in some of these markets another non-stop competitor is also present). When both parties offer connecting service, we predict that, on average, the price paid by someone traveling on the merging carriers will decrease reflecting the fact that the merger may have limited effects on competition (on average 0.4 other carriers are providing non-stop service and 3.2 carriers are providing connecting service in these markets; there are only seven markets where the merging parties are the only carriers in the market) and the higher presence carrier may have a lower price than the other merging party before the merger.[41]

The main value of our model is that we are able to perform counterfactuals where we look at whether a change in competition will lead to carriers changing the type of service that they provide, when we make different assumptions about selection. We start by considering the four sample markets that would have been affected by the proposed remedy. The first row of Table 13 shows that prices are predicted to increase by an average of \$46 (8.6%) in these markets if there are no service changes.

The remaining rows allow for endogenous service choices. To explain the different scenarios, it is useful to begin with the second row where we allow for both "observable and unobserved selection". In this case we take draws of the market demand random effect[42], and we use this to infer the quality and marginal cost of each carrier given the service choice that it actually chooses. For fixed costs and the service types that are not chosen, we then take draws from the estimated distributions of the incremental quality of non-stop service and fixed costs, and solve the game to check whether the observed choices will be made. We repeat this process until we have 100 sets of draws that rationalize observed choices. By construction, these draws will also predict the prices and market shares for each carrier observed in the data.

Next, we perform the merger counterfactual for each of the sets of draws that we have collected, making the same assumptions about the 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 non-stop service, as we do not observe merging parties ceasing non-stop service on these types of routes, although this is also what it would choose to do in almost all of our simulations even when we ignore selection. We predict that the prices of the merged firm would increase by \$32 dollars on average. Across the market-simulations, the probability that American initiates non-stop service is 0.12 and the expected number of other rival carriers initiating non-stop service is 0.28. Therefore while post-merger

[40] On average consumer surplus falls by \$19 per potential consumer in these markets.

[41] Note the elimination of a product and the fact that the price of the higher presence merging party will rise after the merger means that the merger is predicted to reduce average consumer surplus (by \$5 per consumer) for these routes even though the price reported in the table is falling.

[42] Given their small variances, we do not take draws of the price coefficient and the nesting parameter and instead hold them fixed at their mean values.

Table 13: Predicted Effects of United/US Airways Merger in Markets Where the Merging Carriers Were Non-stop Duopolists Prior to the Proposed Merger

Selection/Service Change Considered	Mean Pre-Merger Average United/US Airways Price	Expected Number of Rivals Initiating Non-Stop Service Post-Merger		Expected Post-Merger "New United" Price
		American	Other Rivals	
1. No service change	$532	-	-	$578 (+8.6%)
Allow Service Change				
2. Observable & Unobservable Selection	$532	0.117	0.283	$564 (+6.0%)
3. Observable Selection	$532	0.243	0.618	$549 (+3.2%)
4. No Selection (others as good as weaker merging non-stop)	$532	0.400	0.868	$500 (-6.1%)
Remedy: American Forced In Non-stop				
5. Observable & Unobservable Selection	$532	1	0.1	$556 (+4.5%)
6. Observable Selection	$532	1	0.530	$542 (+1.9%)
7. No Selection	$532	1	0.700	$498 (-6.4%)

service changes do provide some restriction on the market power of the merged firm, relative to the fixed service case, the merged firm's prices increase by more than 6% on average.

Now consider the third row of Table 13. For this counterfactual we use the same random effect draws that we ended up selecting for our counterfactuals in the second row (this makes comparisons easier), so that we infer the same carrier qualities and marginal costs for the service types that are actually chosen. For the service types that are not chosen we take a random draw of incremental qualities and fixed costs given the observed characteristics of each carrier, and then perform the merger counterfactual without taking into account whether the draws imply a different pre-merger market structure from the one that is actually observed. Compared with the second row we predict that American and other rival carriers are more likely to initiate non-stop service post-merger and a much smaller post-merger price increase (around 3.2% on average). The smaller average price increase reflects not only that the changes in service type choices but also because, in this counterfactual, the carriers who initiate non-stop service will tend to have higher qualities than in the counterfactual in the second row.

In the fourth row, we completely ignore selection by assuming that the carriers that currently provide connecting service would be more similar to the merging parties if they initiated non-stop service than their observable characteristics would imply. In this sense we are ignoring both observable and unobservable forms of selection. We implement this using the same random effect draws that are used in the second and third rows, but assuming that carriers that provide connecting service would have the (imputed) quality and marginal cost as the merging party with the lower average endpoint presence if they initiated non-stop service, and would draw their fixed cost from the same distribution as this party as well. If they provide connecting service their quality would be the same as observed in the data. In this case we estimate that American initiates non-stop service with probability 0.4, and that the expected number of other rivals initiating service is 0.87. Therefore, on average, the merger does not reduce the number of non-stop competitors, and it is also possible that the new service will be attractive to customers. These factors cause the new United to set a lower price than the observed average pre-merger price of United and US Airways (by 6.1%), and predicted average consumer surplus also increases (by around $10 per consumer).

The final three rows of Table 13 we consider the effects of the proposed remedy where American committed to initiate non-stop service. The three rows replicate the analysis in rows 2 to 4, except that now American is forced to provide non-stop service whatever its profitability. However, this does not mean that it is an effective non-stop competitor, which is illustrated by the fact that the effect that American's increased non-stop service has on the equilibrium choices of other rival carriers is quite modest. For example, with both types of selection (row 5), the expected number of rivals initiating non-stop service is 0.1 compared with 0.28 when the probability of American's entry was 90% lower (row 2). We see the same pattern independent of what is assumed about selection. The fact that American is typically ineffective as a competitor when it only provides non-stop service because of the remedy also means that the new United's predicted prices under the remedy are only slightly lower than without the remedy. For example, the new United's predicted price is 4.5% above the pre-merger price with the remedy, compared to 6.0% without the remedy, when we allow for selection.

7.2 The Predicted Effects of Later Mergers Observed in the Data

We have repeated the analysis of predicting post-merger prices and changes in service choices for Delta/Northwest, United/Continental and American/US Airways mergers which were completed in the eight years after our data. As our results are based on market structures in Q2 2006, and we consider the effects of each merger separately rather than their combined effect, we view them primarily as an investigation about whether the patterns that we observe in our results for the UA/US merger are broadly typical of our model's predictions for what would happen after the merger of two legacy carriers.[43]

Table 14 shows the predicted price changes for the different types of routes where we hold service types fixed. For the United/Continental and American/US Airways mergers the largest predicted increases are for routes where the merging parties are non-stop duopolists, which is the same pattern predicted for the United/US Airways merger. The pattern is different for Delta/Northwest, where the predicted price increase between a non-stop and connecting carrier is large. The different predictions reflect different patterns in market shares before the merger. For example, on average in a market where either American or US Airways is connecting and the other is non-stop, these carriers have a combined market share of 9.5% and other carriers have a combined market share of 12%. In contrast, in a market where either Delta or Northwest is connecting and the other is non-stop, 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 in these types of markets. None of the mergers are predicted to lead to large average price increases on routes where both carriers provide connecting service.

Table 15 repeats the analysis of service changes and price changes after mergers in markets where the merging parties are non-stop duopolists under different assumptions about selection. 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 non-stop service. While magnitudes vary, the basic pattern is the same as in Table 13: when we ignore selection, prices are predicted to fall as it is likely that carriers providing connecting service prior to the merger will find it profitable to initiate quite competitive non-stop service after the merger; whereas, when we allow for selection, we predict less initiation of non-stop service by rivals and larger price increases. Interestingly, the predicted rate of new non-stop service when we allow for the most general form of selection is quite similar to the proportion of routes where we observe new non-stop service two to three years after these mergers were completed after our sample period: in Section 2.1, we described how 25% and 35% respectively of duopoly routes experienced new entry within two and three years respectively.

Table 16 considers the effect of different selection assumptions on our merger counterfactuals in markets where both the merging parties and at least one other carrier provide non-stop service prior to the merger.[44] We include United/US Airways in the table as we did not consider the effect of selection on merger counterfactuals for these routes in our earlier analysis. While we assume that the merging

[43] We do not present results for the Airtran/Southwest merger as Airtran is part of the amalgamated other low cost category.

[44] In practice, for all but one of these routes there is exactly one non-stop rival.

Table 14: Price Effects of Mergers with Service Type Choices Held Fixed (with number of routes, and percentage changes)

	Delta/Northwest		United/Continental		American/US Airways		Average	
Merging parties are	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Non-stop duopolists	\$566.39	\$593.20	\$503.75	\$556.17	\$459.13	\$521.15	\$482.25	\$537.86
	2 routes	+4.7%	4 routes	+10.4%	11 routes	+13.5%		+11.5%
Non-stop with non-stop	\$351.26	\$382.04	\$438.08	\$464.98	\$363.11	\$404.84	\$380.37	\$417.03
rivals	2 routes	+8.8%	4 routes	+6.1%	10 routes	+11.5%		+9.6%
Non-stop and connecting	\$472.99	\$524.67	\$502.60	\$513.29	\$447.95	\$478.95	\$465.82	\$499.04
	91 routes	+10.9%	59 routes	+2.1%	158 routes	+6.9%		+7.1%
Both connecting	\$433.26	\$444.63	\$487.04	\$486.86	\$464.20	\$457.77	\$458.60	\$460.43
	479 routes	+2.6%	334 routes	0.0%	471 routes	-1.4%		+0.4%

Table 15: Price Effects and Changes in Service Type Choices Following Mergers Where Merging Parties are Non-Stop Duopolists Prior to the Merger

	Delta/Northwest		United/Continental		American/US Airways		Average	
	Price	Numb. New Non-Stop	Price	Numb. New Non-Stop	Price	Numb. New Non-Stop	Price	NNumb. ew Non-Stop
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%	
Service Types Endogenous								
No selection	$482.57	0.55	$505.80	1.23	$426.78	1.40	$451.94	1.26
	-14.8%		+0.4%		-7.0%		-6.3%	
Observable Selection	$574.64	0.34	$546.83	0.16	$470.05	1.07	$500.42	0.77
	+1.5%		+8.6%		+2.4%		+3.8%	
Obs. And Unobs. Selection	$589.38	0.09	$550.85	0.09	$498.64	0.46	$521.60	0.33
	+4.1%		+9.3%		+8.6%		+8.2%	
Number of Routes	2		4		11		17	

parties will provide non-stop service after the merger, we do not impose this on the rival non-stop carriers, which allows for the possibility that they will choose connecting service given appropriate quality and fixed cost draws.[45] In this case accounting for selection makes it more likely that existing non-stop rivals will maintain their non-stop service, even if it makes it less likely that new firms will initiate non-stop service after the merger. As a result, we may over-estimate how much prices will increase after a merger if we allow service decisions to be endogenous but ignore selection. With no selection, we estimate some changes in which carriers provide non-stop service, and some weakening of the competitive strength of the average non-stop rival, so that merging firms tend to increase prices. With only observable selection, there is less new non-stop service and some of the existing non-stop rivals (those that are weaker based on observables) are predicted to switch to only providing connecting service. When we account for both types of selection we predict that almost all carriers keep their pre-merger service types, so that predicted price increases are almost identical to those when service types are held fixed by assumption.

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 non-stop 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 non-stop 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 than a model that does not allow for selection will predict. We find that, on average, many of the mergers that we consider would appear quite benign (in the sense that they would not raise prices significantly) if we allowed for endogenous service choices, which seems natural in this industry, but did not account for selection, whereas we can explain the limited service changes that we observe 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 in 2001 would likely have been ineffective in constraining market power on routes where the merging parties were non-stop duopolists, suggesting that the Department of Justice's view that the remedy was insufficient, even for these routes, was very 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

[45] If the non-stop rival keeps on offering non-stop service it will have the same quality and marginal cost that we back out from the data (given the draw of the random effect), but we are allowing for the possibility that, for example, its draw of the incremental quality of non-service is small enough that, given the fixed cost draw, it will choose to provide connecting service.

Table 16: Price Effects of Mergers Where Both Merging Parties and at least One Other Carrier Provide Non-Stop Service Prior to the Merger

	Delta/Northwest		United/Continental		American/US Airways		United/US Airways		Average	
	Price	Change in Rival Non-Stop	Price	Change in Rival Non-Stop	Price	Change in Rival Non-Stop	Price	Change Rival Non-Stop	Price	Change in Rival Non-Stop
Pre-merger	$351.26	-	$438.08	-	$363.11	-	$350.02	-	$368.70	-
Post-Merger										
Service Types Fixed	$382.04	-	$464.98	-	$404.84	-	$378.15	-	$402.08	-
	8.8%		6.1%		13.5%		8.0%		9.1%	
Service Types Endogenous										
No selection	$407.32	-0.08	$474.00	0.04	$410.11	0.10	$428.28	0.25	$426.71	0.14
	16.0%		8.2%		12.9%		22.4%		15.7%	
Observable Selection	$391.81	-0.39	$481.85	-0.03	$424.02	-0.06	$387.20	-0.04	$416.28	-0.07
	11.5%		10.0%		16.8%		10.6%		12.9%	
Obs. And Unobs. Selection	$380.17	0.11	$464.94	0.00	$403.41	0.07	$376.33	0.10	$400.67	0.07
	8.2%		6.1%		11.1%		7.5%		8.7%	
Number of Routes	2		4		10		10		26	

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 extensioin 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 constrained airports, who these assets should be sold to and the ways that rivals may use these assets if they purchase them (for example, choosing to serve a small number of routes intensively or maybe providing limited non-stop service on a number of routes).

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.

Athey, S., J. Levin, and E. Seira (2011): "Comparing open and sealed bid auctions: Evidence from timber auctions," *The Quarterly Journal of Economics*, 126(1), 207–257.

Azar, J., M. C. Schmalz, and I. Tecu (2016): "Anti-competitive effects of common ownership," Discussion paper, University of Michigan.

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. (1992a): "Estimation of a Model of Entry in the Airline Industry," *Econometrica: Journal of the Econometric Society*, pp. 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. (1992b): "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 60(4), 889–917.

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," *The 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 market," *The Review of Economic Studies*, 57(4), 531–553.

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

Cabral, L. M. (2003): "Horizontal mergers with free-entry: why cost efficiencies may be a weak defense and asset sales a poor remedy," *International Journal of Industrial Organization*, 21(5), 607–623.

Carlton, D. W., M. A. Israel, I. MacSwain, and E. Orlov (2016): "Are Legacy Airline Mergers Pro-or Anti-Competitive? Evidence from Recent US Airline Mergers," Discussion paper, University of Chicago.

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 models1," *Econometrica*, 75(5), 1243–1284.

Ciliberto, F., C. Murry, and E. Tamer (2015): "Market Structure and Competition in Airline Markets," Discussion paper, University of Virginia.

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," *The 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," *QME*, 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," *The 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.

Gentry, M., and T. Li (2014): "Identification in auctions with selective entry," *Econometrica*, 82(1), 315–344.

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

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

Hüschelrath, K., and K. Müller (2014): "Airline networks, mergers, and consumer welfare," *Journal of Transport Economics and Policy (JTEP)*, 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.

Kwoka, J., and E. Shumilkina (2010): "The Price Effects of Eliminating Potential Competition: Evidence from an Airline Merger," *Journal of Industrial Economics*, 58(4), 767–793.

Li, T., and B. Zhang (2015): "Affiliation and entry in first-price auctions with heterogeneous bidders: an analysis of merger effects," *American Economic Journal: Microeconomics*, 7(2), 188–214.

Li, T., and X. Zheng (2012): "Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions," *Journal of Econometrics*, 168(1), 29–46.

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*, pp. 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*, pp. 237–250.

Nevo, A. (2000): "Mergers with differentiated products: The case of the ready-to-eat cereal industry," *The 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," *The Journal of Law and Economics*, 32(2, Part 2), S179–S202.

Roberts, J. W., and A. Sweeting (2013): "When should sellers use auctions?," *The 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," *The 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," *The 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.

A 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 non-stop to affect its decision to do so. Connecting traffic is especially important in explaining why a large number of non-stop 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 flows that fits the data well[46] and which gives us a prediction of how much traffic that a carrier can generate on a route where it does not currently provide non-stop service, taking the type of service that it supplies on other routes as given. We include this prediction in our model of entry as a variable that can reduce the effective fixed cost of providing non-stop service on the route.[47]

We specify a multinomial logit specification for the share of people traveling using one-stop connecting service on a particular route (e.g., RDU to SFO) who choose to connect via a particular domestic hub-carrier combination (e.g., via DFW on American, or Memphis on Northwest). 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-origin-destination 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.[48] 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, doing so would ignore the selection problem that a particular c, i, od combination might not be available because the carrier would not serve enough traffic to make it profitable to serve both segments o to c and c to d with non-stop service (which it has to do to offer a one-stop connection via c). We therefore introduce an additional probit model, as part of a classic Heckman selection model, to describe the probability that carrier c does serve the full ocd route,

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

[46] 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 non-stop service on each segment, but which could for the entry decisions considered in our model.

[47] We also use the predicted value, not the actual value, on routes where we actually observe non-stop service.

[48] For example, the outside good for RDU to SFO could involve traveling via Nashville on any carrier (because Nashville is not a domestic hub) or on Delta via DFW because, during our data, DFW is not defined as a domestic hub for Delta even if it is on American.

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). We can then use the estimates to calculate a predicted value for each carrier's share of traffic on a particular connecting route if it provides non-stop service on a particular *oc* or *od* segment, holding service on other routes and on other carriers fixed, and we then multiply this by the number of connecting passengers and add up across all *oc* and *od* pairs involving a segment to get our prediction of the number of connecting passengers served if non-stop service is provided. The results are very similar using a two-step approach (which does not impose that the ξ_js are normally distributed), with the correlation of resulting predicted values at the route-level above 0.999. When we aggregate predictions up to this level, the estimated model does a pretty accurate job of predicting how many people travel given the set of routes that airlines currently fly. At the segment level our model performs quite well at predicting traffic for those routes that are served non-stop. For example, for the identified legacy carriers in our primary model (AA, CO, DL, NW, UA, US), 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 also does a good job of predicting some of the variation created by geography. For example, the model predicts that AA should serve 2,247 connecting DB1 passengers on RDU-DFW, 1213 on RDU-ORD and 376 on RDU-STL, which compares with observed numbers of 2,533, 1,197 and 376 (here we are using DB1 numbers). On the other hand, from Boston the model correctly predicts that AA will serve more connecting traffic via ORD (2265, observed 2765) than DFW (2040, observed 2364).

We estimate our model using data from Q2 2005 (one year prior to the data used for our main model) for the top 100 markets. During estimation we do not aggregate the other low cost and other legacy carriers, but instead consider them individually and allow them to have their own domestic hubs (such as Atlanta for Airtran and Denver for Frontier). 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 drop origin-destination routes with less than 25 connecting passengers (adding up across all connecting routes) or where any origin-connection or connection-destination segment is less than 100 miles long.[49] We also drop carrier-origin-destination-connecting airport observations where the carrier (or one of its regional affiliates) is not, based on T100, providing non-stop 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.

[49] 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.

In X_{cj} (for the share equation), we include variables designed to measure the attractiveness of the carrier c and the particular connecting route. Specifically, the included variables are carrier c's presence at the origin and its square, its presence at the destination and its square, the interaction between carrier c's origin and destination presence, the distance involved in flying route j divided by the non-stop distance between the origin and destination (we call this the 'relative distance' of the connecting route), an indicator for whether route j is the shortest route involving a hub, an indicator for whether j is the shortest route involving a hub for carrier c 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 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 non-stop service is offered but not the connecting share equation. In W_{cj} we include origin, destination and connecting airport presence for carrier c; 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[50]; indicators for whether either of the origin or destination airports is an airport with limitations on how far planes can fly (LaGuardia and Washington 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_{cj} and W_{cj} we also include origin, destination and carrier-connecting airport dummies.

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

Table A.2: Estimation Coefficients for Ancillary Model of Connecting Traffic

	(1) Connecting Share	(2) Presence	(3) $\frac{1}{2}log\frac{1+\rho}{1-\rho}$	(4) $log(std.\ deviation)$
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

B 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 simpler model than the model that we estimate once using the real data. In particular, the simplified model involves carriers making a binary entry decision, where a firm that does not enter the market is assumed not to participate in the market at all, in contrast to our full model where we allow for carriers that do not enter directly to provide connecting service. 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.[51]

B.1 Model

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

B.1.1 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.

B.1.2 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.

B.1.3 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 parameters are their true values. Carrier i's marginal cost, $c_{i,m}$, is also drawn from a truncated normal

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

[51] The number of parameters estimated is also large in the Monte Carlos reported in Ciliberto, Murry, and Tamer (2015). However, in their Monte Carlos they do not attempt their full procedure of placing bounds on the parameters.

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 heterogeneity in a linear 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:

$$\begin{aligned}
\eta_m &\sim TRN(0, \underset{0.5}{\sigma^\eta}, -2, 2) \\
\alpha_m &\sim TRN(\underset{0.45}{\mu^\alpha}, \underset{0.1}{\sigma^\alpha}, 0.15, 0.75) \\
\lambda_m &\sim TRN(\underset{0.7}{\mu^\lambda}, \underset{0.1}{\sigma^\lambda}, 0.5, 0.9)
\end{aligned}$$

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.

B.1.4 Order of Entry

We assume that the data is generated by carriers playing a complete information game (i.e., each carrier knows the demand and cost draws of all of its rivals) with sequential entry, where carriers simultaneously set prices once entry decisions have occurred. Legacy carriers are assumed to move first, ordered by $X_{i,m}^D$ (highest first), followed by low cost carriers who are ordered randomly. The firms know the order. We consider Monte Carlos where we assume this order is known to the researcher and when this order is not known. 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[52], the game will have a unique subgame perfect Nash equilibrium.

B.1.5 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 α (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 α 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 σ (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_m^C) increases average prices from 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.2 Monte Carlo Exercises

There are 17 parameters, $\Theta = \{\beta^{D,LEG}, \beta^{D,LCC}, \beta_1^D, \sigma^D, \gamma^{C,LEG}, \gamma^{C,LCC}, \gamma^C, \sigma^C, \theta_1^F, \theta_2^F, \theta_3^F, \sigma^F, \sigma^\eta, \mu^\alpha, \sigma^\alpha, \mu^\lambda, \sigma^\lambda\}$, to be estimated. Label the true parameters Θ_0. We present results for three Monte Carlo exercises below.

[52]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.

B.2.1 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

$$\frac{1}{S}\sum_{s=1}^{S} h(x_m, \Psi_{s,m}) \frac{f(\Psi_{s,m}|x_m,\Theta')}{g(\Psi_{s,m}|x_m)}$$

where, in our setting, $\Psi_{s,m}$ 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 Θ', g is the importance density from which $\Psi_{s,m}$ is drawn, and $h(x_m, \Psi_{s,m})$ is the value of the outcome of interest given observed market characteristics and $\Psi_{s,m}$ (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., $g(\Psi_{s,m}|x_m) \equiv f(\Psi_{s,m}|x_m,\Theta_0)$. 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.[53] 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)[54], and three market outcomes: the 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[55] and the square of the number of entrants.

These outcome measures are then interacted with the 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.[56]

[53] 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 by, for each dataset, drawing 1,000 data 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.

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

[55] 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.

[56] 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.

Column (1) of Table B.2 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.3 reports the expected changes in entry, the cumulative market share of entering carriers and average prices under the true parameters.[57] 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 to fall.[58] The second column reports the mean changes and standard deviations (in parentheses) across the 100 Monte Carlo repetitions.[59] We can see that the Monte Carlo counterfactuals predict the true effects accurately, with small standard deviations.[60]

Column (2) of Table B.2 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.[61]

B.2.2 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

[57] 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.

[58] Average prices are only calculated for markets where entry occurs, so average prices after the increase in fixed costs are calculated based on a subset of the markets used to calculate before change average prices.

[59] 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.

[60] 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.

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

Table B.2: 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 (0.073)	0.473 (0.151)	0.511 (0.089)
Mkt Demand Slope	Mean (μ^{α})	-0.45	-0.421 (0.024)	-0.429 (0.026)	-0.425 (0.022)
	Std. Dev. (σ^{α})	0.1	0.038 (0.039)	0.075 (0.056)	0.051 (0.044)
Nesting Parameter	Mean (μ^{λ})	0.7	0.694 (0.033)	0.689 (0.035)	0.701 (0.054)
	Std. Dev. (σ^{θ})	0.1	0.051 (0.039)	0.062 (0.033)	0.089 (0.072)
Carrier Quality	Legacy ($\beta^{D,LEG}$)	0.2	0.189 (0.064)	0.190 (0.103)	0.188 (0.076)
	LCC ($\beta^{D,LCC}$)	0	0.000 (0.064)	-0.031 (0.087)	0.003 (0.069)
	$X^{D}_{i,m} * LEG_i$ (β^{D}_{1})	0.3	0.295 (0.067)	0.295 (0.142)	0.293 (0.097)
	Std. Dev. (σ^{D})	0.2	0.176 (0.043)	0.209 (0.064)	0.170 (0.050)
Carrier Marginal Cost	Legacy Constant ($\gamma^{C,LEG}$)	0	0.031 (0.111)	0.040 (0.133)	0.054 (0.130)
	LCC Constant ($\gamma^{C,LCC}$)	-0.5	-0.507 (0.135)	-0.470 (0.141)	-0.483 (0.158)
	X^{C}_{m} (γ^{C})	0.5	0.500 (0.034)	0.479 (0.047)	0.489 (0.042)
	Std. Dev. (σ^{C})	0.2	0.216 (0.069)	0.169 (0.081)	0.205 (0.072)
Carrier Fixed Cost	Constant ($\theta^{F}_{1}/10,000$)	0.75	0.738 (0.096)	0.725 (0.131)	0.743 (0.101)
	$X^{F}_{i,m}$ ($\theta^{F}_{2}/10,000$)	0.1	0.110 (0.081)	0.118 (0.166)	0.121 (0.130)
	X^{F}_{m} ($\theta^{F}_{3}/10,000$)	0.5	0.556 (0.126)	0.599 (0.163)	0.548 (0.142)
	Std. Dev. ($\sigma^{F}/10000$)	0.25	0.210 (0.065)	0.246 (0.084)	0.209 (0.086)

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

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 (3) of Table B.2 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.

As pointed out by Geweke (1989), the properties of an importance sampling estimator of an expected value will depend on whether the variance of $h(x_m, \Psi_{s,m})\frac{f(\Psi_{s,m}|x_m,\Theta')}{g(\Psi_{s,m}|x_m)}$ is finite. 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 quite wildly as S rises. Figure B.2 shows these recursive estimates of the sample variance for the moments associated with the three market level outcomes (the sum of squared market shares, the squared number of entrants and the weighted average price paid), where the IS draws come from the wider densities that we have just described and the parameters are equal to the average parameters estimated across the 100 repetitions (i.e., those from column (3) of Table B.2).[62] Relative to examples in Koopman, Shephard, and Creal (2009), the jumps in the estimated sample variance seem quite small for $S > 500$.

B.2.3 Monte Carlo Exercise 3: Estimation When the Econometrician Only Knows that a Pure Strategy Nash Equilibrium is Played

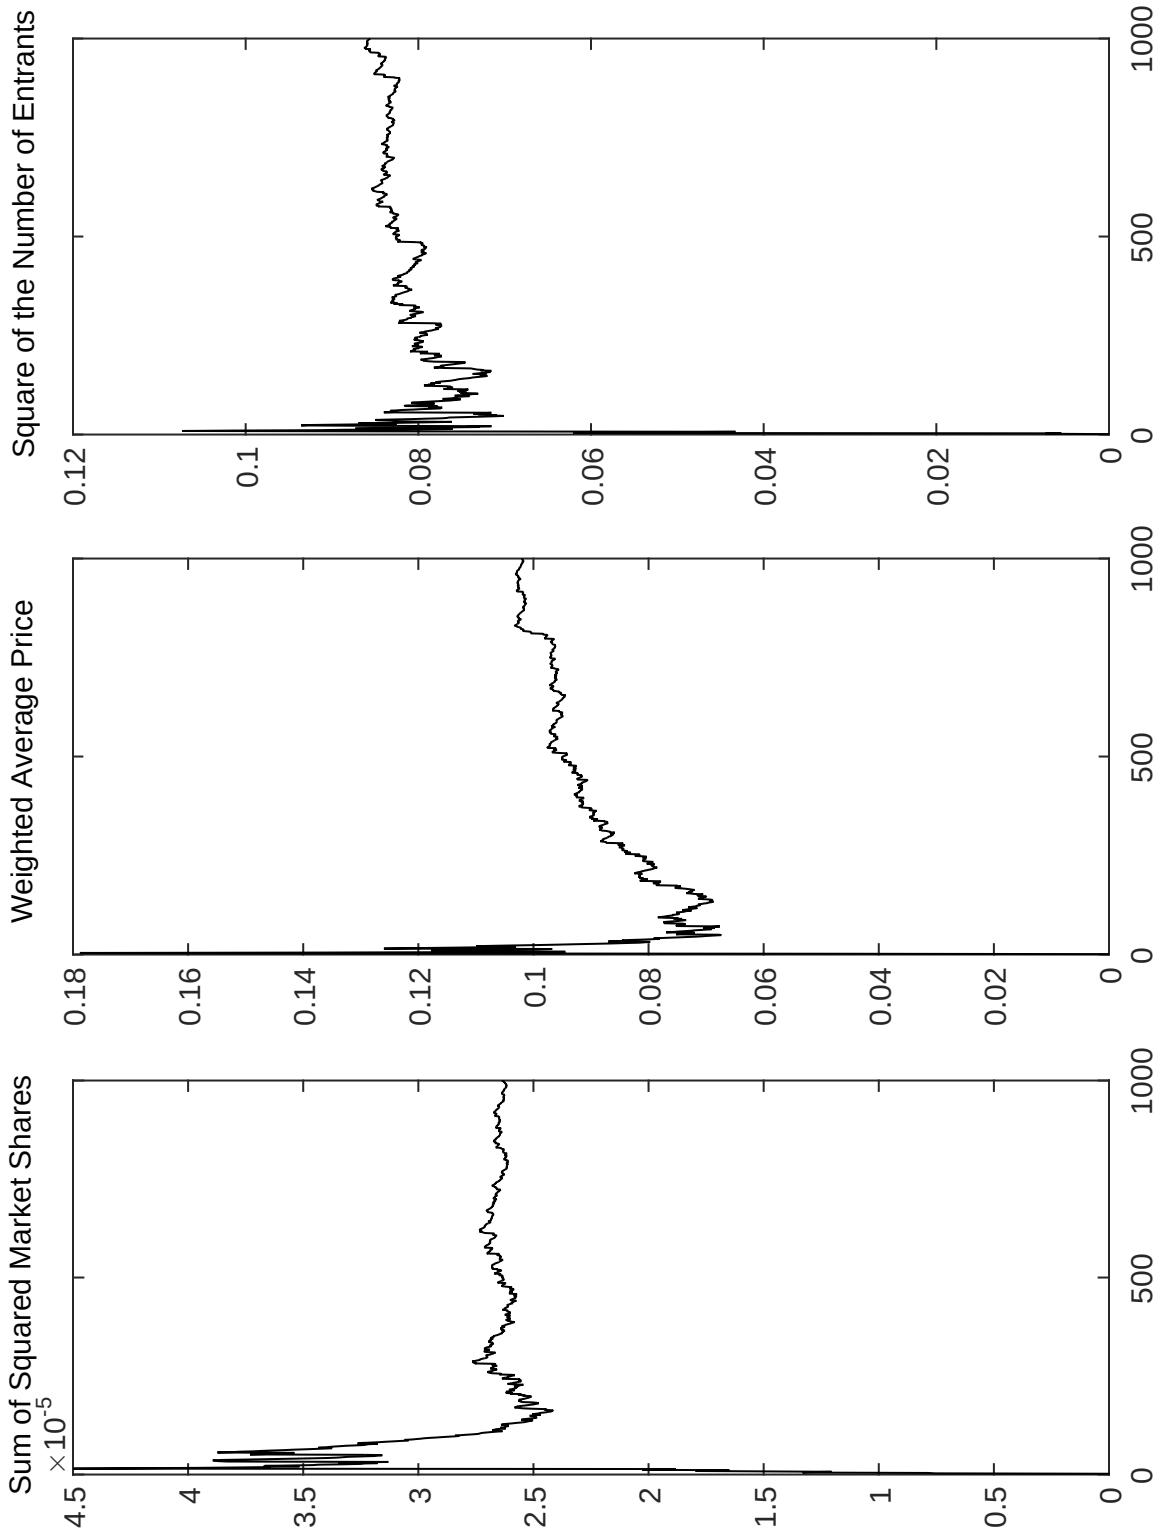
Our third exercise considers estimation when we relax the assumption that entry decisions are made in a known sequential order. Instead, we follow the strand of the literature (most notably, Ciliberto and Tamer (2009)) that has based estimation on moment inequalities formed under the assumption that firms play some pure strategy Nash equilibrium in a simultaneous move game, where the equilibrium is known to the players but not to the researcher.[63] The idea is that, as long as the set of equilibria 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 of the parameters that make inequalities based on these lower and upper bounds hold. For example,

$$\mathbb{E}(m(h, X, Z.\Theta)) = \mathbb{E}\left[\begin{array}{c} h_m^{DATA} - \underline{\mathbb{E}(\widehat{h_m(X,\Theta)})} \\ \\ \overline{\mathbb{E}(\widehat{h_m(X,\Theta)})} - h_m^{DATA} \end{array} \otimes Z_m\right] \geq 0$$

[62] The x-axis shows the number of simulations per market and the y-axis is the variance of $\frac{1}{M}\sum h(x_m, \Psi_{s,m})\frac{f(\Psi_{s,m}|x_m,\Theta')}{g(\Psi_{s,m}|x_m)}$ across simulations $s = 1, .., S$ where S is the value on the x-axis.

[63] A closely related exercise that we have also experimented with assumes that entry takes place sequentially but that the order in which players move is unknown. For given quality and cost draws, the set of outcomes that can be supported as equilibria in a simultaneous move game and as equilibrium outcomes in a sequential move game with some order is often the same, but one can construct examples, even with binary (enter/stay out) choices where this is not the case.

Figure B.2: Recursive Estimate of the Variance of Market-Level Moments



where h_m^{DATA} are observed outcomes in the data and Z_m are instruments. $\underline{\widehat{\mathbb{E}(h_m(X,\Theta))}}$ and $\overline{\widehat{\mathbb{E}(h_m(X,\Theta))}}$ are predicted lower and upper bounds on the expected value of h_m given a set of parameters Θ. For example, suppose that the outcome is whether firm A enters the market. The lower bound would be formed by assuming that whenever there are pure strategy Nash equilibria where A does **not** enter, one of them will be played, whereas the upper bound would be formed by assuming that whenever there are equilibria where A enters, one of them is played. Intuitively outcome moments based on the data should lie between these upper and lower bounds.[64] In performing a Monte Carlo exercise, we follow Exercise 1 in using the true distributions of the parameters when taking our importance sample draws. The main point of this exercise is to illustrate that it is quite possible to adapt the importance sampling estimate to allow for the researcher to be agnostic about equilibrium selection.

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 four 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 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 equilibria in 24.2% of the 2,000 sample data markets. In almost all of these cases, the equilibria differ only in the identity of entrants rather than the number of firms that enter.

For the games that we draw to form our importance sampling set, we solve for all of the pure strategy Nash equilibria for each draw ($\Psi_{s,m}$), and then calculate lower and upper bounds on each of the outcomes ($\underline{h(x_m,\Psi_{s,m})}$, $\overline{h(x_m,\Psi_{s,m})}$). The lower and upper bounds will be identical if there is only one equilibrium. During estimation we re-weight the lower and upper bounds from each simulated game as before to calculate their expectation given parameters Θ. For example,

$$\underline{\widehat{\mathbb{E}(h_m(X,\Theta))}} = \frac{1}{S}\sum_{s=1}^{S} \underline{h(x_m,\Psi_{s,m})} \frac{f(\Psi_{s,m}|x_m,\Theta)}{g(\Psi_{s,m}|x_m)}$$

The values of $\underline{\widehat{\mathbb{E}(h_m(X,\Theta))}}$ and $\overline{\widehat{\mathbb{E}(h_m(X,\Theta))}}$ will be different as long as there are multiple equilibria (which differ in their hs) for at least some of the importance sample draws.

The objective function that is minimized is

$$Q(\Theta) = \min_{t\preceq 0} [m(\widehat{h,X,Z},\Theta) - t] W [m(\widehat{h,X,Z},\Theta) - t]$$

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.

There are now many papers that propose approaches for inference for moment inequality models (for

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

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 minimized values of the objective function that lie slightly above, but close, to zero (its mean value is 0.0026, with a standard deviation of 0.001 across our Monte Carlo runs).[65] Table B.4 reports the Monte Carlo results, using 1,000 markets and 1,000 IS draws for each market in each Monte Carlo run.[66]

Comparing the results to those from column (1) of Table B.2 (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 equilibria in a higher proportion of markets.

Figures B.3-B.5 show the shape of the objective function from 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. The smoothness of the objective function reflects one of the main advantages of the importance sampling approach as it allows us to use minimization methods that use derivatives, which is important when trying to minimize over a large number of parameters. On the other hand, for three parameters (γ^C, $\beta^{D,LEG}$, $\beta^{D,LCC}$) it is also clear that 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.

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

[66] 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.

Figure B.3: Objective Function Around the Estimated Parameters for the First Monte Carlo Run

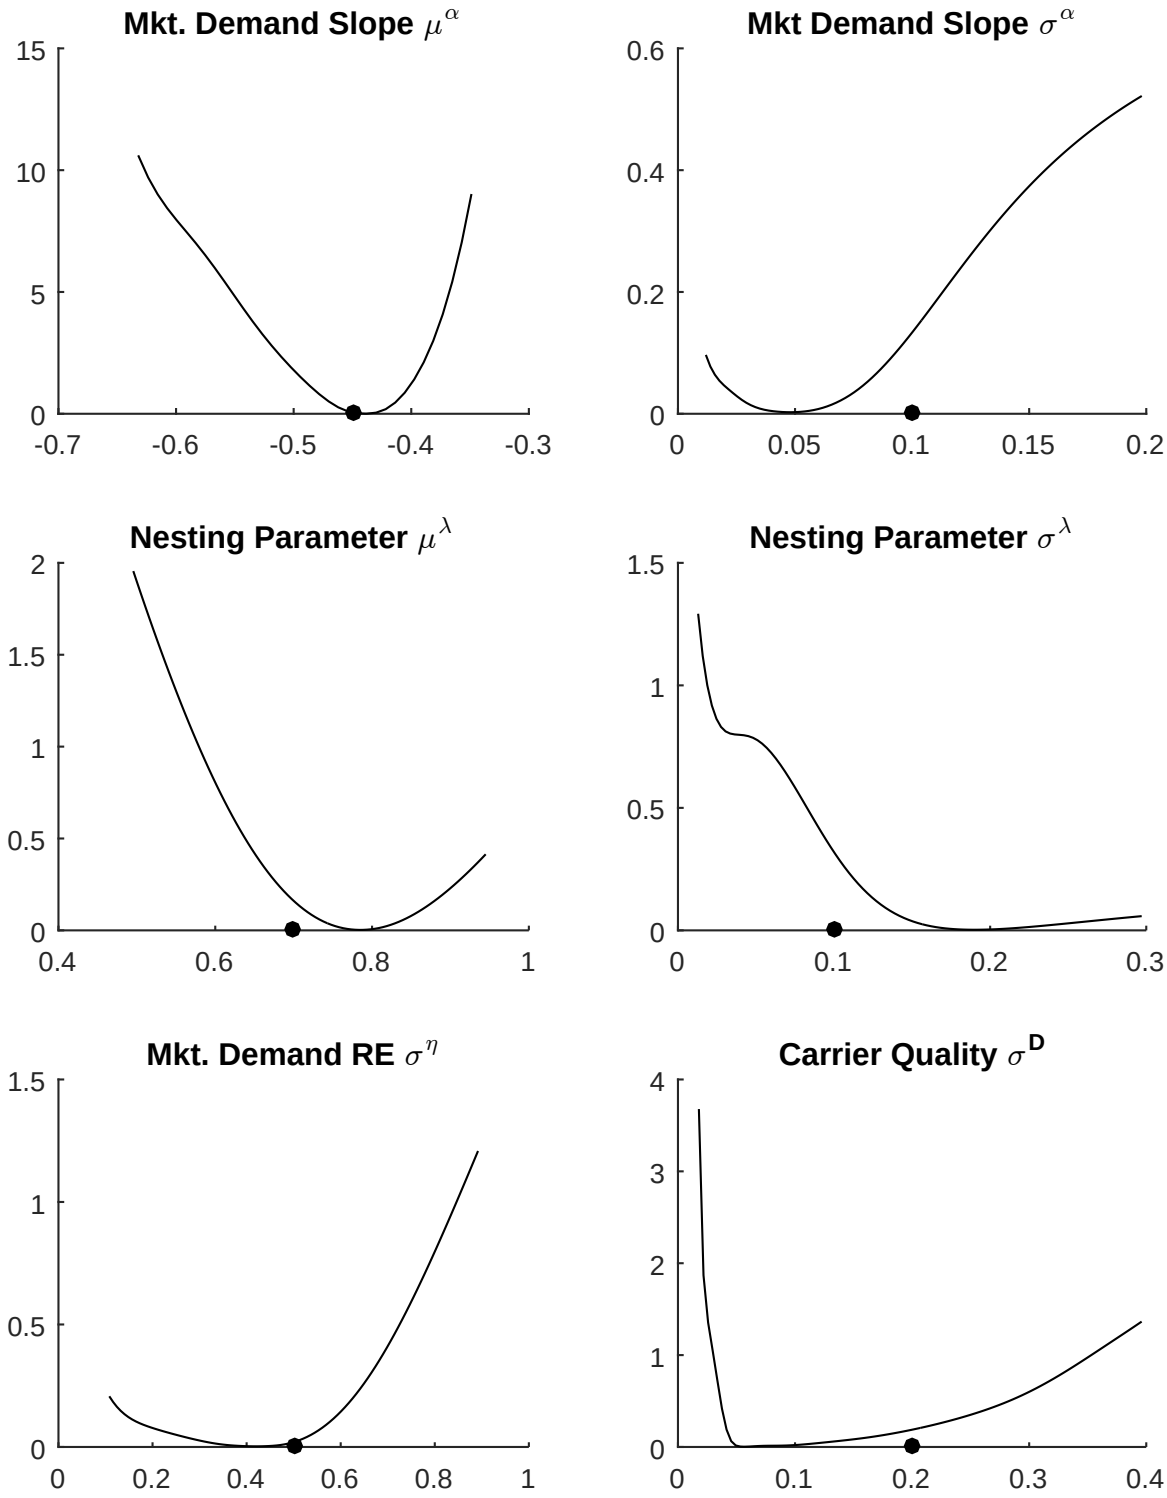


Figure B.4: Objective Function Around the Estimated Parameters for the First Monte Carlo Run

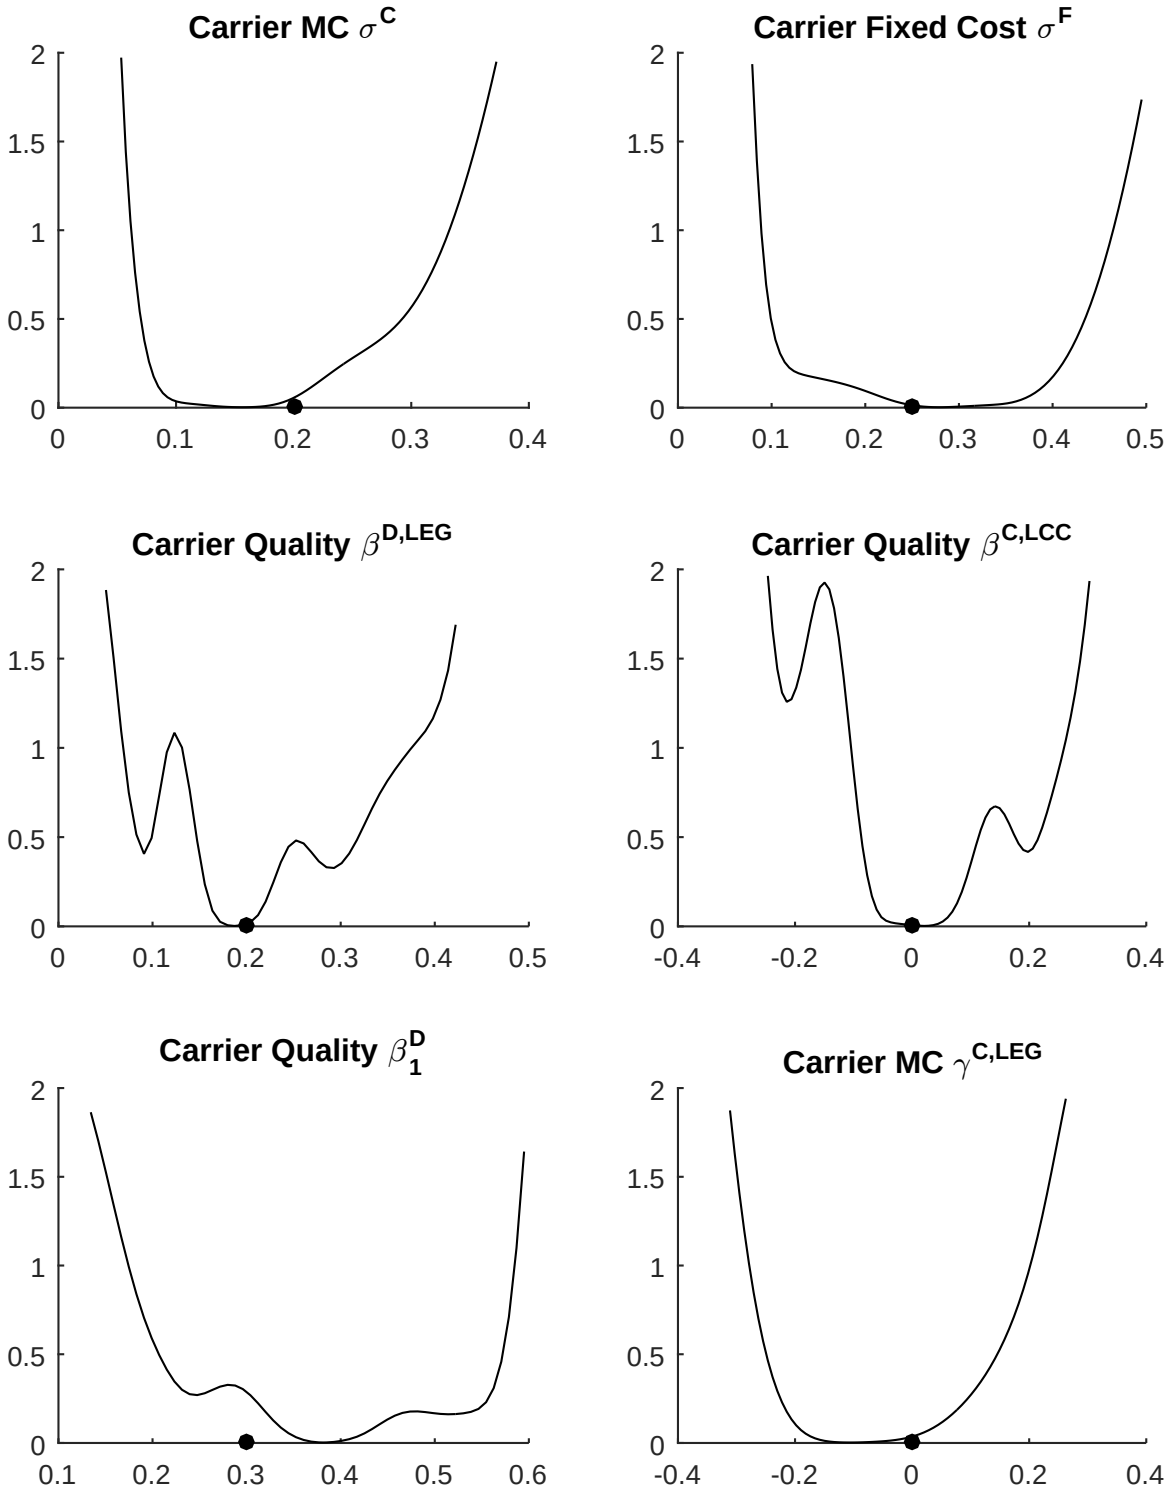


Figure B.5: Objective Function Around the Estimated Parameters for the First Monte Carlo Run

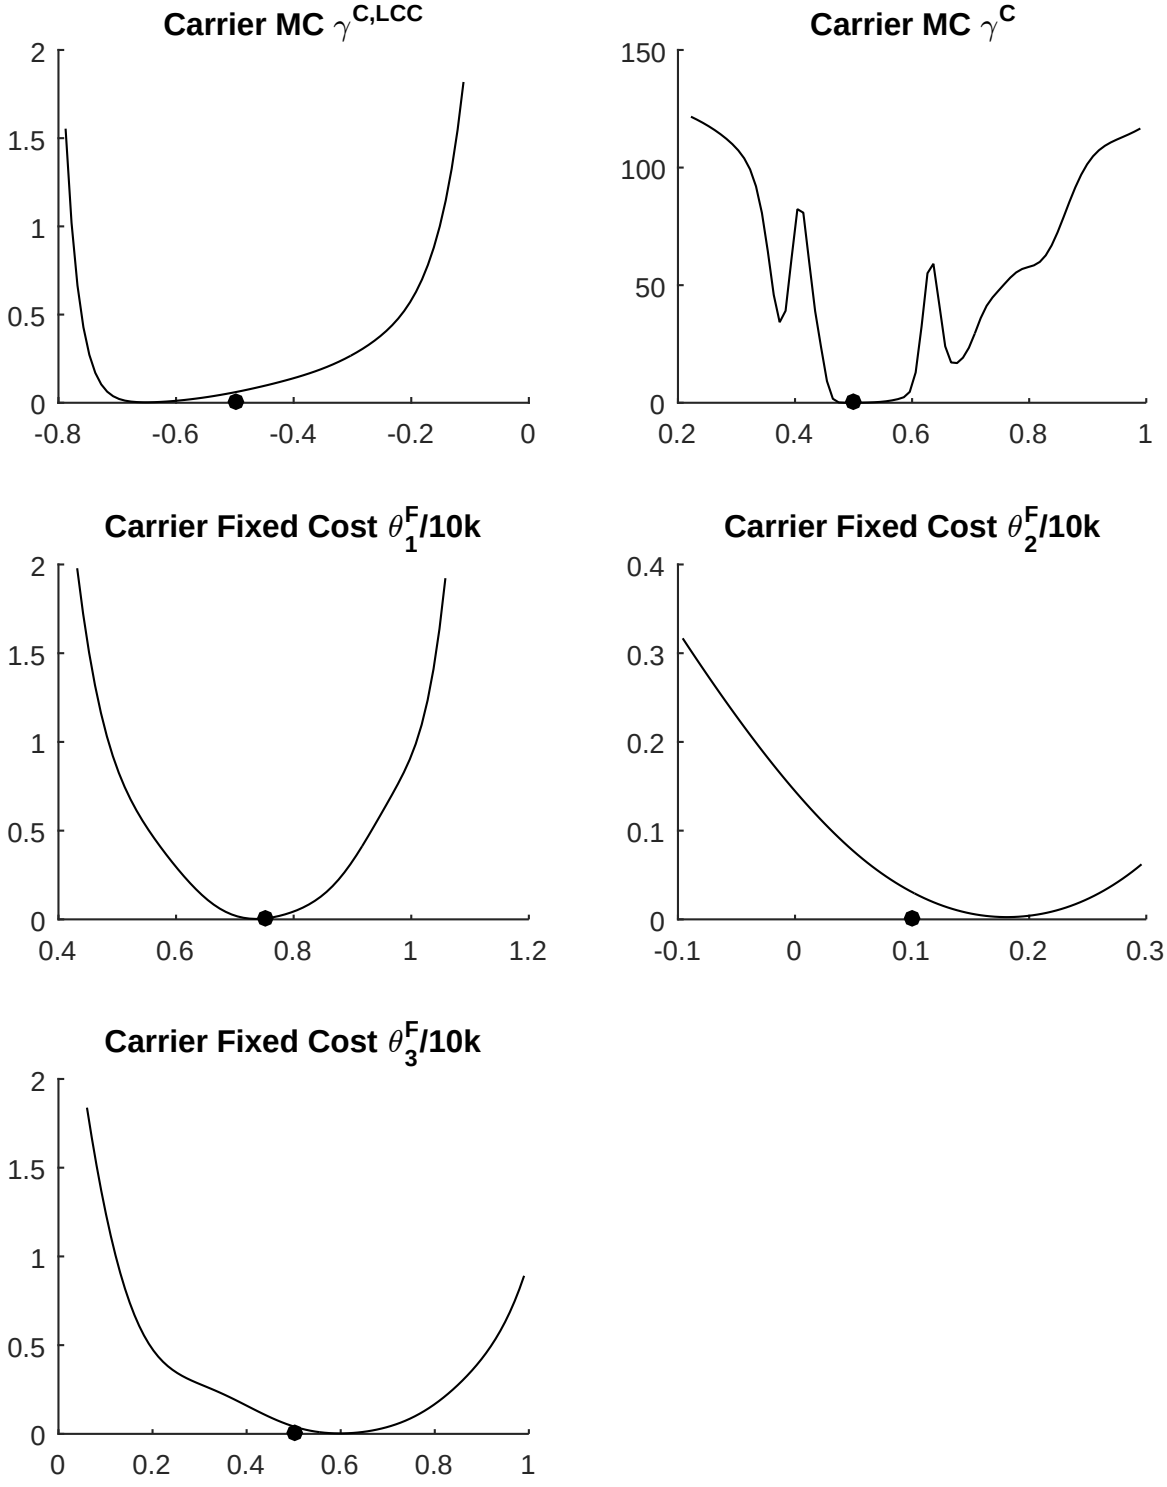


Table B.3: Illustrative Counterfactual: The Effects of Increasing the Fixed Entry Costs of Legacy Carriers Using Parameters Estimated Using IS Distributions that are the Same as True Distribution of the Parameters and 1,000 Markets

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 at least one firm entering)	+0.228	+0.219
		(0.056)

Table B.4: 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)
	$X^{D}_{i,m} * LEG_i$ (β^{D}_{1})	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^{C}_{m}	0.5	0.488
			(0.032)
	Std. Dev. (σ^{C})	0.2	0.189
			(0.059)
Carrier Fixed Cost	Constant ($\theta^{F}_{1}/10,000$)	0.75	0.696
			(0.104)
	$X^{F}_{i,m}$ ($\theta^{F}_{2}/10,000$)	0.1	0.213
			(0.147)
	X^{F}_{m} ($\theta^{F}_{3}/10,000$)	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.