Repositioning and Market Power After Airline Mergers*

Sophia Li[†], Joe Mazur[‡], Yongjoon Park[§], James Roberts[¶],
Andrew Sweeting[‖] and Jun Zhang[**]

January 2021

Abstract

We estimate a model of route-level competition in the airline industry in which carriers choose whether to offer nonstop or connecting service before setting prices. Carriers have full information about the quality, marginal cost and fixed cost unobservables of all carriers throughout the game, so that carriers' service choices will be selected on these residuals. We conduct merger simulations that allow for repositioning and account for the selection implied by the model and the data. Accounting for selection substantially affects predictions about the likelihood of repositioning and the magnitude of post-merger price changes, and it allows us to match what has been observed after consummated mergers.

Keywords: product repositioning, market power, endogenous market structure, selection, horizontal mergers, remedies, discrete choice games, multiple equilibria, airlines.

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

*We thank two referees, Co-Editor Katja Seim, numerous discussants, and seminar and conference participants for very useful comments. Early stages of the research were supported by NSF Grant SES-1260876 (Roberts and Sweeting). Earlier versions of this paper were circulated as "Airline Mergers and the Potential Entry Defense" and "Endogenous and Selective Service Choices After Airline Mergers". The revision was prepared when Sweeting was serving as the Director of the Bureau of Economics at the Federal Trade Commission (FTC). All of the views expressed are the authors' own and do not reflect the views of the FTC or any individual Commissioner. The usual disclaimer applies.

†Li: ying.sophia@gmail.com.

‡Mazur: Krannert School of Business, Purdue University, mazur3@purdue.edu.

§Park: Department of Resource Economics, University of Massachusetts Amherst, yongjoonpark@umass.edu.

¶Roberts: Department of Economics, Duke University and NBER, jr139@duke.edu.

‖Sweeting (corresponding author): Department of Economics, University of Maryland and NBER, atsweet@umd.edu.

**Zhang: School of Agricultural Economics and Rural Development, Renmin University of China, junzhang@ruc.edu.cn.

1 Introduction

Market power created by a horizontal merger may be limited if the merger induces either new entry or existing rivals to reposition to compete more directly with the merging firms. Since 1992, the *Horizontal Merger Guidelines* have specified that the agencies should try to test whether entry or repositioning will be "timely" and "likely", in the sense of being profitable for rivals, and therefore likely to happen, and "sufficient", in the sense of preventing prices from rising (Shapiro (2010), p. 65).[1] While economists accept these criteria, they are rarely assessed in the rigorous and quantitative way that estimated or calibrated merger simulations are used to predict price changes with a fixed set of products.[2] This paper presents a quantitative framework for assessing the likelihood and the sufficiency of repositioning in differentiated product markets. Our empirical analysis models service choices and pricing in route markets after airline mergers, motivated by how several airline mergers in the 1980s were approved based on ease-of-entry/repositioning arguments (Keyes (1987)), and by the suggestions of Fisher (1987) and Schmalensee (1987) that airline mergers provide a setting where repositioning could offset anticompetitive effects.

We use a two-stage model where carriers first choose their discrete service types (nonstop or connecting) and then choose prices. As motivation, suppose that nonstop service has significantly higher quality than connecting service, with similar marginal costs but with a higher fixed cost. A market has two carriers providing nonstop service and two other carriers that provide connecting service via other airports. The nonstop carriers propose to merge, and an analyst has to evaluate whether the merger will raise prices. The answer may (and, in our results, it often does) depend on whether the merger will create an incentive for a connecting carrier to initiate nonstop service (i.e., reposition), and, if one does, the quality that its nonstop service is likely to have.

We make two assumptions that distinguish our analysis from the literature. First, we assume that all elements of qualities and costs, including those unobserved by the analyst, are known to all carriers when they make service choices. We describe this assumption as "full information", and

[1] This formulation was a reaction to courts allowing mergers to proceed based on claims that potential entrants did not face entry barriers that were higher than incumbents had faced, without an evaluation of whether this would be sufficient to prevent anticompetitive effects (*United States v. Waste Management, Inc.* (743 F.2d 976, 978, 983-84, 2d Cir. 1984), *United States v. Baker Hughes Inc.* (908 F.2d 981, 988-89, D.C. Cir. 1990) and *United States v. Syufy Enterprises*, (903 F.2d 659, 661, 9th Cir. 1990)).

[2] Instead, as in the period prior to 1992, both court decisions and agency analysis continue to focus on barriers to entry or repositioning without clear connections to profitability or price effects. For example, Coate (2008) describes the FTC's conclusions about the likelihood of entry in internal memoranda as lacking a "solid foundation" in the evidence, while Kirkwood and Zerbe (2009) classify only one of 35 post-1992 court opinions as reviewing the criteria in the *Guidelines* systematically. Some decisions, such as *Oracle* (331 F.Supp. 2d 1098 (N.D. Cal. 2004)) discuss new entry but are primarily decided on prior questions of market definition.

it contrasts with a "limited information" assumption where only the distributions of quality and marginal cost unobservables are known when discrete choices are made. The second assumption is that the unobservables of the non-merging carriers after a merger will be the same as before the merger.[3] We will consider alternative synergy assumptions for the merged firm. We provide evidence that carrier demand and marginal cost unobservables are persistent in the data, which is consistent with both of our assumptions.

The full information assumption implies that carriers' service choices will be selected based on the unobservables. For example, nonstop carriers will tend to have higher nonstop quality unobservables. This creates a challenge for estimation which we overcome by simultaneously estimating the demand, marginal cost and fixed cost equations, using the importance sampling method proposed by Ackerberg (2009) to reduce the computational burden.[4] The combination of the two assumptions implies that we need to calculate conditional distributions of the unobservables that are consistent with observed service choices to perform counterfactuals. The novel methodological contribution of our paper comes from providing a routine that implements this conditioning, and we show that conditioning impacts our counterfactual predictions.[5] As we discuss in Section 7, our calculation of conditional distributions that are consistent with pre-merger data distinguishes our approach from Ciliberto, Murry, and Tamer (2020) (CMT), who also estimate a full information model using airline data.

Our counterfactuals consider three mergers that were completed after the period of data that we use to estimate our model (Q2 2006) and one merger, between United and US Airways, that was proposed but blocked in 2001. We focus on routes where the merging carriers were both nonstop as these are the markets where merger simulations with fixed products predict the largest price increases. We find that when we condition on pre-merger service choices, our predictions are consistent with what happened after completed mergers: specifically, with conditioning, we predict that rivals launch nonstop service on 18% of nonstop duopoly routes (i.e., routes where the merging firms were the only nonstop carriers), which is close to the 25% rate observed for such

[3]This assumption is standard in the literature that treats products as fixed. However, the literature that endogenizes product choices under limited information has assumed that firms expect to receive new unobservable draws after a merger.

[4]Ackerberg's Example 2 explains how the method could be applied to this type of game. While the method has been used by Laffont, Ossard, and Vuong (1995), Roberts and Sweeting (2013) and Wang (2015), amongst others, we believe that we are the first to apply the method in the context of a discrete choice-and-price competition game with up to nine players and several player-specific unobservables. We find that this estimation approach works well in practice.

[5]Conditioning captures the essence of a frequent agency argument that courts should be skeptical that rivals will enter or reposition after a merger when they have chosen not to do so previously (Baker (1996), p. 364).

routes within two years of a completed merger. In contrast, we predict three times as many nonstop launches when we do not condition on pre-merger choices (i.e., we assume that carriers draw new unobservables post-merger). Conditioning also leads to mergers appearing to be more profitable.

Before discussing the related airline literature, we acknowledge several restrictive features of our approach. First, our model is static rather than dynamic. This is consistent with the short-run focus of most merger analysis (Carlton (2004)), but we provide a comparison to the dynamic airline models of Aguirregabiria and Ho (2012) (AH) and Benkard, Bodoh-Creed, and Lazarev (2020) (BBCL) in Section 7. Second, we focus on whether non-merging carriers will initiate nonstop service on particular routes after a merger, taking carriers' network structures (e.g., which airports are hubs) as fixed. This ignores how a merger might lead carriers to eliminate or add hubs. Third, the non-merging carriers that we focus on are those that provide connecting service prior to the merger. We therefore ignore the possibility of new entry. The primary reason is that, while a model that allows for a "no service" option can be estimated (Li, Mazur, Roberts, and Sweeting (2015)), there are additional unobservables that make it costly to implement our conditioning routine for counterfactuals. However, our approach is consistent with how most carriers that initiate nonstop service previously provide connecting service (see Section 6.3.1) and we also consider whether adding an additional carrier post-merger affects our counterfactual results, motivated by how slot divestitures, that have sometimes been required as part of the merger approval process, may allow new carriers to enter the airports at the endpoints of a route. Fourth, we only model carriers' choice of service types and a single price for traffic originating at each endpoint, ignoring choices of capacity, schedules and the allocation of seats to different price bins. A more complete model would include these choices, which would introduce additional unobservables.[6] Finally, our baseline assumption will be that carriers make service choices sequentially, which guarantees a unique equilibrium, whereas much of the literature allows for multiple equilibria in discrete choice simultaneous move games. We will explain why this assumption does not materially affect our results.

Two related literatures use airline data. Many merger retrospectives have evaluated the price effects of carrier mergers in airport-pair or city-pair markets, both in the 1980s (summarized in Ashenfelter, Hosken, and Weinberg (2014)) and more recently (Hüschelrath and Müller (2014), Hüschelrath and Müller (2015), Israel, Keating, Rubinfeld, and Willig (2013) and Carlton, Israel, MacSwain, and Orlov (2017)). Most studies have estimated price increases, but some results are

[6]Park (2020) uses a model that includes capacity choices at one airport to address the effectiveness of slot divestitures.

sensitive to the chosen control group and time-window.[7] There are no retrospective analyses of how post-merger repositioning by rival firms or how this affects price changes in any industry.[8] We will discuss our own estimates of what happened to prices and repositioning after recent airline mergers, and we find that they are quite similar to the predictions of our model. This result contrasts with Peters (2006) who found that merger simulations with fixed products could not explain price changes after several mergers in the 1980s.

The second literature has estimated route-level entry or service choice models using airline data (Reiss and Spiller (1989) (RS), Berry (1992), Ciliberto and Tamer (2009), AH, BBCL and CMT). CMT and RS also assume full information and consider both service choice (RS) or market entry (CMT), and price competition. RS recognized "that entry introduces a selection bias in equations explaining fares or quantities" (p. S201) and they simplified their analysis by imposing symmetry and allowing for only one nonstop carrier, restrictions that we relax. We will discuss CMT's analysis in more detail in Section 7.

Sections 2, 3 and 4 detail our model, data and estimation procedure respectively. Section 5 presents the parameter estimates, model fit and implied selection. Section 6 presents the method and the results of our counterfactuals. Section 7 compares our approach to alternatives, including the estimation of a limited information model with fixed effects. Section 8 concludes. The Online Appendices provide some additional details of the data, estimation approach and analysis of alternative assumptions.

2 Model

We model carrier service choices at the route level, where a route is denoted by m and connects two airports (A and B).[9] The carriers playing the game in route m are denoted $j = 1, ..., J_m$. We will assume that, conditional on observables, all of the unobservables in our model are independent across routes.

[7] For example, Borenstein (1990), Werden, Joskow, and Johnson (1991), Morrison (1996) and Peters (2006) find different signs for short-run or long-run price effects after the 1986 TWA/Ozark and Northwest/Republic mergers.

[8] Hüschelrath and Müller (2015) provide an analysis of entry in airline routes but without tying entry to pre-merger market structures.

[9] Papers in the airline literature either consider airport-pair or city-pair markets, where the latter aggregates airports in the same city. The appropriate treatment depends on how willing consumers are to substitute between airports. A focus on airports is consistent with the fact that analysis of airline mergers has often focused on overlap at particular airports, such as Washington National airport, but we recognize that this assumption may be restrictive.

Figure 1: Timing of the Game

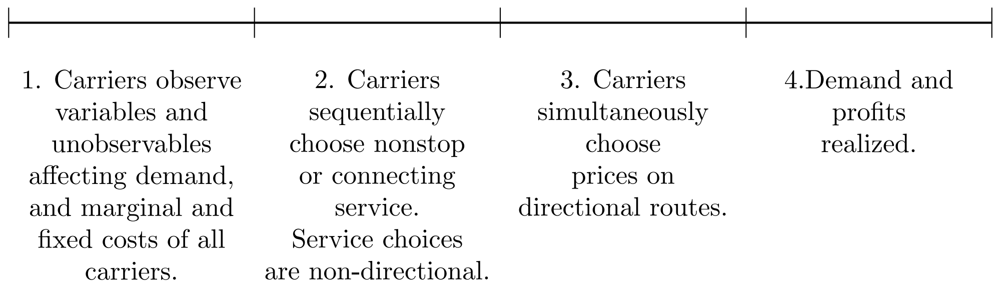

2.1 Overview

Figure 1 shows the timing of the game. The assumption that discrete service choices are made before prices are chosen is standard. Two assumptions are less standard. First, we assume that service choices are made sequentially. We will discuss this assumption in detail below. Second, we assume that carriers observe all demand and cost variables, for all carriers, before choosing their service types. This is our "full information" assumption. It is stronger than a "complete information" assumption, which simply requires firms have the same information.[10]

For each route, we model demand and price competition in two markets ($A \to B$, $B \to A$), one for passengers originating at each endpoint. We assume passengers make round-trips (a passenger making a one-way trip in the data will count as a half passenger). We use directional markets because a carrier's presence at the originating airport (we measure presence as the number of nonstop routes that a carrier serves from an airport, divided by the number of nonstop routes served by any carrier) has a strong correlation with its market share.[11]

2.2 Service Types

We assume that the carriers playing the game make a binary choice between providing nonstop service and providing connecting service via one of its hub airports. We define a carrier as providing nonstop service in our Q2 2006 data if a carrier has at least 65 nonstop flights on the route in each

[10] For example, Eizenberg (2014) and Wollmann (2018) assume that firms choose product portfolios knowing only the distributions from which demand and marginal cost unobservables will be drawn. This is consistent with complete information, but not full information.

[11] Differences in shares and prices are consistent with presence having a large effect on demand, which may reflect frequent-flyers preferring to travel on one carrier. For example, in a route fixed effects regression, a one standard deviation increase in the difference in a carrier's presence across the endpoints increases the difference in the carrier's directional market shares by 20% of the average directional share. Differences in origin presence also have significant, although smaller, effects on directional differences in average fares (Luttmann (2019)). Of course, one should interpret these types of regressions with caution when service choices are endogenous.

direction, and at least 50% of its passengers in the DB1 database are identified as not changing planes. Our main specification assumes that nonstop carriers only offer nonstop service, rather than nonstop and connecting options, thereby reducing the number of prices per carrier from four to two. This is a simplification, but when a carrier has nonstop service, a large majority of its passengers typically travel nonstop (for example, for 78.4% of nonstop observations, less than 10% of passengers make connections). We will show that our estimates are very similar when we assume that nonstop carriers also offer connecting service.

2.3 Demand

Demand is determined by a nested logit model, with all carriers in a single nest. For consumer k originating at endpoint A of route m, the indirect utility for a return-trip on carrier j is

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

where $p_{jm}^{A\to B}$ is the price charged by carrier j for a return trip from A to B. The first term represents carrier quality associated with j's service type (CON for connecting and NS for non-stop), $\beta_{jm}^{A\to B} = \beta_{jm}^{CON,A\to B} + \beta_{jm}^{NS} \times \mathcal{I}(j \text{ is nonstop})$ with $\beta_{jm}^{CON,A\to B} \sim N(X_{jm}^{CON}\beta_{CON}, \sigma_{CON}^2)$ and $\beta_{jm}^{NS} \sim TRN(X_{jm}^{NS}\beta_{NS}, \sigma_{NS}^2, 0, \infty)$, so that quality can depend on observed carrier-origin and route characteristics, and on a random component that is unobserved to the researcher. TRN denotes a truncated normal distribution and the lower truncation of β_{jm}^{NS} at zero implies that the nonstop service is always preferred to connecting service on the same carrier.[12] Note that estimation will require some more restrictive support conditions on draws (see Section 4 and Appendix C.2). As qualities are directional, each carrier has four β_{jm} draws ($\beta_{jm}^{NS,A\to B}$, $\beta_{jm}^{NS,B\to A}$, $\beta_{jm}^{CON,A\to B}$ and $\beta_{jm}^{CON,B\to A}$).

The random element of the carrier-specific quality draws (i.e., the parts not explained by the observables) will reflect differences in passenger tastes (for example, local loyalty developed from choosing the carrier historically), and also differences in carrier schedules and the types of planes that they use. A complete model would endogenize schedules and plane choices, but this would

[12] One could specify the model without restricting a carrier's incremental nonstop quality to be non-negative. For example, some passengers might rather change planes than fly on a very small plane nonstop. Our specification, which reflects the preferences of a representative consumer, assumes that this is not the case. This is consistent with the existing airline demand literature (Berry, Carnall, and Spiller (1996), Berry and Jia (2010) and Ciliberto and Williams (2014)) that has found that both business and leisure travelers have strong preferences for nonstop service. The restriction also helps to explain why some carriers serve almost all routes nonstop at their hubs. Section 4 notes that we make additional support restrictions to estimate the model.

require additional equations and unobservables. However, if carriers can predict the choices that other carriers will make, then this will be broadly consistent with our full information assumption.

The price and nesting parameters are assumed to be the same for all consumers on a given route, but we allow them to vary across routes, with $\alpha_m \sim TRN(X^{\alpha}\beta_{\alpha}, \sigma_{\alpha}^2, -\infty, 0)$, where X^{α} will include a measure of the importance of business travel on the route, and $\tau_m \sim N(\beta_{\tau}, \sigma_{\tau}^2)$. $\nu_m \sim N(0, \sigma_{RE}^2)$ is a route-specific random effect in demand, i.e., a demand shock that is common across carriers. $\varepsilon_{kjm}^{A \to B}$ is a standard logit error for consumer k and carrier j.

While we allow the price and nesting coefficients to vary across routes, demand has a nested logit, rather than a random coefficient structure, within each market. The nested logit model implies strong restrictions on cross-price elasticities, but it is convenient when many pricing games have to be solved to estimate the model and perform counterfactuals.[13]

2.4 Marginal Costs

Each carrier has a constant marginal cost draw for each type of service, $c_{jm} \sim TRN(X_{jm}^{MC}\beta_{MC}, \sigma_{MC}^2, 0, \infty)$, where $X_{jm}^{MC}\beta_{MC}$ allows costs to depend on the type of carrier, the type of service and the distance traveled. As passengers make round-trips, the marginal cost is non-directional, so each carrier has two marginal cost draws (c_{jm}^{NS} and c_{jm}^{CON}). The unobserved variation in marginal costs may reflect, for example, variation in a carrier's fuel efficiency on different routes (which will depend on plane type) and its cost of handling bags.

2.5 Fixed Costs and The Value of Connecting Traffic on Routes to Hubs

We assume that carriers have to pay a fixed cost, F_{jm}, to offer nonstop service on route m. This could include the opportunity cost of assigning gates and planes to a route, as well as airport gate rental and landing fees, which may vary in unobserved ways across routes and carriers. There is no fixed cost to providing connecting service. We assume that $F_{jm} \sim TRN(X_{jm}^F\beta_F, \sigma_F^2, 0, \infty)$ where X_{jm}^F includes a dummy for a slot-constrained airport where opportunity costs may be higher.

In the data, it is common for more than 60% of passengers on routes to or from a carrier's hub to be making connections. A model is only likely to be able to predict a hub carrier's service choices if it accounts for the size of these connecting passengers flows in some way. We take a relatively simple

[13] Berry, Carnall, and Spiller (1996), Berry and Jia (2010) and Ciliberto and Williams (2014) estimate demand models with mixtures of two types of customers, described as business and leisure travelers, with different price sensitivities and different preferences for nonstop service. Our approach assumes that we can adequately capture the effect of this type of heterogeneity by allowing the expected value of α and β_{jm}^{NS} to vary with our business travel index.

approach of assuming that a carrier's fixed costs can be offset by a linear function of three variables, which we will call "network variables": dummy variables for domestic and international hubs, and a third variable that is (the log of) a prediction of the total number of connecting passengers that a carrier will serve when it provides nonstop service on a route that involves a domestic hub (for non-hub routes, the variable is zero).[14] Appendix B.2 describes the model used to construct the prediction, which captures the geographic convenience of different connections on different routes, and it is estimated using data from one year prior to our estimation sample to reduce endogeneity concerns. Appendix D provides descriptive regressions showing that, together with market size, the variables included in our fixed cost specification can predict service choices quite accurately.

One might be concerned that a failure to model connecting traffic in more detail will make our counterfactuals less informative.[15] However, our counterfactuals are focused on whether, when two nonstop carriers merge, their connecting rivals will introduce nonstop service. While the pre-merger nonstop carriers are often serving their hubs, this is never true for the connecting rivals on the routes that we consider, and, as a result, their network variables are all zero. On the other hand, the merging carriers will be assumed to maintain nonstop service, which is what we observe in the data, so that changes in any of their fixed cost variables have no effect on our predictions.

2.6 Price Competition

Given service choices, carriers play static, simultaneous Bertrand Nash pricing games for passengers originating at each endpoint. Our assumptions of nested logit demand, constant marginal costs and single product firms imply that there will be unique equilibrium prices and directional variable profits, $\pi_{jm}^{A\to B}(s)$, given service choices, cost and quality draws (Mizuno (2003)). j's market-level variable profits are $\pi_{jm}(s) = \pi_{jm}^{A\to B}(s) + \pi_{jm}^{B\to A}(s)$, as service choices are assumed to be the same in both directions.

Our assumption that carriers only choose a single price in each direction abstracts away from how carriers sell tickets at many different prices because of price discrimination and revenue man-

[14] We use the log because the standard deviation of the variable in levels is very large. We require that the net fixed cost is non-negative as this reduces the range of the importance draws that we need to take. We show that this does not prevent us from accurately matching service choices at major hubs.

[15] It is not feasible to estimate a full information model that would capture a carrier's simultaneous choice of connecting prices across a very large number of routes or the correlations in costs across routes that connecting traffic might create. For example, there are over 6,000 directional domestic routes which Delta served via a change of planes Atlanta hub. When we simulate data from our estimated model, the implied margin on passengers that use connecting service by the named legacy carriers does not vary too much across markets (median \$94, with 50% of the predictions between \$83 and \$108), suggesting that an assumption of a constant average markup at the route-level should be fairly accurate.

agement incentives. There are no oligopoly revenue management models that it would be feasible to incorporate within the current model.[16]

2.7 Service Choices

In the first stage, carriers choose whether to commit to the fixed cost required for nonstop service, or to provide connecting service. Their realized profits in the full game are therefore $\pi_{jm}(s) - F_{jm} \times \mathcal{I}(j \text{ is nonstop in } m)$ where F_{jm} is a fixed cost draw associated with providing nonstop service. Our baseline specification assumes that carriers make their service choices sequentially in order of their average presence (see Section 2.1 for the definition) at the endpoints. This assumption guarantees a unique predicted outcome for the whole game. We will show that our estimates are robust to making weaker assumptions.

2.8 Solving the Game

Conditional on service choices, Nash equilibrium prices, shares and profits can be found by solving the system of pricing first-order conditions. One approach to finding equilibrium service choices would be to compute equilibrium profits for all combinations of service choices, and then to apply backwards induction to the branches of the extensive-form game tree. However, we can reduce computation by testing whether a carrier would make positive profits from nonstop service if all later movers were not in the game at all. If it would not, we know that the carrier will never choose nonstop service and we can delete branches where it would.[17] See Appendix C.1 for more discussion.

2.9 Full Information, Selection and Market Structure

Our full information assumption implies that service choices will be correlated with demand and marginal cost unobservables of all firms, leading to non-linear form of selection. Correlations between the unobservables could introduce additional non-linearities. Our baseline assumption is that, with the exception of the common route-level demand effect, unobservables are independent,

[16] Papers that have estimated revenue management models using airline data (Lazarev (2013) and Williams (2020)) have only considered monopoly markets. Carriers typically sell tickets at the same list prices in either direction, but the average realized prices may differ due to differences in demand. In our setting this outcome will be treated as if the carrier sets a different price in each direction.

[17] For example, suppose that the first moving carrier would have variable profits as a nonstop monopolist (i.e., with no other carriers in the game at all) that are lower than its fixed cost. This implies that it can never find nonstop service to be profitable on the route, and we can immediately eliminate one-half of the game tree of the full extensive form game.

although, as we note in footnote 28, observed covariates lead to quite strong correlations between a carrier's nonstop service quality and its costs of nonstop service. Our robustness checks allow for correlations in the unobservables and we will find the correlations to be small and statistically insignificant.

Full information and selection also have implications for market structure, which we investigate in Appendix A by comparing outcomes in full and limited information models with the same parameters. The probability that more than one carrier will be nonstop in equilibrium tends to be significantly lower under full information, because when a carrier expects to have a nonstop rival, it will expect that rival to be a stronger competitor, reducing its own profits from nonstop service, when the rival knows its own demand and marginal cost unobservables. An additional feature of a limited information model is that carriers will also frequently regret their service choices, once their unobservables are revealed. If unobservables are persistent and sunk costs are small, this feature makes it doubtful that market structures predicted by a limited information model, pre- or post-merger, will actually persist in the data. This provides an additional reason for believing that a full information model may provide more accurate predictions of what will happen to market structure over several periods.

3 Data and Summary Statistics

We estimate our model using a cross-section of publicly-available DB1 (a 10% sample of domestic itineraries) and T100 (records of flights between airports) data for the second quarter of 2006. We use relatively old data so that we can make predictions about subsequent mergers and avoid later years when carriers have been alleged to price cooperatively (Ciliberto and Williams (2014)). Appendix B provides additional detail and discussion. Tables 1 and 2 provide summary statistics.

Markets and Carriers. We use data for 2,028 airport-pair markets linking the 79 busiest US airports in the lower 48 states. Excluded routes include short routes and routes where nonstop service is limited by regulation. We model six named legacy[18] carriers, American Airlines, Continental Airlines, Delta Air Lines, Northwest Airlines, United Airlines and US Airways, and one named low-cost carrier (LCC), Southwest. We aggregate other ticketing carriers into composite "Other Legacy" (e.g., Alaska Airlines) and "Other LCC" (e.g., JetBlue and Frontier) carriers.

[18]Legacy carriers are carriers founded prior to deregulation in 1978, and they typically operate through hub-and-spoke networks. Our classification of carriers as LCCs follows Berry and Jia (2010).

Table 1: Summary Statistics for the Estimation Sample

	Numb. of Obs.	Mean	Std. Dev.	10^{th} pctile	90^{th} pctile
Market Variables					
Market Size (directional)	4,056	24,327	34,827	2,794	62,454
Num. of Carriers	2,028	3.98	1.74	2	6
Num. of Nonstop	2,028	0.67	0.83	0	2
Total Passengers (directional)	4,056	6,971	10,830	625	17,545
Nonstop Distance (miles, round-trip)	2,028	2,444	1,234	986	4,384
Business Index	2,028	0.41	0.09	0.30	0.52
Market-Carrier Variables					
Nonstop Indicator	8,065	0.17	0.37	0	1
Price (directional, round-trip $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
Indicator for Low Cost Carrier	8,065	0.22	0.41	0	1
$\geq$ 1 Endpoint is a Domestic Hub	8,065	0.13	0.33	0	1
$\geq$ 1 Endpoint is an International Hub	8,065	0.10	0.30	0	1
Connecting Distance (miles, round-trip)	7,270	3,161	1,370	1,486	4,996
Predicted Connecting Traffic (at domestic hubs)	1,036	8,664	7,940	2,347	52,726

Table 2: Distribution of Market Structures in the Estimation Sample

Number of Nonstop Competitors	Number of Sample Markets	Percentage of Sample Passengers	Average Number of Connecting Carriers
0	1,075	15.0%	3.98
1	614	33.6%	2.91
2	277	35.5%	2.07
3	60	15.2%	1.25
4	2	0.10%	0

We attribute tickets and flights to mainline ticketing carriers when they are operated by regional affiliates.

Service Types, Market Shares and Prices. We define the competitors on a route as carriers ticketing at least 20 DB1 passengers and with at least a 1% share of traffic. On average, there are four competitors, with as many as nine on long routes, such as Orlando-Seattle, with many plausible connections. We define a carrier as nonstop when it has 64 nonstop flights in each direction and 50% of passengers do not make connections, although the exact thresholds have little effect on the classification. The remaining competitors are classified as connecting. Most routes have no nonstop carriers, but routes with at least two nonstop carriers account for a majority of passenger trips. This type of route will be the focus of our counterfactuals. Most routes with multiple nonstop carriers connect large cities or hub airports, but non-hub pairs such as Boston-Raleigh and Columbus-Tampa also have nonstop duopolies.[19]

We measure a carrier's price as the average round-trip price in DB1. A carrier's market share is calculated as the total number of passengers that it carries in DB1, regardless of service type, divided by a measure of market size. We calculate market size as the prediction from a gravity model, which accounts for total endpoint enplanements and route distance (see Appendix B.1). This reduces unexplained heterogeneity in market shares across routes and explains service choices better than alternative measures, such as average endpoint city populations.

We assume that nonstop service has higher quality than connecting service. This is consistent with the results in the existing literature (e.g., Berry and Jia (2010) and Ciliberto and Williams (2014) who find that both nonstop and leisure travelers prefer nonstop service) and it is also consistent with the pattern that in our data, controlling for carrier fixed effects and competition, the average market share of a nonstop carrier is 18%, compared to 4.9% for a connecting carrier (even though connecting carriers with few passengers have been excluded), and nonstop prices are $43 higher than connecting fares.[20]

Exogenous Variables. Carrier presence is calculated using T100 data. Nonstop distance is measured as the great circle distance between two airports, and the distance for connecting service is measured as the distance via the carrier's closest connecting hub airport.[21] Appendix B.2 details which airports are domestic or international hubs and the construction of the connecting traffic variable. The business index variable, which approximates the proportion of business travelers on a route, is based on data provided by Severin Borenstein (Borenstein (2010)).

4 Estimation

We estimate the model parameters, $\Gamma = (\beta, \sigma)$, using a simulated method-of-moments estimator. In this section we outline the algorithm, the moments, identification and possible alternative im-

[19]If we had defined markets using city-pairs, rather than airport-pairs, there would be 192 nonstop duopolies (out of 1,533 city-pair markets), with 90 city-pair markets having three or more nonstop carriers.

[20]These estimates are from regressions of a carrier's weighted (across directions) average fare on a route on nonstop distance, carrier dummies, a dummy for whether the carrier provides nonstop service and interactions between whether a carrier provides nonstop service and the number of nonstop carriers on a route. Of course, these regressions do not account for the endogeneity of service choices, so these differences should not be given a strict causal interpretation.

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

Table 3: Moments Used in Estimation

Exogenous Variables (Z)	Market Level (y_m) Endogenous Outcomes 7 per market	Market-Carrier Level (y_{jm}) Endogenous Outcomes 5 per carrier	Row Total
Market-Level Variables (Z_m) (7 per market)	49	315	364
Carrier-Specific Variables (Z_{jm}) (up to 5 per carrier)	280	200	480
"Other Carrier"-Specific Variables (Z_{-jm}) (5 per "other carrier")	315	225	540
Column Total	644	740	1,384

Notes: Z_m = {constant, market size, market (nonstop) distance, business index, number of low-cost carriers, tourist dummy, slot constrained dummy.}
Z_{jm} = {presence at each endpoint airport, our measure of the carrier's connecting traffic if the route is served nonstop, connecting distance, international hub dummy} for named legacy carriers and for Southwest (except the international hub dummy). For the Other Legacy and Other LCC Carrier we use {presence at each endpoint airport, connecting distance} as we do not model their connecting traffic. Carrier-specific variables are interacted with all market-level outcomes and carrier-specific outcomes for the same carrier.
Z_{-jm} = {the average presence of other carriers at each endpoint airport, connecting passengers, connecting distance, and international hub dummy} for each other carrier (zero if that carrier is not present at all in the market).
y_m = {market level nonstop price (both directions), connecting price (both directions), sum of squared market shares (both directions), and the square of number of nonstop carriers}.
y_{jm} = {nonstop dummy, price (both directions), and market shares (both directions)} for each carrier.

plementations. Appendix C provides additional details, including evidence on the performance of the algorithm and its underlying assumptions.

4.1 Objective Function and Moments

The objective function is defined as

$$h(\Gamma)'Wh(\Gamma)$$

where W is a weighting matrix, and $h(\Gamma)$ is a vector of moments where each element has the form $\frac{1}{M}\sum_{m=1}^{m=M}\left(y_m^{data} - E_m\widehat{(y|\Gamma, X_m)}\right) Z_m$. y_m^{data} are observed outcomes and Z_m are a set of observed exogenous variables that serve as instruments. $E_m\widehat{(y|\Gamma, X_m)}$ are the predicted outcomes of the model for market m given the parameters Γ and observed variables X_m. We describe the moments that we use before describing how we compute $E_m\widehat{(y|\Gamma, X_m)}$.

Standard demand estimation with fixed product characteristics (e.g., Berry (1994)) uses moments that are based on the assumed orthogonality of a structural unobservable and instruments. However, the selection implied by the full information assumption implies that the structural errors for the service type that is chosen will not have mean zero and will be correlated with all of the exogenous variables in the model. Instead, we create moments using the fact that, for the true parameters, the expected value of the observed outcomes should match the expectation of predicted outcomes from the model.[22] The moments are summarized in Table 3. The outcomes include both market-carrier outcomes (e.g., Delta's price, its share and an indicator for whether it enters nonstop) and market/route outcomes (such as the sum of squared market shares, and the squared number of nonstop carriers). In principle, any function of the observed variables that are assumed to be exogenous can be used as instruments. The ones that we use can be broken into three groups: market-level variables (e.g., market size and the business index), market-carrier characteristics (e.g., endpoint presence, and distance of connecting service) and the characteristics of rival carriers (e.g., Delta's presence when we are looking at an outcome for a carrier other than Delta). One robustness check will use a subset of the instruments.

4.2 Computation of the Moments Using Importance Sampling.

A nested fixed point algorithm would re-compute $E_m(y|\Gamma, X_m)$, by resolving simulated games for each market, each time a parameter is changed. This would be computationally expensive and would require the minimization of a discontinuous objective function. We instead approximate $E_m(y|\Gamma, X_m)$ using importance sampling following Ackerberg (2009).

The idea is straightforward. Denoting a particular realization of all of the draws as θ_m,

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

where $y(\theta_m, X_m)$ is the unique equilibrium outcome given our baseline assumptions. This integral cannot be calculated analytically, but we can exploit the fact that

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

where $g(\theta_m|X_m)$ is an "importance density" chosen by the researcher.[23]

[22]Moments where outcomes are matched are usually used to estimate endogenous entry models (e.g., Berry (1992)), but here we are also applying them to prices and market shares.

[23]Ackerberg describes his approach as requiring a "change of variables". The change is implicit in the way we

This leads to a two-step estimation procedure. In the first step we take many draws, indexed by s, from densities $g(\theta_{ms}|X_m)$ and solve for the equilibrium outcome, $y(\theta_{ms}, X_m)$, for each of these draws. In the second step we estimate the parameters, approximating $E_m(y)$ using

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

where we only need to recalculate $f(\theta_{ms}|X_m, \Gamma)$ when the parameters change. The objective function is smooth because the $f(\theta_{ms}|X_m, \Gamma)$ densities are smooth in the parameters. We minimize the objective function using the `fminunc` function in MATLAB.

Appendix C details our selection of the parameters of the g density functions that we use in estimation and also our specification of the supports of the random variables (quality draws, costs, nesting and price parameters). As suggested by Ackerberg, the choice of g comes from initially estimating the model using gs that place weight on a broad set of draws.[24]

We form the W matrix by using the results from initially estimating the model using an identity weighting matrix. However, rather than using the inverse of the full covariance matrix, our final estimation uses a diagonal weighting matrix, with equal total weight on the groups of moments associated with price, share and service choice outcomes and, within each group, the weight on each moment is proportional to the reciprocal of the variance of that moment from previous estimates. We choose this approach because, with many moments relative to the number of observations (16,130 carrier-market-directions), the estimated covariances are likely to be inaccurate, and, in practice, some estimates are less stable if we use the full covariance matrix.

4.3 Identification.

As explained above, standard identification arguments for demand and marginal cost parameters will fail as selection implies that carrier demand and marginal cost residuals for chosen service types will neither have mean zero, nor be uncorrelated with exogenous demand and marginal cost variables. Identification therefore requires accounting for the exact form of selection implied by

have written down our model. For example, in a traditional entry model a firm's fixed cost might be written as $F_{i,m} = X_{i,m}\beta_F + u^F_{i,m}$, and a NFXP estimation routine would integrate over the distribution of the us. An importance sampling approach requires a change of variables by taking draws of $F_{i,m}$ rather than draws of $u^F_{i,m}$. This is consistent with how we wrote down the model in terms of random draws of costs (e.g., $F_{im} \sim TRN(X^F_{im}\beta_F, \sigma^2_F, 0, \infty)$) and qualities in Section 2.

[24] Unlike the choice of starting values, the chosen g will always matter for the exact values of the estimated parameters. However, we have found that alternative gs, or using additional rounds of estimation, leads to very similar results.

the full model. However, in our setting, the observed exogenous variables, including market size and the network variables affecting fixed costs, are able to predict the service choices of a large proportion of carrier-route observations almost perfectly (see Appendix D). This implies that, for these observations, there should be almost no selection on demand and marginal cost unobservables (i.e., the expected value of the unobservables should be close to zero, and they should be almost uncorrelated with the exogenous variables), in which case standard identification arguments should approximately apply.

CMT estimate the demand and marginal cost parameters by adjusting the standard demand and marginal cost moments to account for selection, with all of the observed exogenous variables (i.e., observed fixed cost shifters for a firm and its rivals, as well as demand and marginal cost shifters) as valid instruments.[25] We take a different approach, using moments conditions that directly match observed and predicted carrier and market price, share and service choice outcomes, but we can also use the observed exogenous variables as valid instruments.

Given consistent estimates of the demand and marginal cost parameters, identification of the fixed cost parameters follows from variation in carrier entry decisions with variation in variable profits, due to variation in market size, a carrier's own exogenous characteristics and the exogenous characteristics of rivals. Our carrier service choice moments are similar to those used in the literature on discrete choice games (e.g., Berry (1992)).

4.4 Alternative Estimation Methods for a Full Information Model.

While our importance sampling method provides one approach to estimating a full information model, alternatives are possible. CMT re-solve a large number of games every time a parameter is changed. For a given number of simulations, CMT's approach will be more efficient as importance sampling reduces efficiency. On the other hand, it leads to a discontinuous objective function and the computational burden that it creates requires simplifying the game.[26]

It has been suggested to us that alternative implementations might make our estimator more efficient. This is possible, but at least two suggested alternatives cannot be used. For example, Roberts and Sweeting (2013) use importance sampling within a simulated maximum likelihood estimator, which, all else equal, will tend to be more efficient. However, in their incomplete in-

[25] CMT allow for multiple equilibria so that, for some market structures, the moments take the form of inequalities, although there are moment equalities for outcomes where no firms or all firm enter. Given that we assume sequential entry, similar moments for our model would all be moment equalities.

[26] For example, CMT's model has up to six players and assumes demand and price competition are non-directional. We allow up to nine players and directional demand and pricing.

formation auction and bidding game, any possible outcome has positive likelihood for any set of unobservable draws (for auction characteristics), whereas, in our game, many outcomes would have a zero simulated likelihood even with many simulations. The Geweke–Hajivassiliou–Keane (GHK) estimator (for example, Keane (1994)) provides an efficient method for estimating models with multiple normally distributed errors using sequential conditioning. However, in our setting, the dependence of a firm's variable profits on the draws of every carrier in the market would make sequential conditioning infeasible.

5 Parameter Estimates

The first columns of Tables 4 and 5 present our baseline estimates. The coefficients are consistent with expected patterns. All else equal, consumers have a strong preference for nonstop service, legacy carriers and carriers with greater originating airport presence. Demand is less elastic on routes with more business travelers.[27] The average own price demand elasticity is -4.25, and the elasticity of demand for air travel (i.e., when all prices rise by the same proportion) is -1.3, close to the literature average reported by Gillen, Morrison, and Stewart (2003). For the average nonstop carrier, consumers' preference for nonstop service is $118 dollars.

LCCs have lower marginal costs, and costs increase with distance. To illustrate, consider the 3,000 mile round-trip Miami-Minneapolis route. For the named legacy carriers, the expected nonstop marginal cost is $345, compared to an average of $367 for (longer-distance) connecting service. Marginal costs for Southwest (and Other LCC) are lower and, for this route, Southwest's expected nonstop and connecting (via Chicago Midway) costs are almost identical ($303 and $298 respectively). On a non-hub route, the average nonstop fixed cost is $924,000, but, on domestic hub routes, the average fixed cost, once we offset the value of the network variables, is $399,000.

5.1 The Role of Unobservables

Accounting for selection on unobservables will be a key feature of our counterfactual analysis. It is therefore natural to look at the relative importance of observable covariates and unobservables in different parts of our model. To assess this, we simulate each market 20 times using the baseline coefficients, and compute how much of the variation (across carrier-market simulations) in a particular type of draw is accounted for by variation in the observed Xs. For example, the standard

[27] The expected price coefficient (α) for Dayton-Dallas-Fort Worth, which has the highest business index, is -0.34 compared to the cross-market average of -0.57.

Table 4: Parameter Estimates: Demand

				(1) Independent Unobservables	(2) Correlation Specific. 1	(3) Correlation Specific. 2	(4) Nonstop and Connecting
Route-Level Parameters							
Demand RE	S.D.	σ_{RE}	Constant	0.311	0.538	0.469	0.369
$[-2, 2]$				(0.138)	(0.151)	(0.122)	(0.135)
Nesting	Mean	β_τ	Constant	0.645	0.634	0.640	0.617
Parameter				(0.012)	(0.013)	(0.015)	(0.013)
$[0.5, 0.9]$	S.D.	σ_τ	Constant	0.042	0.005	0.050	0.020
				(0.010)	(0.010)	(0.008)	(0.009)
Price	Mean	β_α	Constant	-0.567	-0.542	-0.612	-0.602
Coefficient				(0.040)	(0.045)	(0.031)	(0.041)
(price in $100 units)			Business	0.349	0.189	0.435	0.382
$[-0.75, -0.15]$			Index	(0.110)	(0.118)	(0.088)	(0.113)
	S.D.	σ_α	Constant	0.015	0.043	0.013	0.035
				(0.010)	(0.011)	(0.013)	(0.010)
Carrier-Level Parameters							
Carrier	Mean	β_{CON}	Legacy	0.376	0.322	0.465	0.291
Connecting			Constant	(0.054)	(0.064)	(0.047)	(0.054))
Quality			LCC	0.237	0.336	0.150	0.223
$[-2, 10]$			Constant	(0.094)	(0.086)	(0.094)	(0.113)
			Presence	0.845	0.674	0.524	0.835
			at Origin	(0.130)	(0.125)	(0.127)	(0.196)
	S.D.	σ_{CON}	Constant	0.195	0.208	0.201	0.255
				(0.025)	(0.027)	(0.028)	(0.026)
Incremental	Mean	β_{NS}	Constant	0.258	0.192	0.560	0.519
Quality of				(0.235)	(0.214)	(0.221)	(0.181)
Nonstop			Distance	-0.025	-0.057	-0.009	-0.061
Service				(0.034)	(0.037)	(0.036)	(0.044)
$[0, 5]$			Business	0.247	0.841	-0.396	0.288
			Index	(0.494)	(0.455)	(0.479)	(0.372)
	S.D.	σ_{NS}	Constant	0.278	0.241	0.213	0.257
				(0.038)	(0.042)	(0.034)	(0.045)

Notes: 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 a pool of 2,000 draws) for each selected market. Distance is measured in thousands of miles. See Table 5 for estimates of the cost and covariance parameters. The supports of the random variables are indicated in square brackets. For example, the nesting parameter can lie between 0.5 and 0.9.

Table 5: Parameter Estimates: Marginal Costs, Fixed Costs, Network Effects and Covariances

				(1) Independent Unobservables	(2) Correlation Specific. 1	(3) Correlation Specific. 2	(4) Nonstop and Connecting
Carrier Marginal Costs ($100 units) [0, 6]	Mean	β_{MC}	Legacy Constant	1.802 (0.168)	1.350 (0.146)	1.847 (0.190)	1.389 (0.229)
			LCC Constant	1.383 (0.194)	0.961 (0.169)	1.344 (0.207)	1.100 (0.247)
			Conn. X Legacy	0.100 (0.229)	0.443 (0.211)	0.040 (0.251)	0.629 (0.295)
			Conn. X LCC	-0.165 (0.291)	0.288 (0.255)	0.140 (0.273)	0.388 (0.322)
			Conn. X Other Leg.	-0.270 (0.680)	-0.213 (0.166)	-0.228 (0.160)	0.051 (0.188)
			Conn. X Other LCC	0.124 (0.156)	0.046 (0.152)	-0.173 (0.167)	0.171 (0.168)
			Nonstop Distance	0.579 (0.117)	0.823 (0.101)	0.510 (0.128)	0.865 (0.155)
			Nonstop Distance^2	-0.010 (0.018)	-0.044 (0.016)	-0.001 (0.019)	-0.059 (0.023)
			Connecting Distance	0.681 (0.083)	0.661 (0.096)	0.675 (0.091)	0.524 (0.083)
			Connecting Distance^2	-0.028 (0.012)	-0.018 (0.013)	-0.026 (0.013)	0.000 (0.012)
	S.D.	σ_{MC}	Constant	0.164 (0.021)	0.191 (0.016)	0.143 (0.018)	0.148 (0.020)
Carrier Fixed Costs ($1m. units) [0, 5]	Mean	β_F	Legacy Constant	0.887 (0.061)	0.897 (0.056)	0.855 (0.075)	1.104
			LCC Constant	0.957 (0.109)	1.008 (0.118)	0.857 (0.100)	0.922 (0.124)
			Slot Const. Airport	0.568 (0.094)	0.424 (0.099)	0.514 (0.085)	0.411 (0.105)
	S.D.	σ_F	Constant	0.215 (0.035)	0.275 (0.029)	0.220 (0.030)	0.195 (0.033)
Carrier Network Variables (offset fixed costs)			Dom. Hub Dummy	-0.058 (0.127)	-0.302 (0.157)	-0.205 (0.193)	0.000 (0.212)
			Log $\widehat{\text{(Conn. Traff.)}}$	-0.871 (0.227)	-1.000 (0.207)	-0.602 (0.257)	-0.972 (0.287)
			Intl. Hub	-0.118 (0.120)	-0.144 (0.090)	-0.107 (0.093)	-0.261 (0.137)
Covariances	Incremental Nonstop Quality & Fixed Cost			-	0.012 (0.010)	0.018 (0.010)	-
	Connecting Quality & Connecting Marginal Cost			-	-	0.006 (0.007)	-

Notes: see notes below Table 4. The Log(Predicted Connecting Traffic) variable is zero for routes that do not involve a domestic hub, and for hub routes it is re-scaled with mean 0.52 and standard deviation 0.34. Supports are in square brackets.

Table 6: Model Fit: Predictions of Service Decisions at Raleigh-Durham

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

Notes: Predictions from the model calculated based on twenty simulation draws from each market from the relevant estimated distributions.

deviation of $F_{j,m}$ is \$301,912, and the standard deviation of $X_{i,m}^F\widehat{\beta_F}$ is \$259,481, so that unobserved heterogeneity provides only 14% of the variation. Similarly, unobserved heterogeneity accounts for only 3% of the variation in marginal costs and 15% of the variation in the price sensitivity of demand. However, it accounts for 26% and 34% of variation in connecting and nonstop carrier quality respectively, and our estimates also indicate that the variance in the unobserved route-level demand effect is quite large. These results suggest that accounting for selection on demand unobservables may be particularly important.

5.2 Model Fit

We use the 20 sets of draws to assess how well our model predicts observed service choices (discussed here) and variation in prices and market shares across service types (Appendix C.5). We correctly predict a carrier's service choice for 87.5% of draws (with standard error 1.1%), and for 82.6% (2.2%) of observations where a majority of our simulations predict a carrier will be nonstop, the carrier is nonstop in the data. We accurately predict carrier choices at hubs (Appendix Table C.2, e.g., Delta serves 96.5% of routes at Atlanta nonstop compared to a prediction of 92.5% (2.3%)) and non-hub airports. Table 6 illustrates the non-hub fit for routes with Raleigh-Durham (RDU) as an endpoint. The proportion of nonstop routes is served accurately for each carrier. The prediction is least accurate for United, as our simulations predict that United should serve Denver and San Francisco nonstop. United has launched nonstop service on both routes since 2006.

5.3 Robustness Checks

We now discuss what happens when we relax some of the assumptions imposed on our baseline estimates.

Correlations Between the Unobservables. Our baseline specification imposes that demand and cost unobservables are independent.[28] Columns (2) and (3) of Tables 4 and 5 present our estimates when we allow for correlations between the unobserved incremental quality of nonstop service and the fixed cost of providing nonstop service, and between connecting quality and connecting marginal costs. The estimated covariances are small, and only one of them is statistically significant at the 10% level.[29]

Nonstop Service Includes Connecting Service. Our baseline model assumes that a carrier that offers nonstop service only provides nonstop service. However, carriers often offer both nonstop and connecting service. Column (4) of Tables 4 and 5 presents our estimates when we assume nonstop carriers provide connecting service and set four prices on each route.[30] For the vast majority of coefficients are close to their baseline values.

Reduction in the Number of Moments. Our baseline estimation uses 1,384 moments, which is large relative to the sample size, creating the possibility of bias. Appendix C.6 presents estimates, an analysis of fit and some example counterfactual results using only the 740 carrier-specific moments. All of the results are similar to the baseline.

Relaxing the Known, Sequential Order Assumption. Our baseline estimates assume that service choice decisions are made in a known sequential order, which guarantees a unique equilibrium and point identification.[31] In contrast, Ciliberto and Tamer (2009), Eizenberg (2014)

[28] However, the coefficients on observed covariates lead to strong correlations between demand and costs. For example, based on the 20 sets of draws used to examine model fit, the correlation between a carrier's nonstop quality and its fixed costs of nonstop service is -0.56.

[29] When we allow unrestricted correlations we find that the objective function has multiple local minima. We have used a grid search on the covariance parameters to confirm that values close to zero minimize the objective function. We find no obvious improvements in fit when we allow for correlations. CMT estimate a more flexible covariance structure, and they find that some of the covariances are large. This may reflect the fact that the unobservables in their model have to pick up the large differences in demand and costs between between nonstop and connecting carriers, whereas we explicitly model these differences.

[30] The difference in this model is that, when solving our simulated games, we solve for four prices and quantities (two types of service in each direction) for carriers that choose to be nonstop. These are then matched, as carrier-direction-nonstop service type averages, to the moments from the data that are calculated in the same way. Note that this assumes that in the data all nonstop carriers provide connecting service. In the DB1 data 31% nonstop carriers have no connecting passengers, but in many cases, at least for legacy carriers, this seems likely to reflect passengers' choices rather than no availability of a connecting product.

[31] On the other hand, the parameters can be point identified even if some equilibria are not unique, because an outcome such as "no firms are nonstop" will always be unique. In our data the most common outcome is that no firms are nonstop.

and Wollmann (2018) estimate using inequalities assuming that moves are simultaneous and that any pure strategy equilibrium can be played. Ciliberto and Tamer (2009) report that there are multiple equilibria in over 95% of simulations of their airline entry game. We now explain why our equilibrium selection assumption is not particularly restrictive in our setting.

A discrete choice game can support multiple equilibrium outcomes only if at least two players do not have dominant strategies. Without estimating a model, one indicator that many carriers are likely to have dominant service choice strategies is that simple probit regressions using market and a carrier's own characteristics are able to predict carriers' observed choices with very high probability. As discussed in Appendix D, two or more carriers have predicted nonstop probabilities between 0.1 and 0.9 in only 15% of markets. In contrast, if we use the same covariates to predict whether carriers that serve the endpoints provide either connecting or nonstop service (which is similar to Ciliberto and Tamer's outcome of interest), the corresponding figure is 96%. When we simulate data from our estimated baseline model, we find that multiple equilibrium outcomes can be supported (either by pure strategies in a simultaneous move game, or a sequential game with any order of moves) for only 1.6% of simulated games. Appendix C.7 also reports the parameters that minimize the objective function when we extend our estimation methodology to use moment inequalities to allow for simultaneous moves or unknown orders. These parameters are very similar to the baseline estimates, and the percentage of games that support more than one equilibrium outcome is almost identical.

6 Merger Counterfactuals

We now present our counterfactuals. Section 6.1 describes the mergers that we consider. Section 6.2 describes the assumptions and results of standard merger simulations with fixed products, which leads us to focus on routes where both merging parties are nonstop. Section 6.3 describes the assumptions, method and results when we allow repositioning by rivals, including a comparison to what happened after three completed mergers. Section 6.4 uses the model to examine how far two alternative remedies can constrain post-merger market power. We use the baseline demand and cost estimates throughout.

6.1 The Set of Mergers and Routes Considered

We examine the three legacy mergers (Delta/Northwest (2008), United/Continental (2010), American/US Airways (2013)) completed after 2006 and a United/US Airways merger that was proposed in 2000, but abandoned when the Department of Justice opposed it.[32] We do not consider the consummated merger between Southwest and Airtran, because Airtran is part of our composite "Other LCC".

The United/Continental and American/US Airways were allowed to proceed on the condition that the parties divested slots and other facilities at major airports to low-cost carriers. In the United/US Airways merger, the parties proposed a remedy where a third carrier, American, would commit to provide nonstop service for ten years on several routes where the merging parties were nonstop duopolists. The Department of Justice did not accept the remedy on the grounds it was insufficient to restore pre-merger competition.[33] Section 6.4 will discuss both types of remedy.[34]

6.2 Merger Counterfactuals with Fixed Products

We first present results from a set of standard merger counterfactuals (e.g., Nevo (2000)) that do not allow for repositioning. We make the following assumptions when we resolve for equilibrium prices.

Assumption 1 (Merger Counterfactuals with Fixed Products) *We assume*

1. The products owned by the merging parties are eliminated and replaced by a single product of the merged carrier ("Newco"). We consider two alternative assumptions about Newco's demand and costs:

(a) ("baseline assumption") Newco has the quality and marginal cost of the merging party with the higher average endpoint presence before the merger when both parties have the same service

[32] BBCL consider the United/US Airways, Delta/Northwest and United/Continental mergers, and CMT consider the American/US Airways merger.

[33] R. Hewitt Pate, Deputy Assistant Attorney General, discussed the merger and the remedy in a speech, "International Aviation Alliances: Market Turmoil and the Future of Airline Competition", on 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 nonstop service on over 30 routes. We concluded that ... American Airlines' promise to fly five routes on a nonstop basis [was] 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."

[34] Park (2020) uses a model that allows for the allocation of slots across routes to provide a detailed analysis of the effectiveness of the American/US Airways divestiture.

Table 7: Predicted Effects of Mergers with Service Choices Held Fixed

	Delta/Northwest		United/Continental		American/US Airways		United/US Airways		Average	
	Data	Post	Data	Post	Data	Post	Data	Post	Data	Post
1. Alternative Market Structures & Merger Eliminates Lower Presence Carrier										
Merging parties nonstop duopolists	\$566.39	\$593.20	\$503.75	\$556.17	\$459.13	\$521.15	\$479.32	\$549.49	\$481.40	\$541.25
	2 routes		4 routes		11 routes		7 routes		24 routes	
Merging parties nonstop with nonstop rivals	\$351.26	\$382.04	\$438.08	\$464.98	\$363.11	\$404.84	\$350.02	\$378.15	\$368.70	\$402.08
	2 routes		4 routes		10 routes		10 routes		26 routes	
One party nonstop, other connecting	\$472.99	\$524.67	\$502.60	\$513.29	\$447.95	\$478.95	\$443.30	\$462.53	\$458.02	\$486.40
	91 routes		59 routes		158 routes		163 routes		471 routes	
Both parties connecting	\$433.26	\$444.63	\$487.04	\$486.86	\$464.20	\$457.77	\$484.25	\$479.62	\$466.00	\$465.97
	479 routes		334 routes		471 routes		521 routes		1,805 routes	
2. Merging Parties Nonstop Duopolists & Merger Eliminates Lower Presence Carrier										
Numb. of Routes	2 routes		4 routes		11 routes		7 routes		24 routes	
Merging Carrier Prices	\$566.39	\$593.20 (1.17)	\$503.75	\$556.17 (1.57)	\$459.13	\$521.15 (1.79)	\$479.32	\$549.49 (1.44)	\$481.40	\$541.25 (1.55)
Combined Mkt. Share	18.9%	14.3% (0.1)	29.1%	21.7% (0.0)	26.9%	18.8% (0.1)	20.8%	12.9% (0.0)	24.8%	17.2% (0.0)
Combined Variable Profit (\$k)	1,287 (47)	1,235 (43)	6,112 (208)	6,225 (208)	4,336 (201)	4,006 (187)	4,023 (124)	3,907 (104)	4,287 (164)	4,116 (151)
Combined Fixed Costs (\$k)	689 (108)	194 (55)	1,143 (144)	536 (66)	1,564 (90)	621 (47)	1,248 (100)	503 (61)	1,321 (96)	537 (51)
Average Rival Prices	\$235.90	\$237.48 (0.10)	\$455.57	\$457.29 (0.03)	\$400.44	\$404.18 (0.14)	\$280.19	\$282.09 (0.15)	\$360.85	\$363.53 (0.08)
Change in Cons. Surp. Per Traveler	−\$51.36 (1.64)		−\$62.81 (2.21)		−\$64.06 (2.89)		−\$80.03 (2.19)		−\$67.04 (2.47)	
3. Merging Parties Nonstop Duopolists & Merged Firm Receives Highest Qualities and Lowest Costs of the Merging Parties										
Merging Carrier Prices	\$566.39	\$598.77 (1.14)	\$503.75	\$558.63 (1.58)	\$459.13	\$513.34 (2.18)	\$479.32	\$537.60 (1.57)	\$481.40	\$535.09 (1.75)
Combined Mkt. Share	18.9%	15.1% (0.0)	29.1%	21.9% (0.0)	26.9%	19.9% (0.01)	20.8%	14.3% (0.0)	24.8%	18.2% (0.0)
Combined Variable Profit (\$k)	1,287 (47)	1,343 (414)	6,112 (208)	6,359 (209)	4,336 (201)	4,418 (162)	4,023 (124)	4,566 (103)	4,287 (164)	4,528 (141)
Combined Fixed Costs (\$k)	689 (108)	162 (39)	1143 (144)	375 (57)	1,564 (90)	565 (39)	1,248 (100)	446 (64)	1,321 (96)	465 (47)
Rival Carrier Prices	\$235.90	\$237.87 (0.06)	\$455.57	\$457.14 (0.03)	\$400.44	\$403.46 (0.14)	\$280.19	\$281.36 (0.04)	\$360.85	\$362.91 (0.08)
Change in Cons. Surp. Per Traveler	−\$44.88 (1.72)		−\$60.83 (2.18)		−\$56.94 (3.21)		−\$68.25 (2.19)		−\$59.74 (2.63)	

Notes: for routes where the merging carriers are nonstop duopolists, standard errors for measures not directly observed in the data are reported in parentheses, and the share, fixed cost and profit numbers are for the merging carriers combined. Prices averaged across directions, and pre-merger prices are averages across carriers. For other pre-merger market structures, the table shows the number of affected routes, and merging carrier prices with no standard errors. All calculations assume that the price and nesting parameters have their expected values for each market. Each merger is considered separately, not cumulatively.

type, and otherwise it has the quality and marginal cost of the nonstop party.

(b) ("best case assumption") Newco has the higher quality and the lower marginal cost of the merging parties.

2. the nesting and price parameters are equal to their expected values for each market given observed market characteristics and the baseline parameter estimates.

3. the products of the non-merging carriers remain the same, with the same service type and the same demand and marginal cost draws as in the data.

The second assumption reduces the computational burden, although we have verified that the results are almost identical if we relax it, consistent with the small estimated standard deviations of the price coefficient and the nesting parameter. The best case assumption parallels the assumption of Li and Zhang (2015) concerning valuations and hauling costs in the context of timber auctions, and it tends to increase the profitability of the merger relative to the baseline assumption. We can implement the merger simulation by inverting the demand of each carrier (for its offered service type) and its marginal cost from the observed price and market share data in each market, given the demand system parameters and the Nash pricing assumption.

6.2.1 Comparison of Merger Price Effects Across Different Market Structures.

The first panel of Table 7 reports pre-merger and post-merger prices for four different groups of markets under the baseline assumption about the merging parties. The reported pre-merger price is the average of the share-weighted average prices for the merging carriers across directions, and the post-merger price is the average for Newco across directions. To save space, we do not report standard errors for the post-merger prices, although, as can be seen when we report them for nonstop duopoly routes in other parts of the table, they are small. Looking at the cross-merger average prices in the final column, the price increases are largest when the merging parties are both nonstop and they are the only nonstop carriers (nonstop duopolies) (12.4%), and, next, when they are both nonstop and there is at least one nonstop rival (9.1%). Price increases are smaller when one (6.1%) or both (no change) of the merging carriers are connecting.[35] Given these results, we will focus the rest of our analysis on the 50 routes where both the merging carriers are nonstop

[35] When two connecting carriers merge, consumer surplus still falls because the disappearance of a choice, but the drop is much smaller than for nonstop duopolies (average \$5 per pre-merger traveler, compared to \$67 for nonstop duopolists).

prior to the merger, and particularly the 24 routes where they are nonstop duopolists.[36]

Some have suggested that our focus on routes where the merging firms are already nonstop means that we are missing the possibility that the merger might expand the set of routes where the merging parties offer nonstop service. We have two responses to this concern. First, the focus of merger analysis is always on the markets where the merger raises anticompetitive concerns, and while the agencies may credit "out of market" efficiencies when deciding whether to allow a merger to proceed, they are not required to do so (Rybnicek and Wright (2014)). Therefore, an analysis that focuses on where anticompetitive effects are most likely is appropriate. Second, even though the simulations in CMT and BBCL both suggest that the merging firms will increase the set of routes that they serve (nonstop in the case of BBCL), empirically we do not observe significant expansions of nonstop service, at least in the short-run. For example, three quarters before their merger, United and Continental served a total of 258 of our sample routes nonstop, and the merged carrier served 259 routes nonstop seven quarters after the merger.[37]

6.2.2 Detailed Analysis for Nonstop Duopoly Markets.

The second and third panels of Table 7 report more detailed results for the 24 nonstop duopoly markets, under the baseline (second panel) and best case (third panel) assumptions. We report standard errors, and all of the predictions are precise. On two routes (one for Delta/Northwest and one for United/US Airways) there are no connecting rivals, so the mergers create monopoly.

In the baseline case, all of the considered mergers are predicted to raise the merging carriers' average prices by between 5% and 15%, and the parties' market shares are predicted to fall by between 25% and 30%, reflecting both the price increases and the elimination of a product. The next rows allow us to examine the profitability of the merger. Although the decision to merge is taken at the network level, not the route level, the predicted profitability of a merger can be used

[36]One might be concerned that the 24 routes are not representative of nonstop duopoly routes as a whole, and so could give a misleading impression of what would happen in nonstop duopoly markets if different mergers were completed. Appendix B.3 provides a comparison of these 24 routes with the remaining sample routes where two non-merging legacy carriers were nonstop duopolists and routes where Southwest was one of two nonstop duopolists. The table shows that, compared to other legacy duopoly routes, the 24 routes have more connecting carriers with higher connecting market shares. Therefore, our finding that mergers create significant market power for our routes would likely be even clearer if we considered mergers on these other routes. On the other hand, prices are significantly lower where Southwest is nonstop. We also find (Appendix B.4) that we observe a different pattern of post-merger outcomes on nonstop duopoly routes when Southwest merged with Airtran.

[37]The equivalent numbers for the Delta/Northwest merger are 336 and 296, but with a declining trend before the merger, and for American/US Airways they are 291 and 302, with a slight upward trend before the merger. Therefore, in none of the consummated mergers that we look at does the merger have a clear or immediate effect on the number of nonstop routes served by the merging parties.

to understand whether the assumptions are plausible. While the elimination of a product and the lack of synergies means that variable profits tend to fall, total profits tend to increase because a fixed cost of nonstop service is eliminated. Connecting rivals are predicted to raise their prices, although the increases are small. Consumer surplus, measured in dollars per pre-merger traveler, tends to fall quite significantly.[38]

The directional changes in the predictions when we make the best case assumption are intuitive, with the merging parties losing fewer passengers, and their profits increasing. However, the magnitudes of the changes are quite small, because the higher presence carrier, which survives under the baseline assumption, will usually be the carrier with the highest quality, and our estimates imply that nonstop legacy carriers have very similar marginal costs, so choosing one rather than the other makes little difference. For example, we predict that the merged firm's prices increase by 11.2%, rather than 12.4% under the baseline assumption.

6.3 Mergers Counterfactuals where Rivals Can Reposition

We now analyze counterfactuals where we allow rivals to change their service choices after the merger. Our discussion of our assumptions and method will assume that we are considering markets where the merging carriers are nonstop duopolists, but we will also present results where there are additional nonstop rivals.

6.3.1 Assumptions.

Assumption 2 (Merger Counterfactuals with Repositioning) *We assume*

1. the nonstop products owned by the merging parties are eliminated and replaced by a single nonstop product of the merged carrier ("Newco"). We consider two alternative assumptions:

(a) ("baseline assumption") Newco has the quality and marginal cost of the merging party with the higher average endpoint presence before the merger .

(b) ("best case assumption") Newco has the higher quality and the lower marginal cost of the merging parties.

2. the nesting and price parameters are equal to their expected values for each market given observed market characteristics and the baseline parameter estimates.

3. the non-merging carriers have the same quality and cost draws for both types of service as they

[38] We measure consumer surplus per pre-merger traveler because the markets considered vary quite dramatically in size, and our definition of market size is imperfect.

do before the merger, which should, therefore, be consistent with their pre-merger service choices. They choose their service type in the same sequential order as before the merger, knowing that Newco will be nonstop. We assume that no additional carriers, that do not provide some type of service in the data, can enter.

These assumptions follow the assumptions made in the fixed product case as closely as possible, except for allowing connecting rivals to reposition. It is worth asking, however, why assuming that qualities and costs remain the same is a reasonable assumption for either type of merger simulation. An empirical justification is that, when we apply our estimated demand model to panel data on prices and market shares, the implied carrier demand and marginal cost unobservables are highly persistent. Persistence is also consistent with our assumption that carriers will know the values of the unobservables.[39]

We measure persistence by examining the correlation between unobservables using a regression where the implied demand or marginal cost unobservable in the second quarter of 2006 (our estimation period) is regressed on a constant and the carrier's unobservable in the second quarter of 2005. The unobservables are backed out using the expected values of the price and nesting coefficients in the demand system and the pricing first-order conditions, under the assumption that the same demand system is appropriate in both years. The serial correlation coefficient for demand unobservables for carriers that are nonstop in both quarters is 0.638 (standard error 0.023) for a specification without market fixed effects, and 1.007 (0.042) when we include market fixed effects to control for differences in the level of demand across markets. For connecting carriers, the average coefficients are lower (0.410 and 0.690, respectively). We also observe persistence for marginal cost unobservables. For nonstop carriers, the serial coefficients are 0.889 (s.e. 0.014) and 0.802 (0.028) without and with market fixed effects, and 0.798 (0.008) and 0.419 (0.015) for connecting carriers.[40] Carriers' service choices are also highly persistent, consistent with persistence of demand, marginal cost and fixed cost unobservables.[41] Of course, persistence in service choices could also be explained by the addition of nonstop service requiring a large sunk cost, even if fixed costs and

[39]Of course, a richer model could allow for different components of the unobservables to be persistent and others to be unanticipated innovations, and this could also rationalize the empirical persistence that we observe. Estimating a model that combines service type and pricing choices with this type of flexibility is beyond the current literature.

[40]The lower correlations for connecting carriers are probably due to sampling error in the DB1 data, causing small connecting shares to be measured noisily. We have also estimated specifications where implied residuals from earlier years are used as instruments. The resulting correlation coefficients for both demand and marginal costs are between 0.9 and 1.25 for connecting carriers.

[41]We have identified all cases where the named carriers added nonstop service, other than through mergers, after Q1 2001 but before 2006, and then followed their service choices over subsequent quarters. On average, these carriers maintained nonstop service for 27 consecutive quarters.

variable profits are not persistent. However, the existence of large sunk costs is inconsistent with the fact that carriers serve many smaller routes nonstop on a seasonal basis and the fact that they have responded to short-run demand spikes in 2020 by offering nonstop service temporarily on some routes.[42]

We view the assumption that no other carriers can enter the route as more restrictive, and we will consider one analysis where we allow for an additional competitor. However, we view the assumption as being reasonable, given the computational costs of the alternative (see footnote 44), because, over the period from the first quarter of 2005 to the first quarter of 2008, 86.2% of carriers that began nonstop service on a route offered connecting service in the previous quarter, and, after mergers affecting nonstop carriers, three-quarters of rivals that began nonstop service were previously connecting carriers.

6.3.2 Method: The Simulation of Conditional Distributions.

To calculate equilibrium post-merger service choices, we need to infer the qualities and costs that rival carriers would have if they changed their service choices (i.e., qualities and costs that are not observed in the data). We do this by forming "conditional distributions" of qualities, including the route-level demand effect, and costs which are consistent with both the estimates and pre-merger service choices.[43] The natural interpretation of the conditional distributions is they are posteriors with the estimated distributions treated as priors, with the conditioning being on the service choices, prices and market shares observed in the data.

We form the conditional distributions using the following steps. We first specify a discrete set of possible values for the route-level demand effect. For each value, we calculate the qualities and marginal costs implied by observed prices and market shares for the chosen service types. We then take draws of the remaining random components of the model (carrier qualities and marginal costs for the non-chosen service types, and the fixed costs of nonstop service) from their estimated distributions and, for each set of draws, we keep (accept) those draws which would support the observed service choices as an equilibrium outcome. We weight the accepted draws using the estimated densities of the route-level demand effect and the implied carrier qualities and costs for chosen service types, to form the conditional joint distribution of the route-level demand effect,

[42] Wall Street Journal article, October 6, 2020, "How Are Legacy Airlines Surviving Covid-19? By Borrowing from the Low-Cost Playbook".

[43] We note that one could also choose to condition, for example, on the profitability of the merger, and to also form a conditional distribution for the merger synergy.

Figure 2: Selection of Marginal Conditional Distributions for Philadelphia-San Francisco

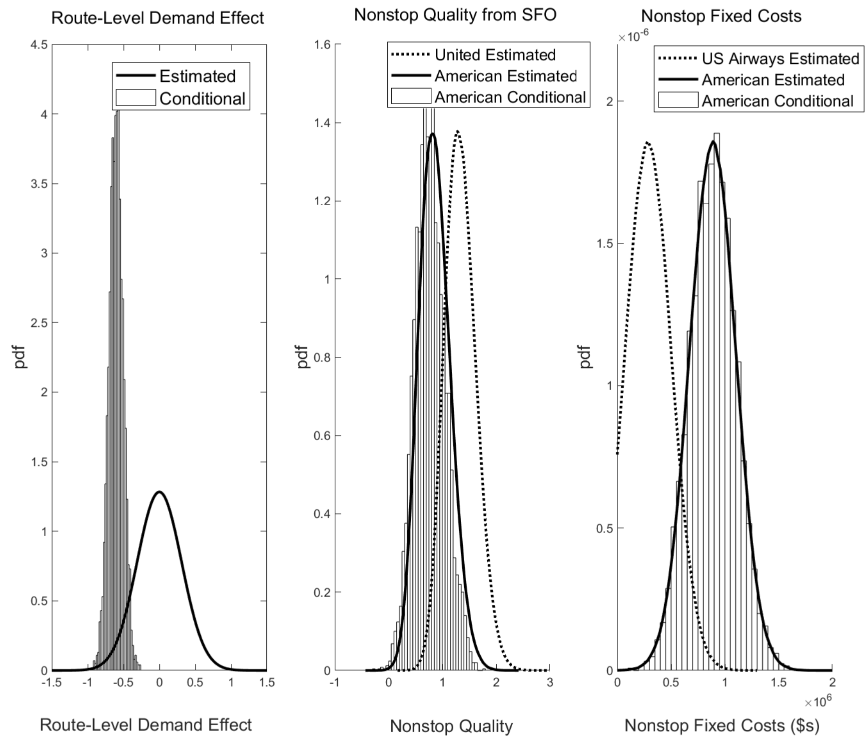

carrier qualities, marginal costs and fixed costs for all of the carriers in the market.[44]

We illustrate the effect of conditioning in Figure 2 for the Philadelphia (PHL)-San Francisco (SFO) route, one of the nonstop duopoly routes affected by the United/US Airways merger. The solid line in the left panel shows the estimated density of the route-level demand effect, while the histogram shows the simulated marginal conditional density (50,000 simulation draws). The conditional distribution has a lower mean, reflecting the fact that the number of observed passengers, across all carriers, is relatively low (combined market share is 28.3%, averaged across directions) given market size and the observed covariates. As a comparison, the mean of the conditional distribution of the demand effect for Las Vegas-Miami, where combined market shares equal 42.5%, is 0.5.

Nonstop quality is the sum of a carrier's connecting quality and the incremental quality of nonstop service. The lines in the middle panel show the density of nonstop quality for passengers originating at SFO for United and American based on the estimates (i.e., not conditioning on what is

[44]The acceptance rate drops when more unobserved variables are added to the model or we add additional players. For example, if we considered a model where carriers choose between {do not enter, enter connecting, enter nonstop}, as in Li, Mazur, Roberts, and Sweeting (2015), we would have to calculate the conditional distribution of four qualities, two marginal costs and two fixed costs for each carrier that does not enter.

observed). United's expected quality is higher, because of its high presence at SFO. The histogram shows the conditional density for American's nonstop quality. This distribution is similar, but with a slightly lower mean, than the distribution implied by the estimates. The intuition is that given observed shares and prices and the likely value of the random effect, we need to shift our expectation of American's nonstop quality down, by a small amount, to explain why it chooses connecting service. The posteriors for carrier quality would be identical to those implied by the estimates in a limited information model. The third panel shows the densities for the fixed cost of nonstop service for American and US Airways. US Airways has a lower expected effective fixed cost because of its domestic and international hubs at PHL. The estimated and conditional distributions for American's fixed costs look essentially identical.

We note that the step of forming conditional distributions of qualities and costs that are consistent with observed outcomes in the data, is an important difference between our counterfactual analysis and the one provided by CMT. CMT estimate a full information entry and price competition model, and then simulate market outcomes, and perform counterfactuals by implementing an American/US Airways merger using the simulated data, where all of the demand and cost draws are known to the researcher, including for carriers that do not enter. Their counterfactuals tell us what happens in simulated markets, not actual markets. Instead, we take the approach that an agency has to take when analyzing a proposed merger, where they need to predict what will happen in one or more actual markets, with known competitors and observed market shares, if a merger takes place. Our approach is also the only way to make sure that our predictions are consistent with those from a standard fixed product merger simulation, which predicts changes from observed (not simulated) prices and market shares.

6.3.3 Results.

Predicted Effects of a United/US Airways Merger. We start by presenting our results for the United/US Airways merger on four routes where the merging carriers were nonstop duopolists and American was a connecting competitor, so that we can connect our discussion to our calculation of conditional distributions for the Philadelphia-San Francisco route, and our discussion of the proposed service remedy where these are the affected routes.

The upper panel in Table 8 presents results under our baseline merger assumption. We expect the merged firm's prices to increase by 8.3% on these routes if service types are held fixed with a significant predicted decline in consumer surplus. Such predictions would usually lead an antitrust

Table 8: Predicted Effects of United/US Airways Merger in Four Nonstop Duopoly Markets Allowing Repositioning By Rivals

Counterfactual	Pre-Merger United/US Airways Price	Exp. Numb. of Rivals Launching Nonstop Service: American	Others	Post-Merger Merged Carrier Price	Change in Consumer Surplus
Baseline Merger Assumption					
1. Service Types Fixed	$531.97	-	-	$576.18 (0.77)	-$48.07 (1.69)
Allow Rival Service Changes					
Connecting Rivals Nonstop Quality and Costs Drawn from:					
2. Conditional Distns.	$531.97	0.035 (0.024)	0.063 (0.055)	$573.37 (2.36)	-$42.96 (4.88)
3. Estimated Distns.	$531.97	0.190 (0.062)	0.325 (0.106)	$559.56 (5.08)	-$16.22 (11.22)
4. Average of Merging Parties	$531.97	0.678 (0.062)	1.915 (0.106)	$531.79 (1.97)	+$62.36 (8.43)
Best Case Merger Assumption					
5. Service Types Fixed	$531.97	-	-	$562.82 (0.94)	-$37.76 (1.77)
Allow Rival Service Changes					
Connecting Rivals Nonstop Quality and Costs:					
6. Conditional Distns.	$531.97	0.020 (0.015)	0.043 (0.042)	$560.73 (1.96)	-$33.80 (4.00)

Notes: predictions with endogenous service choices are averages from 1,000 draws from the appropriate distributions. Pre-merger prices are averages across the merging parties. Implementation of rows 3 and 4 explained in the text. Standard errors reported in parentheses.

agency to challenge a merger unless offsetting synergies or repositioning are likely.

The second row reports our predictions when we allow rivals' service types to change after the merger, using 1,000 draws from the conditional distributions for each market. The expected number of rivals initiating nonstop service, a measure of the likelihood of repositioning, is small, leading to the result that, in expectation, the merged carrier's price increases by $41 (7.8%). We also find that the merger is, on average, profitable for the merging firm despite the repositioning that takes place, with its profits increasing by an average of $279k (s.e. $78k) per route. We will return to rows 3 and 4 below.

The lower panel (rows 5 and 6) report the results under the best case assumption. As Newco now tends to have slightly lower marginal costs, the predicted price increase is smaller, but there is also less repositioning by rivals. The merger now appears to be much more profitable, raising profits by an average of $1.1m (s.e. $85k) per route.

To understand the predictions for the baseline assumption, Table 9 provides more detail for the PHL-SFO route. On this route, United, the lower average presence carrier that is assumed to be eliminated by the merger, has a particularly large share, so that the merger potentially creates a significant opportunity for a connecting carrier that launches nonstop service. The results in the table use 5,000 draws so we can measure the probability of different outcomes accurately.

Table 9: Predictions for the Philadelphia-San Francisco Market Allowing Repositioning By Rivals Following a United/US Airways Merger

Carrier (pre-merger service type, price and share)	No Service Changes 3,267/5,000 Draws		American Nonstop 570/5,000 Draws		Delta Nonstop 483/5,000 Draws	
	Price	Share	Price	Share	Price	Share
US Airways/Newco	\$691.53	15.4%	\$661.67	14.1%	\$661.46	14.0%
(NS, \$649.74, 13.0%)	(1.17)	(0.0)	(0.66)	(0.1)	(1.64)	(0.1)
United	-	-	-	-	-	-
(NS, \$613.54, 12.1%)						
American	\$478.98	1.2%	\$554.64	8.1%	\$477.30	0.8%
(CON, \$476.52, 0.5%)	(0.05)	(0.0)	(9.70)	(0.4)	(0.07)	(0.0)
Delta	\$666.89	0.6%	\$666.08	0.4%	\$550.98	7.9%
(CON, \$665.77, 0.3%)	(0.03)	(0.0)	(0.04)	(0.0)	(8.74)	(0.5)
Northwest	\$307.35	3.5%	\$302.51	2.4%	\$302.47	2.4%
(CON, \$300.60, 1.9%)	(0.18)	(0.0)	(0.23)	(0.1)	(0.23)	(0.1)
Other LCC	\$377.27	1.1%	\$375.82	0.7%	\$375.80	0.7%
(CON, \$375.27, 0.6%)	(0.06)	(0.0)	(0.07)	(0.0)	(0.07)	(0.0)

Notes: predictions are averages from 5,000 draws from the conditional distributions. Standard errors in parentheses based on the same bootstrap estimates used for the parameter estimates. The merger assumed to eliminate United (lower presence carrier). NS denotes nonstop and CON denotes connecting pre-merger.

For two-thirds of the draws, no connecting rival launches nonstop service, and the merged carrier's price increases by 9.5% (from the pre-merger average) and its market share falls by 38%. The non-merging carriers, with small connecting shares pre-merger, increase their prices slightly and double their combined market share. Reflecting the loss of a large carrier, consumer surplus falls by an average of \$72.91 per pre-merger traveler.

The remaining columns show what happens when one of American or Delta launch nonstop service, which are the most common outcomes involving repositioning (for 0.9% of draws more than one rival launches nonstop service). The increased competition reduces (but does not eliminate) the equilibrium price increase for US Airways, and the new nonstop carrier usually has a market share that is significantly smaller than United's prior to the merger, causing consumer surplus to fall by around \$30 per pre-merger traveler in both cases. Repositioning by rivals, when it happens, does tend to make the merger unprofitable for this route: for example, the merged firm's profits fall by \$920k when American becomes nonstop.[45]

[45] We have also calculated what happens under the best case assumption. In this case, there is no repositioning

This route provides an example where there can be multiple equilibrium outcomes in the counterfactuals depending on timing assumptions about service choices. For example, there are 27 (out of 5,000) draws where either American launching nonstop service or Delta launching nonstop service (but not both) are equilibrium outcomes. However, the different outcomes typically have very similar welfare implications. For example, the average within-draw-across-outcome standard deviation in the predicted US Airways price is $3.

Predicted Effects Using Alternative Assumptions About Rival Qualities. Rows 3-4 of the upper panel of Table 8 shows what happens if we make assumptions about rivals' qualities and costs that may not be consistent with their pre-merger service choices. We make the baseline assumptions about the merger.[46]

Row 3 uses new draws from the estimated (i.e., not conditional) cost and incremental nonstop quality distributions for the nonstop qualities and costs of the connecting carriers. We therefore account for differences in the observed characteristics of the connecting carriers, but do not account for the additional information in pre-merger service choices. Row 4 assumes that if any connecting rival becomes nonstop then it would have the average quality and marginal costs of the merging nonstop carriers and draw its fixed cost from a distribution that has a mean equal to the average of the means for the merging carriers. This approach ignores observable differences between carriers.[47] In both cases, we continue to draw the route-level demand effect from its conditional distribution and we use the qualities and marginal costs for observed service types that are implied by observed prices and market shares, so that we can isolate the effects that arise from making alternative assumptions about how competitive connecting rivals will be if they launch nonstop service.[48]

Compared to our preferred results using the conditional distributions, the estimated distributions imply it is more likely that rivals will launch nonstop service (the expected number of nonstop launches is 0.52, rather than 0.1), leading to a smaller expected price increase, and a smaller, statistically insignificant, decrease in consumer surplus of $16.22 per pre-merger traveler (s.e. $11.22).

---

for 78% of draws (rather than 65%), the US Airways price increases by an average of 4.3% (rather than 6.4%) when there is no repositioning and the merger is only marginal unprofitable when repositioning occurs (for example, profits fall by $106k when American becomes nonstop rather than falling by $920k).

[46] The results are similar if we make the best case assumption about the merger: for example, the expected number of carriers launching nonstop service are 0.46 (row 3) and 2.4 (row 4), rather than 0.52 and 2.6.

[47] This assumption might be viewed as consistent with the logic of the District Court's decision in *Waste Management* which considered only whether rivals would face entry barriers higher than those that had been faced by the merging parties.

[48] A rationale for using the conditional distribution of the route-level demand effect is that we include this component of the model to address thew fact that our market size measure may be imperfect. The parties and the agencies would likely be able to construct a better measure in a merger investigation.

These results also imply that the merger is likely to be unprofitable: average profits fall by \$105k (s.e. \$150k), compared to the \$279k (s.e. \$78k) increase using the conditional distributions.

Assuming that connecting carriers can offer nonstop service on similar terms to the merging parties leads to a prediction that, on average, 2.6 of them would launch nonstop service[49] and that, because consumers prefer nonstop service, consumer surplus is predicted to increase after the merger. However, if we use the same assumption to solve for equilibrium outcomes *before the merger*, we often predict that several connecting carriers should have chosen to offer nonstop service (e.g., American's probability of launching nonstop service would be 0.6 pre-merger), which is inconsistent with the observed data. This illustrates the importance of considering whether assumptions about the post-merger competitiveness of repositioning firms, or new entrants, are consistent with their pre-merger choices.

Predicted Effects of Completed Legacy Mergers on Nonstop Duopoly Routes. The upper panel of Table 10 summarizes our baseline merger assumption predictions for repositioning and post-merger prices for the 17 nonstop duopoly routes affected by the consummated mergers, under our different assumptions about the nonstop quality and costs of connecting carriers.

The qualitative patterns are very similar to Table 8, although magnitudes vary across mergers reflecting differences in conditions across routes. When we use our preferred conditional distributions, we expect 0.18 rivals to launch nonstop service on each affected route, and the merged carriers' prices are predicted to increase by an average of just under 10%, which is only 2 percentage points smaller than if service types are held fixed. Using the estimated distributions we predict more than three times as much repositioning by rivals and smaller, although still economically significant, price increases.[50] If we assume that connecting carriers could provide nonstop service with similar quality and costs to the merging parties, we predict that, on average, the mergers would have no anti-competitive effects.

6.3.4 Comparing Predictions to What Happened After Legacy Mergers on Nonstop Duopoly Routes.

It is natural to compare these predicted changes to what we observe actually happening after these mergers, albeit with the caveat that market conditions may have changed between 2006, the year

[49] If we assumed that connecting carriers would be similar to the eliminated carrier, rather than the average of the merging carriers, we would expect 1.5 of them to launch nonstop service.

[50] Under the best case merger assumption, we predict two-and-a-half times as much repositioning using the estimated distributions, so that the comparisons we make below to repositioning in the data still hold.

Table 10: Predicted Price and Service Changes for Subsequent Completed Mergers on Routes where Merging Parties are Nonstop Duopolists (Baseline Assumption)

	Delta/Northwest		United/Continental		American/US Airways		Average for Completed Mergers	
	Price	Exp. Numb. New Nonstop	Price	Exp. Numb. New Nonstop	Price	Exp. Numb. New Nonstop	Price	Exp. Numb. New Nonstop
Pre-merger	\$566.39	-	\$503.75	-	\$459.13	-	\$482.25	-
Post-Merger								
Service Types Fixed	\$593.20	-	\$556.17	-	\$521.15	-	\$537.86	-
Allow Rival Service Changes								
Connecting Rivals Nonstop Quality and Costs Drawn from:								
Conditional Distributions	\$590.34	0.07	\$547.65	0.14	\$511.33	0.21	\$529.17	0.18
Estimated Distributions	\$584.20	0.19	\$534.08	0.35	\$488.45	0.73	\$510.45	0.57
Average of Merging Parties	\$573.83	0.93	\$454.36	2.62	\$460.25	2.10	\$472.23	2.08
Number of Routes	2		4		11		17	

Table 11: Predicted Price and Service Changes Where Merging Parties and at Least One Rival are Nonstop

	Delta/Northwest		United/Continental		American/US Airways		United/US Airways		Average	
	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals	Price	Δ in # Of NS Rivals
Pre-merger	\$351.26	-	\$438.08	-	\$363.11	-	\$350.02	-	\$377.51	-
Post-Merger										
Service Types Fixed	\$382.04	-	\$464.98	-	\$404.84	-	\$378.15	-	\$412.27	-
Allow Rival Service Changes										
Connecting Rivals Nonstop Quality and Costs Drawn from:										
Conditional Distributions	\$378.90	0.16	\$464.86	0.01	\$404.41	0.03	\$377.24	0.05	\$411.07	0.06
Estimated Distributions	\$386.40	-0.51	\$466.18	-0.03	\$403.55	-0.27	\$375.17	-0.11	\$413.33	-0.28
Average of Merging Parties	\$374.37	0.66	\$455.64	0.61	\$398.85	-0.03	\$367.68	0.48	\$404.95	0.34
Number of Routes	2		4		10		10		26	

Notes: see notes to Table 8. All predictions make the baseline merger assumption and, when service types are endogenous, use 1,000 draws from the relevant distribution. Standard errors not reported. In the case where we assume that connecting rivals would have the same nonstop quality as the merging nonstop parties, we use the observed nonstop quality and marginal costs for the nonstop rival(s), and draw its (their) connecting qualities and marginal costs, and fixed costs, from the estimated distributions.

of our analysis, and the year that the mergers were consummated.

Appendix B.4 uses panel data to estimate what happened to rivals' service choices, prices and shares after the three completed legacy mergers, on routes where the carriers were nonstop duopolists prior to the merger. We summarize the findings here. On the nonstop duopoly routes, the merged carrier always maintained nonstop service. Within two years of the merger closing (the Department of Transportation explicitly used two years when considering repositioning (Keyes (1987)), a rival launched nonstop service on no routes, out of five, for Delta/Northwest, one route, out of five, for United/Continental and three routes, out of six, for American/US Airways.[51] There were two additional nonstop launches in the third year following these mergers. The Appendix also presents analyses of changes in the prices and market shares of the merging firms on routes where the merging firms were nonstop duopolists for three years before the merger, using a comparison set of routes where one of the parties was nonstop and the other was either absent or a connecting carrier with a small share.[52] On routes where no rivals initiated nonstop service, we find that the merged carrier increased its prices by an average of 10%, with its number of local passengers (i.e., those only flying the route itself) falling by almost 30%. On routes where rival nonstop service was launched, the merged carriers' prices did not rise, although they did lose market share, presumably reflecting the new competition. These patterns suggest that rivals tend not to launch nonstop service because they would not be competitive, rather than because the merged carrier enjoys large synergies.

These patterns are broadly consistent with the predictions of our model when we use draws from the conditional distributions of qualities and costs for the rival carriers. In particular, our analysis predicts that, on average, 0.18 rival carriers will initiate nonstop service, compared with 0.25 in the data, and that prices will increase by around 12% when there is no repositioning, compared with 11% in the data. It is also the case that we observe the most nonstop launches after the American/US Airways merger, consistent with our prediction. Our predictions of changes in merging carrier market shares when there is no repositioning are also close to the data. While the numbers of mergers and routes are too small to claim that the close match proves that our approach is correct, we view the match as at least encouraging. It stands in contrast to the conclusion of Peters (2006) that fixed product merger simulations poorly predict outcome changes after airline

[51]There is no overlap in the routes across these mergers.

[52]Estimated price changes may be affected by using different control groups or time windows, as suggested by the contrasting results of Hüschelrath and Müller (2015) and Carlton, Israel, MacSwain, and Orlov (2017) for recent mergers.

mergers.

6.3.5 Predictions for Markets with Additional Nonstop Rivals Pre-Merger.

Merger simulations with fixed products also indicate that prices would rise significantly on routes where the merging parties are nonstop but have at least one nonstop rival (there is one route with two nonstop rivals). Table 11 presents our predictions for these routes. When we simulate counterfactuals allowing for repositioning, we assume that the merged firm will be nonstop and make the same assumptions about connecting rivals that we made for nonstop duopoly routes. However, we also now endogenize the service choice of the nonstop rival(s). The nonstop quality and marginal costs of this type of carrier are observed, but we need to make assumptions about the quality and marginal costs of its connecting service, and its fixed costs of providing nonstop service.

When we use conditional distributions, we predict that the nonstop rival(s) will always continue to provide nonstop service and that connecting carriers will rarely introduce nonstop service. As a result, predicted price changes are almost identical to those where service types are assumed fixed. This is consistent with our earlier results. However, differences to our earlier results emerge for the other assumptions, because it becomes likely that the nonstop rival, which is usually quite an effective nonstop competitor, may cease nonstop service and this type of repositioning can lead to price increases. For example, a nonstop rival ceases nonstop service for around one-third of simulations in the results reported in the final ("Average of Merging Parties") row of the table. As a result, we now predict significant price increases under all three approaches, and the largest predicted prices increases and the greatest probability of post-merger nonstop monopoly are when we use the estimated distributions. Therefore while the intuition that the conditional distributions will tend to predict the largest prices increases when nonstop duopolists merge is fairly clear, there are additional nuances for other market structures that are relevant for merger analysis.

6.4 Remedies.

Remedies are often negotiated when only a small part of a transaction is likely to have anticompetitive effects. The agencies have a well-known preference for structural remedies, such as divestitures, but, in some circumstances, they also accept behavioral remedies or remedies that involve some

Table 12: Predicted Effects of the American Service Remedy in United/US Airways Merger

Service Change Considered	Pre-Merger United/US Airways Price	Exp. Numb. of Rivals Launching Nonstop Service: American	Exp. Numb. of Rivals Launching Nonstop Service: Other Rivals	Post-Merger Merged Carrier Price	Change in Consumer Surplus
No Remedy					
1. Service Types Fixed Fixed	$531.97	-	-	$577.72	-$48.07
2. Allow Rival Service Changes (Condit. Distns.)	$531.97	0.035	0.063	$573.37	-$42.96
American Nonstop Remedy					
3. Allow Rival Service Changes (Condit. Distns.)	$531.97	1	0.030	$566.34	-$31.29

Notes: see notes to Table 8. The merger is assumed to eliminate the party with the lowest presence on the route. Consumer surplus changes measured per pre-merger traveler. For American, the expected number of rivals launching nonstop service is the probability that American launches nonstop service. Standard errors not reported.

ongoing relationship between the merging firm and third parties.[53] We use our model to consider, in a stylized way, the effectiveness of two different types of remedies that have been proposed or used in airline mergers.

The Service Remedy Proposed in the United/US Airways Merger. The results presented so far suggest that when rivals launch nonstop service, the merged carrier can only increase prices by a small amount. This might be interpreted as implying that the remedy proposed in the United/US Airways merger, where American would guarantee to initiate nonstop service on routes where the parties were nonstop duopolists (see footnote 33), so that the number of nonstop carriers would not have changed, would have been effective. However, this logic implicitly assumes that American's nonstop service would constrain the merged carrier's prices even when it is unprofitable.[54]

The first two rows of Table 12 repeat the results from Table 8 for the four routes where United and US Airways were nonstop and American was a connecting competitor. The third row repeats the analysis under the remedy so that, whatever its draws from the conditional distribution, American is nonstop and other carriers then make their service choices taking this into account. We see that the effect of the remedy on expected post-merger prices is small. The insignificance of American as a nonstop competitor when its nonstop service is not profitable is also illustrated by how other rival carriers' service decisions are largely unaffected by the remedy.

[53] See September 2020 Department of Justice "Merger Remedies Manual" (https://www.justice.gov/atr/page/file/1312416/download, accessed November 11, 2020).

[54] The parties did not claim that nonstop service on the affected routes would be profitable for American: instead the attraction for American was that it would receive a package of assets on the East Coast if the merger was completed.

Figure 3: Distribution of American Incremental Profits (in $00s) from Nonstop Service on PHL-SFO and the Predicted Increase in the Merged Carrier's Price if American Launches Nonstop Service (Relative to Pre-Merger Average Prices) Given American's Profitability. The grey area marks the interquartile range of price outcomes.

(a) Full Information (b) Limited Information

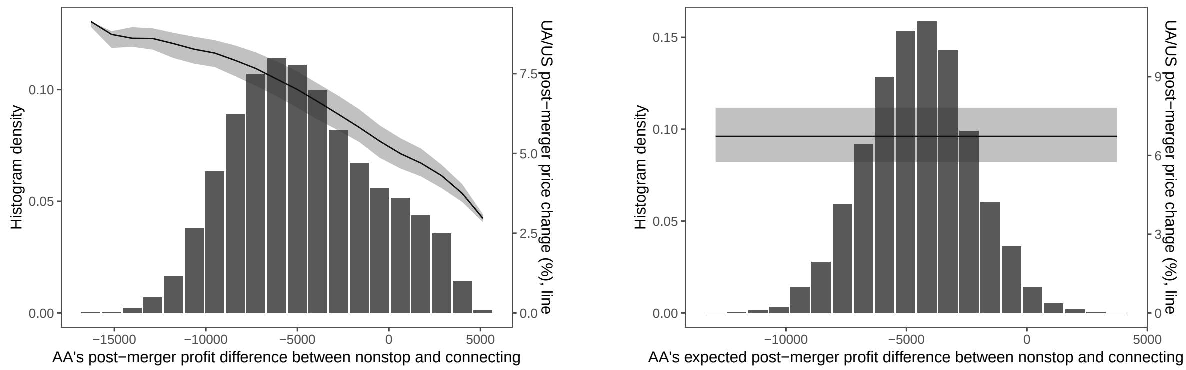

Figure 3(a) provides additional insight into what happens. The histogram shows the distribution of the difference between nonstop and connecting profits for American on the PHL-SFO route. For simplicity, we draw the figure assuming that American knows no other connecting carriers will launch nonstop service. The line on the figure shows the median simulated post-merger price increase for US Airways (relative to the average of United's and US Airways's pre-merger prices) when we force American to provide nonstop service given this level of profitability (the shaded area indicates the interquartile range generated by our simulations). There is a monotonic relationship between American's profitability and its effectiveness at reducing increases in the US Airways's prices, and there is only a significant constraining effect on those prices when nonstop service is at least close to being profitable for American.

To illustrate the effects of our assumption that demand and cost shocks are known when making service choices ("full information"), Figure 3(b) shows the same figure assuming that American has no information about its quality or marginal cost unobservables when making its service choice (for comparability, we assume American does know its fixed costs and the qualities and costs of other carriers). The variance of the (expected) profit distribution is reduced, as it now reflects only the distribution of fixed costs. As fixed costs will not affect the prices that carriers set, there is no link between the level of profit that American expects when it launches nonstop service and how much this will constrain the market power of the merging carriers.

Table 13: Predicted Effects of a Remedy When an Additional Other Low-Cost Carrier is Added as a Competitor on Nonstop Duopoly Routes

		Post-Merger Predictions with Repositioning			
				Pre-Merger Rivals	
	Pre-Merger	Pre-Merger Rivals		+ Addn. LCC	
Merger	Price	Exp. New NS	Price	Exp. New NS	Price
Delta/Northwest (2 routes)	$ 566.39	0.07	$590.34	0.05	$556.69
United/Continental (4 routes)	$503.75	0.14	$547.65	0.14	$530.03
American/US Airways (11 routes)	$459.13	0.21	$511.33	0.25	$492.24
All United/US Airways (7 routes)	$479.32	0.08	$546.74	0.32	$496.18
United/US Airways with AA connecting (4 routes)	$531.97	0.10	$573.37	0.14	$564.66
Average (24 routes)	$481.40	0.15	$ 534.30	0.24	$505.03

Notes: see notes to Table 8. The additional LCC carrier receives unconditional draws from the estimated distributions, has the characteristics of the average "other LCC" carrier (e.g., presence 0.17 at both endpoints), and is assumed to make its service choice last in the sequential move order. The merger is assumed to eliminate the party with the lowest presence on the route. Standard errors are not reported.

The Effect of Adding An Additional Low-Cost Competitor. We consider the effects of introducing an additional low-cost competitor as an alternative remedy. The Department of Justice allowed the United/Continental and American/US Airways mergers to proceed when the parties agreed to divest slots and gates to low-cost carriers at major airports.[55] The aim of these divestitures was to increase competition at the affected airports. While our model does not formally include slots, we can use it to ask the question of whether the addition of a low-cost carrier as a competitor would offset the anticompetitive effects of a merger.[56] As the new carrier was not on the route prior the merger, we assume that its quality and cost draws are not selected (i.e., they are new draws from the estimated distributions), and that it takes on the observed characteristics of the average "other LCC" carrier in the data. We then repeat our conditional distribution counterfactuals for routes where the merging parties were nonstop duopolists, assuming that the new carrier is last in the sequential order.

Table 13 compares the predictions of price and the number of new nonstop carriers (in total) when we add the new competitor to our baseline predictions for the nonstop duopoly routes. The pattern in the results varies across the mergers, reflecting differences in market structure. For example, on the two routes where the merged firms are the only competitors, adding the new carrier has a significant pro-competitive effect. However, the addition of a new rival, by reducing profitability of the remaining carriers, can actually lead to fewer carriers initiating nonstop service than in the baseline. Averaging across the 24 routes, the remedy reduces, but does not eliminate, the expected post-merger price increase (the average increase is 5% rather than 11%).

For the four United/US Airways routes where American offers connecting service, we can compare the effectiveness of the two remedies. We see that they are roughly equally effective in the sense that the expected post-merger prices are similar ($564.66 with an additional LCC competitor compared to $566.34 with the service remedy). We also investigated what would happen if the additional LCC is a stronger competitor, by increasing the assumed presence of the new carrier from 0.17 to 0.5 (which would be equivalent to the new carrier establishing some type of focus airport presence at both endpoints). In this case, the new carrier is more likely to add nonstop service, and the merged carrier's expected price increase is smaller, but still economically significant. For

[55] The settlement in the American/US Airways case also required divestitures of slots at Washington Reagan and New York LaGuardia, and of ground facilities at seven airports. The settlement in United/Continental required divestitures at Newark.

[56] Our stylized analysis will miss the fact that the additional LCC and the merging parties will need to choose how to allocate their scarce slots across routes. Park (2020) explicitly includes this type of slot allocation decision for a single carrier at a single airport.

example, on the four United/US Airways routes where American was a competitor the expected post-merger price is \$557.70. Therefore, we conclude that neither remedy is necessarily effective at preventing anticompetitive harm from the merger on nonstop duopoly routes. On the other hand, when slot divestitures bring additional competitors into major airports, they may create pro-competitive effects on a large number of additional routes, so that, if consumer surplus was aggregated across routes, the harmful effects of the merger might be eliminated to a much greater extent. In contrast, the service remedy is only likely to have pro-competitive effects on routes where the remedy applies.

7 Comparison to Alternative Approaches

In this section, we compare our approach to several alternatives.

7.1 Limited Information Static Models

RS, CMT and our paper assume that demand and marginal cost unobservables are known when service or entry choices are made, whereas the rest of the recent literature that estimates two-stage static models (Draganska, Mazzeo, and Seim (2009) Eizenberg (2014), Wollmann (2018) and Fan and Yang (2020a)) assume that only their distributions are known ("limited information"), so that the researcher and the firms have the same information about demand and costs in the first stage. This allows demand and marginal cost to be estimated separately from the discrete choice game. Assuming that a static model is to be used, there are three reasons why we believe that it is appropriate to assume full information in our setting.

The first reason is that demand and marginal cost unobservables, particularly for nonstop service, are quite persistent (see evidence discussed in Section 6.3.1). Persistence, combined with our view that the unobservables are likely to reflect differences in local tastes and carrier operations that other carriers are likely to understand, makes it plausible that carriers will be able to predict the unobservables quite accurately. The second reason is that limited information implies that firms may make choices that turn out not to be profitable (in the Appendix A example 48% of carriers choosing nonstop service would have made higher profits with connecting service). Given the persistence of the unobservables, it makes sense to focus on repositioning choices that are actually profitable when trying to predict whether post-merger repositioning will prevent consumer harm over a number of periods.

The third reason is that, while it is less convenient to estimate, there can be non-trivial differences in the predictions of limited and full information models. This is illustrated using a detailed example in Appendix A. We first show that the models can predict qualitatively different outcomes and reactions to a merger when we use the same parameters. We then consider what happens if we estimate the parameters of a limited information model using data generated from a full information model. In our example, we allow our estimated demand model to include a market fixed effect to capture cross-market differences in demand, as this is a common practice in papers that assume limited information (for example, Aguirregabiria and Ho (2012), Fan and Yang (2020b)). The failure to account for selection leads to significant bias in the demand and fixed cost parameters, and, even though the estimated parameters allow the limited information model to match average service choices in the data, the estimated model overpredicts post-merger repositioning and underpredicts the expected post-merger price increase.

7.2 An Alternative Full Information Static Model

CMT and our paper both model airline competition in a full information framework. However, the papers differ in several respects.

First, we model carriers' service choices rather than their route entry decisions. Our choice reflects the fact that the existing literature (Berry and Jia (2010), Ciliberto and Williams (2014)) has identified that passengers have a strong preference for nonstop travel, so that a merger that leads to a nonstop monopoly is likely to create significant market power, unless a rival initiates nonstop service. In contrast, when any carrier that flies one passenger per day is defined as an entrant, the marginal entrant is likely to have minimal competitive effect. Our focus also has two practical implications. Our model has more unobservables (each carrier has two demand and one marginal cost unobservable for each service type), which potentially increases the computational burden of estimation and counterfactuals, and our observable variables have more explanatory power (Appendix D), which has implications for the fit of our model and the existence of multiple equilibria (see discussion in Section 5.3).

Second, as discussed in Section 4, we use different estimation algorithms. CMT's approach involves resolving simulated discrete choice games for each guess of the parameters, resulting in a non-smooth objective function. We use the importance sampling approach, solving games in advance of estimation, and then re-weighting the simulated outcomes, resulting in a smooth objective function. For a given number of simulations, the CMT approach should be more efficient but it implies a larger computational burden, particularly when global optimization techniques have to be used.[57] A detailed comparison of the performance of our algorithms would be an interesting topic for future research.

Third, we develop an algorithm for accounting for selection in the data, and we show how it affects predicted post-merger changes from observed prices. Our approach is consistent with how merger analysis is actually conducted. In contrast, as explained in Section 6.3.2, CMT's counterfactuals examine how a merger would affect outcomes in simulated markets, so that the researcher knows every unobservable, and there is no selection problem to deal with.

7.3 Dynamic Models of Entry and Repositioning

AH and BBCL use dynamic models to understand carriers' nonstop segment entry and exit choices, allowing the set of carriers that compete on a route to change over time. BBCL use their model to consider the effects of mergers, comparing changes in concentration from a static analysis, which treats service choices as fixed, and a dynamic analysis which allows service choices to evolve for ten years.[58] Our analysis with endogenous product positioning lies between these two extremes: we take the set of nonstop and connecting competitors as fixed, and investigate whether connecting carriers will upgrade their service to nonstop. As BBCL note, our analysis is therefore complementary to theirs, and it is likely to be most useful for understanding positioning changes in a shorter-time window.

A dynamic model can provide insights into the "timely" criterion in the Guidelines, but, in general, dynamic models are rarely sufficiently tractable or transparent to be used in antitrust analysis of a merger. For example, BBCL's analysis of route service choices does not provide predictions of price effects and, as they do not resolve their model, their predictions do not allow for post-merger changes in firms' equilibrium strategy functions. AH do resolve their model, but, as in Sweeting (2013), which considers repositioning in the broadcast radio industry, resolving requires imposing additional simplifications and approximations.[59]

[57] Of course, the lower computational burden could allow more simulations to be used in the importance sampling approach.

[58] This analysis leads BBCL to conclude that the loss of competition from a United/US Airways merger would have been offset after ten years, whereas this would not have been the case after the Delta/Northwest merger.

[59] AH make limited information assumptions so that demand and marginal costs can be estimated separately. Sweeting slightly relaxes the limited information assumptions by assuming that it is innovations in product quality that are unknown when station format choices are made.

8 Conclusions

We have developed a model of endogenous service choices and price competition in airline markets, assuming that carriers have full information about demand and marginal costs when they make their service choices. In this framework, carriers will tend to choose the service type in which they are most competitive, and this naturally has implications for how likely they will be to change their service types in response to a change in their competitive environment, such as when two rivals merge. While it is unlikely to be the right assumption for all industries, we believe the full information assumption is the natural one to use when trying to predict product repositioning by experienced market participants in an environment where demand and marginal cost unobservables are persistent, and when trying to understand whether repositioning will sustainably limit market power after a merger.

We make two contributions. First, we show how a full information model can be estimated without an excessive computational burden. This is a significant result for the academic literature, as researchers have often chosen to estimate models where firms do not have any information on the realization of demand and marginal cost shocks when entry or positioning decisions are made in order to avoid the computational burden that is perceived to be involved with estimating discrete entry/positioning choice and pricing games simultaneously.

Our second, and more important, contribution comes from performing a set of counterfactuals which try to systematically assess the likelihood and sufficiency of repositioning, consistent with the *Horizontal Merger Guidelines*. We show how to account for the selection on unobserved demand and marginal cost shocks that is implied by the model, and we find that doing so is important. When we take selection into account we predict that rivals are much less likely to launch nonstop service when nonstop duopolists merge than if we ignore selection, and we predict larger average price increases and significant decreases in consumer surplus. We find that our predictions are consistent with what has been observed after actual airline mergers only when we account for selection. These results are important both for academic research, where we are not aware of this type of conditioning being used previously, and for the analysis of mergers at antitrust agencies, where it is common to perform merger simulations and other counterfactuals, even when parameters come from documents, expert testimony or simple calibrations rather being econometrically estimated.

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

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

Baker, J. B. (1996): "The Problem with Baker Hughes and Syufy: On the Role of Entry in Merger Analysis," *Antitrust Law Journal*, 65, 353–374.

Benkard, L., A. Bodoh-Creed, and J. Lazarev (2020): "Simulating the Dynamic Effects of Horizontal Mergers: U.S. Airlines," Discussion paper, Stanford University.

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

Berry, S., M. Carnall, and P. T. Spiller (1996): "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity," Discussion paper, National Bureau of Economic Research.

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

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

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.

Carlton, D., M. Israel, I. MacSwain, and E. Orlov (2017): "Are Legacy Airline Mergers Pro-or Anti-Competitive? Evidence from Recent US Airline Mergers," *International Journal of Industrial Organization*.

Carlton, D. W. (2004): "Why Barriers to Entry are Barriers to Understanding," *American Economic Review*, 94(2), 466–470.

Ciliberto, F., C. Murry, and E. T. Tamer (2020): "Market Structure and Competition in Airline Markets," Discussion paper, University of Virginia.

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

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

Coate, M. B. (2008): "Theory Meets Practice: Barriers to Entry in Merger Analysis," *Review of Law & Economics*, 4(1), 183–212.

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

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

Fan, Y., and C. Yang (2020a): "Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market," *American Economic Journal: Microeconomics*, 12(2), 99–134.

——— (2020b): "Merger, Product Repositionng and Firm Entry: The Retail Craft Beer Market in California," Discussion paper, University of Maryland.

Fisher, F. M. (1987): "Horizontal Mergers: Triage and Treatment," *Journal of Economic Perspectives*, 1(2), 23–40.

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

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

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

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

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

Keane, M. P. (1994): "A computationally practical simulation estimator for panel data," *Econometrica: Journal of the Econometric Society*, pp. 95–116.

Keyes, L. S. (1987): "The Regulation of Airline Mergers by the Department of Transportation," *Journal of Air Law & Commerce*, 53, 737–764.

Kirkwood, J. B., and R. O. Zerbe (2009): "The Path to Profitability: Reinvigorating the Neglected Phase of Merger Analysis," *George Mason Law Review*, 17, 39–112.

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

Laffont, J.-J., H. Ossard, and Q. Vuong (1995): "Econometrics of First-Price Auctions," *Econometrica*, 63(4), 953–980.

Lazarev, J. (2013): "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in US Monopoly Markets," Discussion paper, New York University.

Li, S., J. Mazur, Y. Park, J. Roberts, A. Sweeting, and J. Zhang (2018): "Repositioning and Market Power After Airline Mergers," Discussion paper, University of Maryland.

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

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.

Luttmann, A. (2019): "Evidence of Directional Price Discrimination in the US Airline Industry," *International Journal of Industrial Organization*, 62, 291–329.

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

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

Nevo, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, 31(3), 39–421.

Park, Y. (2020): "Structural Remedies in Network Industries: An Assessment of Slot Divestitures in the American Airlines/US Airways Merger," Discussion paper, University of Massachusetts.

Peters, C. (2006): "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry," *Journal of Law and Economics*, 49(2), 627–649.

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

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

Rybnicek, J., and J. D. Wright (2014): "Outside In or Inside Out? Counting Merger Efficiencies Inside and Out of the Relevant Market," Discussion paper, George Mason University.

Schmalensee, R. (1987): "Horizontal Merger Policy: Problems and Changes," *Journal of Economic Perspectives*, 1(2), 41–54.

Shapiro, C. (2010): "The 2010 Horizontal Merger Guidelines: From Hedgehog to Fox in Forty Years," *Antitrust Law Journal*, 77(1), 49–107.

Sweeting, A. (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 (2020): "A Model of Dynamic Limit Pricing with an Application to the Airline Industry," *Journal of Political Economy*, 128(3), 1148–1193.

Wang, E. Y. (2015): "The Impact of Soda Taxes on Consumer Welfare: Implications of Storability and Taste Heterogeneity," *RAND Journal of Economics*, 46(2), 409–441.

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.

Williams, K. (2020): “Dynamic Airline Pricing and Seat Availability,” Discussion paper, Yale University.

Wollmann, T. (2018): “Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles,” *American Economic Review*, 108(6), 1364–1406.

APPENDICES TO "REPOSTIONING AND MARKET POWER AFTER AIRLINE MERGERS" FOR ONLINE PUBLICATION

A Comparison of the Limited and Full Information Models

The model used in the text assumes that carriers know all demand and cost shocks when making service choices. An alternative assumption used in the literature assumes that firms only know the distributions from which these shocks are drawn. In this Appendix we investigate whether this difference in modeling assumptions matters. Section A.1 provides the model and then our discussion is divided into two parts. Section A.2 compares outcomes under the two different assumptions using the same parameters. We show that limited information assumptions imply that, for markets that are sufficiently large, more carriers will tend to choose nonstop service, and that a merger of nonstop carriers will result in a higher probability of repositioning. Section A.3 examines what happens when parameters of a limited information model are estimated using data generated from a full information model.[60] It shows that there are substantial, and intuitive, biases in the estimated parameters: for example, the quality impact of nonstop service is overestimated because selection is not accounted for, and, to explain why fewer carriers choose nonstop service, mean nonstop fixed costs are overestimated. While the upward bias in the estimated fixed cost tends to reduce predicted post-merger repositioning, we continue to find that we overpredict likely repositioning.

A.1 Model

Overview. We consider a single market, although, in Section A.2 we shall vary its size, with six carriers, A,..,F. In the first stage of the game, the carriers choose whether to provide connecting service or higher-quality nonstop service. Nonstop service requires payment of a fixed cost. Having selected their service types they simultaneously choose prices in the second stage.

Demand. Demand is determined by a nested logit model. The indirect utility for consumer k using carrier j is $u_{kj} = \beta_j - \alpha p_j + \tau\zeta_k + (1-\tau)\varepsilon_{kj}$, with $\tau = 0.7$, $\alpha = 0.5$ (for a price measured in hundreds of dollars), and $\beta_j = \beta_j^{CON} + \beta_j^{NS} \times \mathcal{I}(j \text{ is nonstop})$. The mean utility of not traveling is zero.

Marginal Cost. Carriers will have a linear marginal cost of connecting service and a (possibly different) linear marginal cost of nonstop service.

Fixed Cost. To provide nonstop service, carriers pay a fixed cost which is drawn from a normal distribution with mean \$600,000 and standard deviation \$125,000.

[60] For simplicity we will assume that some of the demand parameters, such as the price coefficient and the nesting parameter are known.

Information Structures. We compare outcomes under two alternative information structures, although both are "complete information" in the sense that the firms do not have any private information. Under "full information", all draws are known to all carriers throughout the game. Under "limited information", carriers only know the model parameters and the draws of fixed costs (assumed to be known by all carriers) in the first stage, but the demand and marginal cost draws are revealed before prices are chosen. Limited information is the common assumption in the empirical literature on models with entry or product selection and price competition (Draganska, Mazzeo, and Seim (2009), Eizenberg (2014), Wollmann (2018) and Fan and Yang (2020a)). The method for solving the full information model is the same as the one used in the paper. For the limited information model, we approximate the expected profits of each carrier in every possible market configuration by taking 1,000 draws of marginal costs and qualities.

A.2 Comparing Outcomes from Limited and Full Information Models With Identical Parameters

A.2.1 Parameterization.

β_i^{CON} is drawn from a normal distribution with standard deviation 0.2 and mean values of 0.6, 0.55, 0.5, 0.45, 0.4 and 0.35 for carriers A to F respectively. The carriers make their service choices in alphabetical order, so that the carrier with the highest expected quality moves first. The incremental quality of nonstop service, β_j^{NS}, is a draw from a truncated normal distribution with mean 0.3, standard deviation 0.2 and a lower truncation point of 0. Carrier marginal costs are \$200 for nonstop service and \$220 for (longer) connecting service, plus a carrier-specific random component, common across service types, drawn from a normal distribution with mean zero and standard deviation \$15.

A.2.2 Comparison.

We simulate equilibrium outcomes 50,000 times for each of 30 different market sizes, ranging from 5,000 and 295,000, for both information structures.

Number of Nonstop Carriers and Consumer Surplus Figure A.1 compares the average number of nonstop carriers and consumer surplus in equilibrium. In a small market, nonstop service may only be profitable when a carrier has unusually high nonstop quality or low marginal costs, unless its fixed cost is very low. Knowledge of quality and marginal cost draws can therefore make it more likely that a carrier will be nonstop. Fewer carriers provide nonstop service in larger markets under full information. The intuition comes from the competitiveness of the nonstop rivals that a carrier expects to face. Under full information, a nonstop rival will tend to be a stronger competitor (because it has been selected based on its quality and cost), which lowers the expected nonstop profitability of another carrier considering nonstop service. This reduces the number of

Figure A.1: The Relationship Between Market Size, Expected Consumer Surplus and the Expected Number of Nonstop Carriers Under Different Informational Assumptions

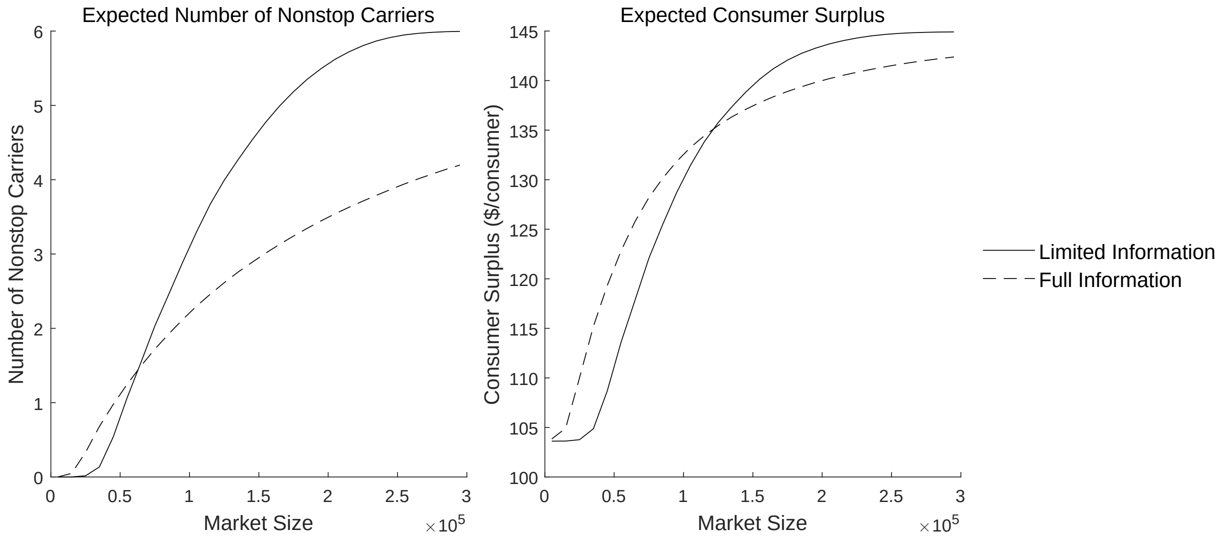

nonstop carriers in equilibrium. However, selection also means that nonstop carriers tend to provide better quality products, which raises expected consumer surplus under full information for a given number of nonstop carriers. Figure A.2 shows that, for a given market size, the *distribution* of the number of nonstop carriers is much tighter under limited information.[61]

Implications for Post-Merger Changes. The tighter distributions for the number of nonstop carriers in a limited information model has implications for what we would predict should happen after an exogenous merger if rival carriers can change their service choices. To illustrate, we consider a market size of 85,000 and collect all sets of draws that result in A and B being nonstop duopolists, which is the most common outcome under either information structure. Now suppose that these carriers merge, eliminating the carrier with the smaller market share, and that the remaining carriers can re-optimize their service choices in the same sequential order.[62] Under limited information, the probability that at least one rival carrier will introduce nonstop service after the merger is 0.8, and the expected reduction in consumer surplus following the merger is just under \$0.3 million. Under full information, the probability that at least one rival will introduce nonstop service after the merger is 31% lower (0.55) and the expected loss of consumer surplus is almost \$1.15 million.[63]

[61] For example, for a market size of 145,000, 97% of simulated outcomes have either three or four nonstop carriers, compared with 69% under full information.

[62] The reader might view it as unreasonable to use the limited information assumption in this case because carriers' pre-merger experience on the route in question would inform them of their quality and costs, even for the type of service that they are not offering. We completely agree, which is one reason why we believe a full information model is the natural model for merger counterfactuals.

[63] The loss in consumer surplus is greater under full information not only because there is less repositioning but also because the pre-merger market shares of the merging nonstop carriers, whose merger we are considering, tend

Figure A.2: The Relationship Between Market Size and Equilibrium Market Structure Under Different Informational Assumptions

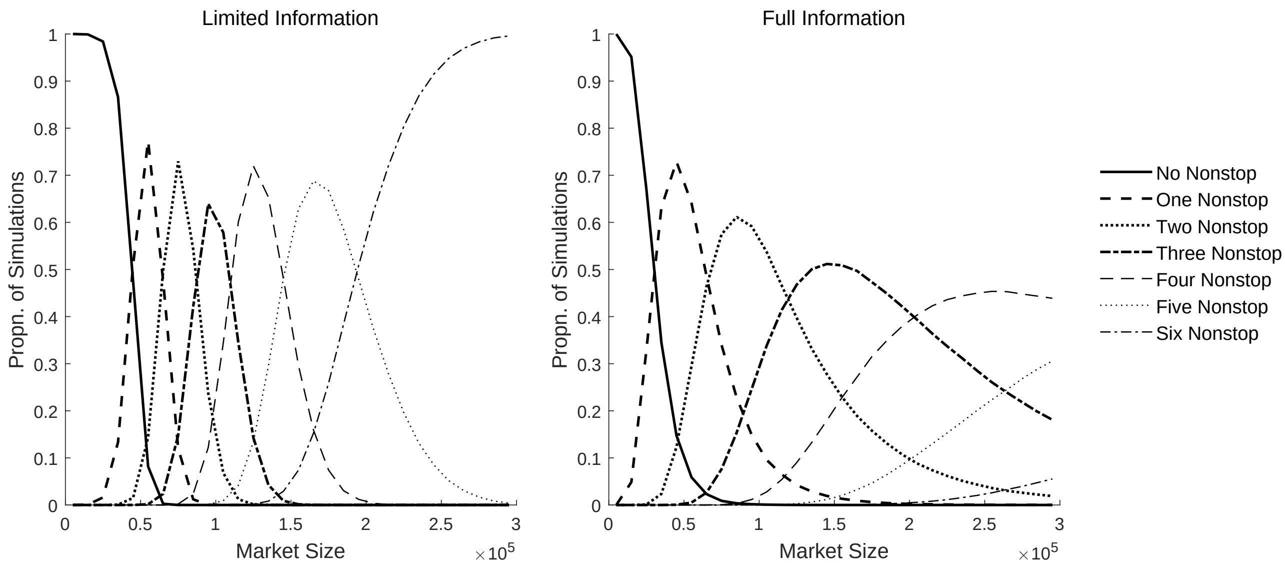

In the limited information case, the merger is also, on average, unprofitable for the merging parties, while it is profitable under full information.

Regret. The example also illustrates the feature that carriers may frequently regret their choices under limited information: for example, for a market size of 55,000, for 48% of the draws where a single carrier chooses to be nonstop, that carrier would have increased its (ex-post) profits by only offering connecting service.

A.3 Estimating a Limited Information Model Using Full Information Data

As explained in the text, the majority of papers in the literature estimate limited information models. One reason for this choice is that the ability to estimate demand separately allows for the inclusion of numerous fixed effects that may help to make "unobservable quality observable". In this Section we therefore use our example model to address the question of whether using fixed effects in a limited information model will lead to accurate estimates and predictions when the true model involves full information. We find that there are large and intuitive biases in the estimated parameters, and that, while the bias in the fixed cost parameters tends to reduce the extent to which a limited information counterfactual overstates how much repositioning will take place in response to a merger, there is still a non-trivial difference in the predictions.

to be higher because of selection.

Parameterization. We use a simpler parameterization than in Section A.2. Specifically we assume that all carriers are symmetric, and that β_i^{CON} is drawn from a normal distribution with mean 0.4 and standard deviation 0.2. β_i^{NS} is a draw from a normal distribution with mean 0.3 and standard deviation 0.2. Given the value of the price coefficient, these coefficients correspond to an average quality premium for nonstop service that is worth \$60, and the standard deviations of both draws of \$40. To simplify estimation of the limited information model, we assume that this draw is not truncated, so that some carriers will have lower nonstop quality than connecting quality. The marginal costs of all carriers for both types of service are set equal to \$200. To provide nonstop service, carriers pay a fixed cost which is drawn from a normal distribution with mean \$600,000 and standard deviation \$125,000.

Parameters to Estimate. We focus our comparison on the estimation of the parameters reflecting average connecting quality (which here corresponds to a market fixed effect, as it is the same for all carriers), the average incremental nonstop quality and the standard deviation of the quality draws for nonstop and connecting service. We treat the price and nesting parameters as known to the researcher, although this does not imply that these parameters would not also be biased.[64] We also report on the bias in the estimated mean and standard deviation of fixed costs.

Data Generation and Estimation of the Demand Parameters. We generate data from 100 simulated markets with market size 90,000. For each market we generate data from 50 independent repetitions of the game, which we will think of as constituting a possible panel structure for the data. For each "simulated market" and $t' = 1,..,T$ we estimate the parameters β_m and β_{NS} in following equation, which corresponds to the equation that would be estimated in a limited information model with a market fixed effect and nonstop dummy.

$$\ln(s_{jmt}) - \ln(s_{0mt}) - \tau \ln\left(\frac{s_{jmt}}{\sum_k s_{kmt}}\right) + \alpha p_{jmt} = \beta_m + \beta_{NS}\mathcal{I}(\text{j is nonstop}) + \xi_{jmt} \ \forall j, t = 1,..t'$$

As α and τ are known, this model can be estimated by OLS *under the limited information assumption.* We then use the estimated residuals to estimate the standard deviations of the connecting and nonstop quality draws.[65]

[64]For example, CMT identify that the price coefficient is biased when selection is not accounted for. However, the magnitude of the bias in their simulations is not as dramatic as some of the biases that we identify here.

[65]We estimate the standard deviation of the connecting draws as the standard deviation of the residuals for the carriers choosing connecting service, and the standard deviation of the nonstop draws as $\sqrt{\sigma^2_{\xi,\mathcal{I}(j=NS)} - \sigma^2_{\xi,\mathcal{I}(j=CONN)}}$ where $\sigma^2_{\xi,\mathcal{I}(j=NS)}$ is the estimated variance of the ϵ_{jmt}s for the carriers that are nonstop.

Figure A.3: Estimated Demand Parameters When A Limited Information Model is Applied to Full Information Data

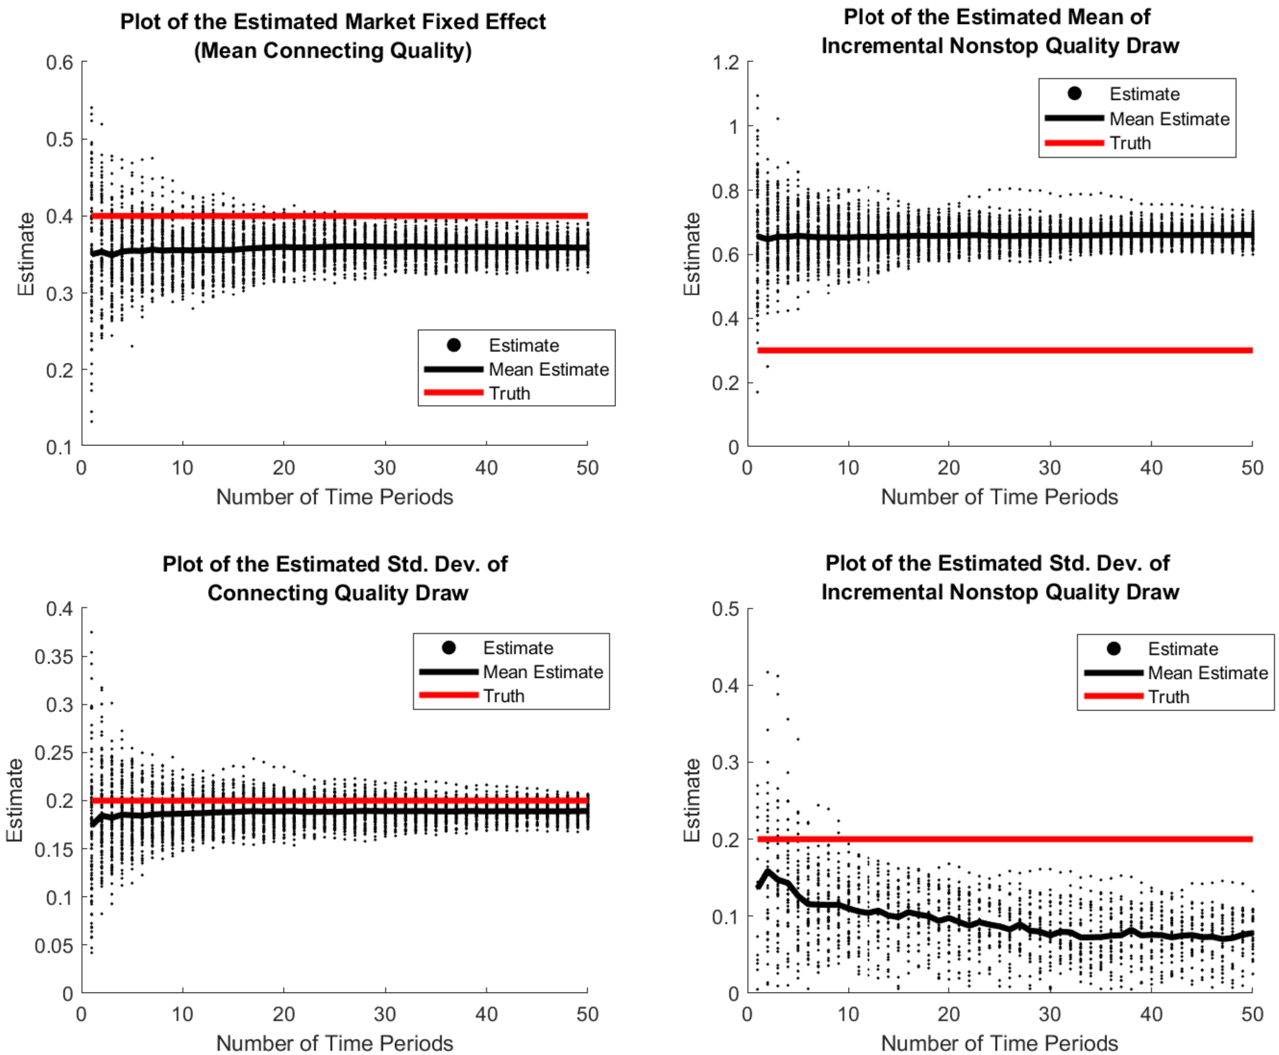

Demand Estimates. Figure A.3 shows the estimates of the demand parameters. On the x-axis is the number of time periods (t') used in estimation, and on the y-axis the estimated value of the parameters. Each dot represents an estimate, and black solid line represents the average estimate, across markets, for panels of length (t').[66] The red lines show the true value of the parameters.

The figures show that the estimates in all of the parameters are biased to some extent, and in a way that reflects the selection that exists in a full information model. Carriers with low connecting quality are also likely to have low nonstop quality, and so will be more likely to choose to offer connecting service. Therefore, the estimated mean connecting quality, which corresponds to the market fixed effect, is biased downwards. The standard deviation of connecting quality is also biased downwards. The biases in the nonstop parameters are larger, partly reflecting the fact that, given the assumed parameters, most carriers choose connecting service (on average, 1.59 carriers are nonstop). The average nonstop quality is biased upwards (its average estimate with a 50-period panel is 0.66, compared with a true value 0.3) and the standard deviation of nonstop quality is biased downwards, reflecting the fact that only the carriers with the highest nonstop quality tend to find nonstop service profitable.

A feature of the estimates is that, as one would expect, the variances of estimates are much larger when we estimate using short panels. This is relevant because some of the papers that assume limited information use short panels and their counterfactuals rely on the point estimates of the values of the fixed effects.[67]

Fixed Cost Estimates. We now consider what these parameters imply for estimated fixed costs. We do so by taking the mean value of the estimated demand parameters across simulated markets based on the longest panels, and then estimate values of the mean and standard deviation of fixed costs using a simulated method of moments estimator where we match the proportion of markets where 1, 2, 3, 4, 5 and 6 carriers are nonstop in the equilibrium of a limited information model and the data generated from the full information model (pooling the data from all of the simulated markets).[68]

This procedure estimates the mean of the fixed cost distribution as $1,553,400 and a standard deviation of $734,910. These compare to true values of $600,000 and $125,000, which is consistent with the fact that, with the same fixed cost parameters, the limited information model would predict more carriers would choose to be nonstop, and a lower variance in the number of nonstop carriers, than a full information model.[69] However, with the estimated parameters the model is

[66] Obviously β_{NS} cannot be estimated if no carriers choose nonstop service or $\sigma^2_{\varepsilon,\mathcal{I}(j=NS)} < \sigma^2_{\varepsilon,\mathcal{I}(j=CONN)}$. For these observations there is no dot and the averages are calculated for the remaining markets.

[67] For example, Aguirregabiria and Ho (2012) use a 4-quarter panel of airline markets. Even if the limited information assumptions are correct and, on average, the estimated fixed effects are unbiased, the average results of counterfactuals could be misleading if predictions are nonlinear functions of the fixed effect coefficients, which will typically be the case.

[68] We solve for expected profits in the limited information model as described in Section A.2, and then simulate 5,000 repetitions of the limited information service choice game.

[69] As discussed in the last sub-section, a limited information model predicts more nonstop carriers when more than

able to match the data on the number of nonstop carriers quite closely: for example, in the data 41.2%, 47.0% and 7.6% of markets have one, two and three nonstop carriers respectively, and the estimated parameters predict 42.1%, 47.9% and 7.9%.

Comparison of the Counterfactuals. We now compare the predictions of merger counterfactuals that use (i) the full information model with the true parameters and (ii) the limited information model with the estimated demand and fixed cost parameters. As the estimated parameters allow us to match the pattern of service choices in the data, we might expect that predictions of repositioning might change after the merger could be quite similar.

We perform the counterfactuals in both cases by running 100,000 simulations for the true model and for the limited information model with the estimated parameters. We identify cases where firms A and B (the first two movers in the sequential service choice game) are nonstop duopolists. We then use the draws that support this outcome to simulate a merger between A and B (which eliminates B). To be consistent with the text, we assume that, after the merger, the merged firm will be nonstop (by setting its fixed costs to zero), and we compare the prices charged by firm A.

A's expected pre-merger prices are similar across the two models: $286.95 for full information and $288.06 for limited information. After the merger, the expected price of A is $305.22 under full information (a 6.4% increase) and $297.49 (a 3.3% increase) under limited information. Therefore there is a non-trivial difference in predicted price changes. This difference results, in part, from a difference in predicted repositioning. Under full information, there is no repositioning by rivals for 54% of draws (in these cases the price increases to $312.58, on average), whereas under limited information at least one rival initiates nonstop service for 57% of draws (in which case post-merger prices are predicted to hardly increase at all, because the new nonstop carrier is just as likely to be a strong nonstop competitor as A or B).

Conclusion. This example illustrates that assuming limited information can result in substantial bias in the parameters when the true model has full information. There is also bias in the counterfactual predictions using either the true or estimated parameters. The example therefore illustrates the importance of incorporating full information into the model, if that is likely to be the correct informational assumption, even if doing so increases computational costs.

B Data Appendix

This Appendix complements the description of the data in Section 3 of the text.

one carrier is nonstop.

B.1 Sample Construction and Variable Definitions

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

- airport-pairs that are less than 350 miles apart as ground transportation may be very competitive on these routes;
- airport-pairs involving Dallas Love Field, which was subject to Wright Amendment restrictions that severely limited nonstop flights;
- airport-pairs involving New York LaGuardia or Reagan National that would violate the so-called perimeter restrictions that were in effect from these airports[70];
- airport-pairs where more than one carrier that is included in our composite "Other Legacy" or "Other LCC" (low-cost) carriers are nonstop, have more than 20% of non-directional traffic or have more than 25% presence (defined in the text) at either of the endpoint airports. Our rationale is that our assumption that the composite carrier will act as a single player may be especially problematic in these situations[71]; and,
- airport-pairs where, based on our market size definition (explained below), the combined market shares of the carriers are more than 85% or less than 4%.

Seasonality. The second quarter is the busiest quarter for airline travel, and one might be concerned that seasonality affects our measures of passenger flows and service choices, and therefore our estimates. We do not believe that this is a first-order concern for our sample of relatively large markets. The website `http://www.anna.aero` (accessed May 29, 2018) provides a formula for measuring the seasonality of airport demand (SVID) which we have calculated for all of the airports in our sample using monthly T100 data on originating passengers.[72] The website classifies seasonality as "excellent" if SVID is less than 2 or "good" if the SVID is between 2 and 10, on the basis that seasonality is costly for an airline or an airport because it requires changes in schedules. All of the airports in our sample are within these ranges, with the highest (most seasonal) values for Seattle (2.4), New Orleans (2.8), Palm Beach (5.2) and Southwest Florida (9.9). In contrast, a non-sample airport with very seasonal demand, Gunnason-Crested Butte (GUC), has an SVID of 65. Applying SVID on a route-level to quarterly traffic, only one sample route (Minneapolis to Southwest Florida) has an SVID greater than 10 (19), and the 95th percentile is 3.12.

[70] 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.

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

[72] The measure is calculated as $\frac{\sum_{m=1,..,M=12}\left(\frac{100\times \text{Traffic}_{a,m}}{\text{Traffic}_a}-100\right)^2}{1000}$.

We also find little evidence of seasonality if we identify routes which a carrier serves nonstop in our data and in the second quarter of 2005, but which they did not serve nonstop in either Q1 2005 or Q1 2006 (i.e., routes where a carrier's nonstop service may be seasonal). We can only identify two such carrier-routes in our sample (United for San Antonio-San Francisco and Sun Country (part of Other Low Cost) for Indianapolis-Kansas City), out of 8,065 carrier-routes.

Definition of players, nonstop and connecting service. We are focused on the decision of carriers to provide nonstop service on a route. Before defining any players or outcomes, we drop all passenger itineraries from DB1 that involve prices of less than $25 or more than $2000 dollars[73], 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). Based on the number of passengers carried, the largest Other Legacy carrier is Alaska Airlines, and the largest Other LCC carriers are JetBlue and AirTran.

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

There is some arbitrariness in these thresholds. However, the 64 flight and 50% nonstop thresholds for nonstop service have little effect because almost all nonstop carriers far exceed these thresholds. For example, Figure B.1 shows that the carriers we define as nonstop typically carry only a small proportion of connecting passengers. For the same reason, we also model nonstop carriers as only providing nonstop service even if some of their passengers fly connecting, although we include the connecting passengers when calculating market shares.

On the other hand, our 1% share/200 passenger thresholds do affect the number of connecting carriers. For example, if we instead require players to carry 300 return passengers and have a 2% share, the average number of connecting carriers per route falls by almost one-third as marginal connecting carriers are excluded.

Market Size. Market size is used to define market shares and to calculate counterfactual quantities and profits. Given the role of market size in the identification and estimation of demand and

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

Figure B.1: Proportion of DB1 Passengers Traveling with Connections, Based on the Type of Service

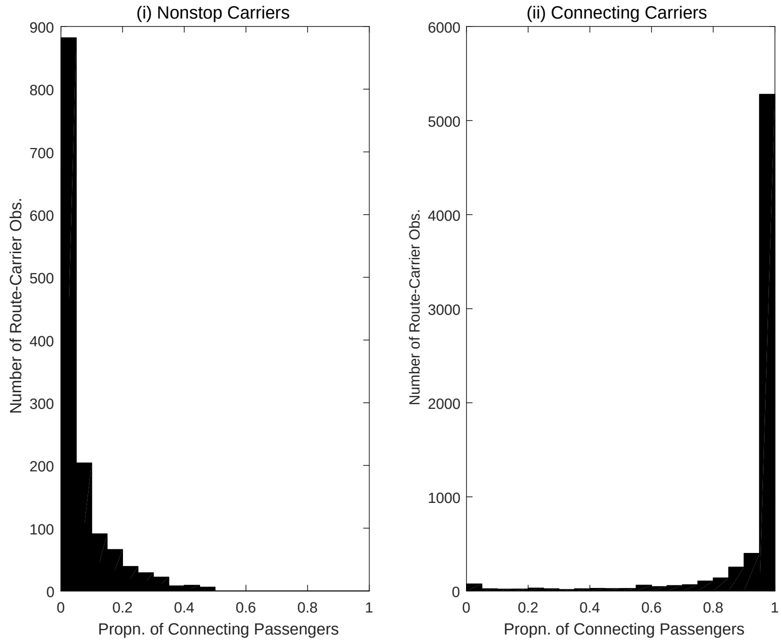

entry-type models, the ideal definition should imply that variation in shares across routes, or across directions, should reflect changes in prices, carrier characteristics and service types, and it should be a good predictor of the number of nonstop firms.

A standard approach in the literature is to use the geometric average of the endpoint MSA populations (e.g., Berry and Jia (2010), Ciliberto and Williams (2014)). However, this performs poorly for airport-pair routes (MSA demand may be split between several airports, it cannot allow for the possibility that demand is systematically different at the endpoints and it does not account for the effects of distance on how much people want to travel.

We therefore consider an alternative definition based on the estimates of a generalized gravity equation, used previously in Sweeting, Roberts, and Gedge (2020). The model specifies that the total number of second quarter passengers on a route varies with a linear function of the log of the count of originating and arriving passengers at each of the endpoint airports (measured for the second quarter of the previous year), log route distance and interactions of these lagged passengers flow and distance variables. The corresponding Poisson regression is estimated using data from 2005-2011, including year, origin and destination fixed effects and interactions between a dummy for long-distance routes, defined as those over 2,300 miles and origin and destination fixed effects.[74] With the estimates in hand, we calculate the expected number of passengers for each directional

[74] The individual coefficients are not especially informative because of the interactions, but combining them shows reasonable patterns. For example, the expected number of passengers declines in route distance, increases with both lagged originating traffic at the origin airport, and lagged arriving traffic at the destination.

market for Q2 2006, based on lagged values of passenger flows in Q2 2005. Our market size measure multiplies this prediction by 3.5.

Two comparisons suggest that our measure provides a superior measure of market size to estimates based on average population. Given that prices and service in each direction on a route tend to be similar we would expect the correlation in the combined market share of all of the carriers to be quite high: using our measure the correlation is 0.86, whereas it is only 0.56 using the geometric average population. Consistent with this difference, if one estimates our model using population-based market size measures, there is much greater unobserved heterogeneity in demand than there is in our estimates. CMT, who use a population-based measure, also estimate much more demand heterogeneity than we do.

Table B.1 examines the ability of the different market size variables to predict the number of nonstop carriers on a route using an ordered probit model. Examination of the reported pseudo-R^2s shows that our gravity measure has much stronger predictive power, and that when we add population-based variables to a specification with a flexible function of our measure (i.e., going from column (2) to column (5)) the R^2 increases by less than 1%. However, because we recognize that our market size measure is still imperfect, we also allow for an additional route-level unobservable that is common to the demand of all carriers, but is unobserved by the researcher.

Prices and Market Shares. As is well-known, airlines use revenue management strategies that result in passengers on the same route paying quite different prices. Even if more detailed data (e.g., on when tickets are purchased) was available, it would likely not be feasible to model these type of strategies within the context of a combined service choice and pricing game. We therefore use the average price as our price measure, but allow for prices and market shares (defined as the number of originating passengers carried divided by market size) to be different in each direction, so that we can capture differences in passenger preferences (possibly reflecting frequent-flyer program membership) across different airports.[75]

B.2 Network Variables.

The legacy carriers in our data operate hub-and-spoke networks. On many medium-sized routes local demand could not generate sufficient variable profits to cover the fixed costs of nonstop service, but nonstop service may be profitable once the value of passengers who will use a nonstop flight as one segment on a longer journey is taken into account. While our structural model captures price competition for passengers traveling only the route itself, we allow for three "network variables" that capture the value of traffic to other destinations to offset the fixed cost of providing nonstop service.

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

Table B.1: Market Size Measures and the Number of Nonstop Carriers

	(1)	(2)	(3)	(4)	(5)
Our Market Size	3.230	11.05			11.04
(/10,000)	(0.110)	(0.440)			(0.482)
Our Market Size^2		-8.933			-8.780
		(0.560)			(0.587)
Our Market Size^3		2.283			2.230
		(0.190)			(0.196)
Geom. Avg. Pop.			2.476	10.48	2.125
(/1 m.)			(0.136)	(0.823)	(0.966)
Geom. Avg. Pop.^{2}				-12.98	-4.835
				(1.536)	(1.757)
Geom. Avg. Pop.^{3}				4.977	2.433
				(0.773)	(0.877)
Ordered Probit Cutoffs					
Cutoff 1	0.730	1.596	0.725	1.801	1.813
	(0.0369)	(0.0604)	(0.0460)	(0.113)	(0.126)
Cutoff 2	2.082	3.350	1.722	2.844	3.571
	(0.0563)	(0.0965)	(0.0548)	(0.120)	(0.146)
Cutoff 3	3.915	4.995	2.761	3.890	5.217
	(0.128)	(0.132)	(0.0789)	(0.133)	(0.171)
Cutoff 4	6.987	6.877	4.134	5.181	7.112
	(0.431)	(0.333)	(0.232)	(0.240)	(0.351)
Observations	2,028	2,028	2,028	2,028	2,028
Pseudo-R^2	0.262	0.368	0.0770	0.109	0.371

Notes: coefficients from an ordered probit regression where the dependent variable is the number of nonstop carriers on the non-directional route. "Our market size" measure is the average of our measure of market size across directions. Standard errors in parentheses. Number of observations is equal to the number of routes.

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 small, even though the carrier serves many destinations from the airport. International hubs are airports that carriers use to serve a significant number of non-Canadian/Mexican international destinations nonstop. Table B.2 shows the airports counted as hubs for each named carrier.

An Ancillary Model of Connecting Traffic The third variable is a continuous measure of how much connecting traffic a carrier is likely to carry if it serves a route to a domestic hub nonstop. We use a reduced-form model of network flows that fits the data well[76] and which gives us a prediction of how much connecting traffic that a carrier can generate on a route where it does not currently provide nonstop service, taking the service that it provides on other routes as given. We include this prediction in our model of service choice as a variable that can reduce the effective fixed cost of providing nonstop service on the route.[77]

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

$$s_{c,i,od} = \frac{\exp(X_{c,i,od}\beta + \xi_{c,i,od})}{1 + \sum_l \sum_k \exp(X_{l,k,o,d}\beta + \xi_{l,k,od})} \tag{2}$$

where $X_{c,i,od}$ is a vector of observed characteristics for the connection (c)-carrier (i)-origin (o)-destination (d) combination and $\xi_{c,i,od}$ is an unobserved characteristic. The Xs are functions of variables that we are treating as exogenous such as airport presence, endpoint populations and geography. The outside good is traveling using connecting service via an airport that is not one of the domestic hubs that we identify.[78] Assuming that we have enough connecting passengers that the choice probabilities can be treated as equal to the observed market shares, we could potentially

[76]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 wanting to avoid excessive complexity, we would face the problem that we would not observe connection times for routes that do not currently have nonstop service on each segment, but which could for alternative service choices considered in our model.

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

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

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

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

estimate the parameters using the standard estimating equation for aggregate data (Berry 1994):

$$\log(s_{c,j,od}) - \log(s_{0,od}) = X_{c,j,od}\beta + \xi_{c,j,od}. \tag{3}$$

However, estimating (3) would ignore the selection problem that arises from the fact that some connections may only be available because the carrier will attract a large share of connecting traffic. We therefore introduce an additional probit model, as part of a Heckman selection model, to describe the probability that carrier i does serve the full *ocd* route,

$$\Pr(i \text{ serves route } ocd) = \Phi\left(W_{c,j,od}\gamma\right). \tag{4}$$

Sample, Included Variables and Exclusion Restrictions. We estimate our model using data from Q2 2005 (one year prior to the data used to estimate our main model) for the top 100 US airports. We use DB1 passengers who (i) travel from their origin to their destination making at least one stop in at least one direction (or their only direction if they go one-way) and no more than one stop in either direction; and, (ii) have only one ticketing carrier for their entire trip. For each direction of the trip, a passenger counts as one-half of a passenger on an origin-connecting-destination pair route (so a passenger traveling RDU-ATL-SFO-CVG-RDU counts as $\frac{1}{2}$ on RDU-ATL-SFO and $\frac{1}{2}$ on RDU-CVG-SFO). Having joined the passenger data to the set of carrier-origin-destination-connecting airport combinations, we then exclude origin-destination routes with less than 25 connecting passengers (adding up across all connecting routes) or any origin-connection or connection-destination segment that is less than 100 miles long.[79] We also drop carrier-origin-destination-connecting airport observations where the carrier (or one of its regional affiliates) is not, based on T100, providing nonstop service on the segments involved in the connection. This gives us a sample of 5,765 origin-destination pairs and 142,506 carrier-origin-destination-hub connecting airport combinations, of which 47,996 are considered to be served in the data.

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

[79] 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.

and SFO, but it should not be directly relevant for the choice of whether a traveler who is going from RDU to SFO connects via Atlanta (or a smaller city such as Charlotte), so these population terms can appear in the selection equation for whether nonstop service is offered but not the connecting share equation. In $W_{c,j,od}$ we include origin, destination and connecting airport presence for carrier i; the interactions of origin and connecting airport presence and of destination and connecting airport presence; origin, destination and connecting city populations; the interactions of origin and connecting city populations and of destination and connecting city populations, a count of the number of airports in the origin, destination and connecting cities[80]; indicators for whether either of the origin or destination airports is an airport with limitations on how far planes can fly (LaGuardia and Reagan National) and the interactions of these variables with the distance between the origin or destination (as appropriate) and the connecting airport; indicators for whether the origin or destination airport are slot-constrained. In both $X_{c,j,od}$ and $W_{c,j,od}$ we also include origin, destination and carrier-connecting airport dummies.

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

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

However we find that the estimated model provides quite accurate predictions of how many connecting travelers use different segments, which makes us believe that it should be useful when thinking about the gain to adding some marginal nonstop routes to a network. For the named legacy carriers in our primary model, there is a correlation of 0.96 between the predicted and observed numbers of connecting passengers on segments that are served nonstop. The model also captures some natural geographic variation. For example, for many destinations a connection via Dallas is likely to be more attractive for a passenger originating in Raleigh-Durham (RDU) than a passenger originating in Boston (BOS), while the opposite may hold for Chicago. Our model predicts that American, with hubs in both Dallas (DFW) and Chicago (ORD), should serve 2,247

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

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

	Connecting Share	Serve Route	$\frac{1}{2}log\frac{1+\rho}{1-\rho}$	$log(std.\ deviation)$
Constant	4.200	-8.712	-0.109	0.308
	(0.338)	(0.823)	(0.0860)	(0.0150)
Presence at Origin Airport	4.135	6.052		
	(0.396)	(1.136)		
Presence at Connecting Airport		11.90		
		(0.721)		
Presence at Destination Airport	2.587	6.094		
	(0.396)	(1.126)		
Origin Presence X Connecting Presence		-5.536		
		(1.311)		
Destin. Presence X Connecting Presence		-5.771		
		(1.303)		
Population of Connecting Airport		-1.20e-07		
		(3.16e-08)		
Origin Population X Origin Presence		-5.09e-08		
		(2.23e-08)		
Destin. Population X 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 X Distance Origin - Connection		0.000569		
		(0.000207)		
Destin. Restricted X 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 X Relative Distance^2	0.479			
	(0.151)			
Most Convenient Hub of Any Carrier X	0.590			
Relative Distance	(0.434)			
Origin Presence X Destination Presence	-5.278			
	(0.513)			
Observations	142,506	-	-	-

Notes: robust standard errors in parentheses.

connecting DB1 passengers on RDU-DFW, 1,213 on RDU-ORD and 376 on RDU-STL (St Louis), which compares with observed numbers of 2,533, 1,197 and 376. On the other hand, from Boston the model predicts that American will serve more connecting traffic via ORD (2,265, observed 2,765) than DFW (2,040, observed 2,364).

B.3 Nonstop Duopoly Routes Used in the Counterfactual Analysis.

Most of our counterfactual analysis will involve 24 routes where nonstop duopolists were involved in four specific legacy carrier mergers. While a focus on routes with multiple nonstop carriers is sensible given that these are the types of mergers that, with no service changes, are predicted to have the largest price increases (see Table 7) and most passengers travel on routes with at least two nonstop carriers (Table 2), one might wonder whether the 24 routes that we focus on are representative of nonstop duopoly routes in general. Table B.4 provides a comparison between three groups of markets.

The first group contains the 24 legacy nonstop duopoly routes that we use in our main counterfactuals. The second group contains all remaining legacy nonstop duopoly routes (for example, American and Northwest might be the nonstop duopolists). The third group are nonstop duopoly routes where Southwest (the named non-legacy carrier in our analysis) is nonstop.

The most noticeable pattern when we compare the first two groups is that on the 24 routes there are more connecting rivals and, together, they account for a larger market share. Therefore one would expect that, holding everything else equal, mergers on our routes would tend to have less anticompetitive effects and that, simply given the larger number of connecting passengers, there might be a higher probability of repositioning. Nonstop prices are significantly lower on routes where Southwest offers nonstop service, consistent with Southwest having lower costs. As noted in Appendix B.4, we also observe different price changes on Southwest-Airtran nonstop duopoly routes after the Southwest-Airtran merger, consistent with greater efficiencies. These differences imply that we should not extrapolate from our results, which fit what happens after legacy mergers well, to what would happen after mergers involving low-cost carriers.

B.4 An Analysis of Changes to Prices and Service After Airline Mergers Post-2006

We use our model to predict the effects of three legacy carrier mergers that took place after the period of our data (Delta/Northwest merger (closed October 2008), United/Continental (October 2010) and American/US Airways (December 2013)). In this section we describe an analysis of what happened to the prices and quantities of the merging parties and the service decisions of rivals on routes where the merging parties were nonstop duopolists. Holding service types fixed, one would expect that the merger might create significant market power on these routes. We also consider the Southwest/Airtran merger (May 2011) although we do not perform counterfactuals for that

Table B.4: Characteristics of the Nonstop Duopoly Routes

	Legacy Nonstop Duopoly Routes in Our CF Analysis		Legacy Nonstop Duopoly Routes Not in Our CF Analysis		Nonstop Duopoly Routes with Southwest NS	
	Mean	SD	Mean	SD	Mean	SD
Route distance (miles)	2,437	1,304	1,885	940	1,932	1,111
Number of Conn. Carriers	2.833	1.435	1.778	1.321	2.314	1.954
Combined NS Mkt Shares A to B	25.20%	10.50%	27.00%	9.90%	32.80%	12.50%
Combined NS Mkt Shares B to A	24.40%	8.70%	36.30%	9.20%	31.80%	12.90%
Average NS Price A to B	\$488	\$107	\$450	\$110	\$322	\$80
Average NS Price B to A	\$484	\$102	\$449	\$113	\$319	\$79
Combined Conn. Mkt Shares A to B	3.10%	2.20%	1.80%	2.10%	2.50%	2.40%
Combined Conn. Mkt shares B to A	3.10%	2.30%	1.70%	2.20%	2.60%	2.60%
Number of Routes	24		99		86	

N ote: NS=nonstop, Conn.= connecting, CF = Counterfactual.

merger as Airtran is part of our composite Other LCC carrier. To perform the analysis, we created a panel dataset that runs from the first quarter of 2001 to the first quarter of 2017 using the same definition of nonstop service, but without aggregating smaller carriers into composite Other Legacy and Other LCC rivals.

B.4.1 Frequency of Rivals Launching Nonstop Service

On routes where the merging firms are nonstop duopolists before the merger, the merged firm always maintains nonstop service until the end of our data. We calculate the number of routes where at least one rival carrier, including carriers that were not providing any service prior to the merger, initiated nonstop service within two (or three) years of the merger closing. A two year window is often considered when examining entry and repositioning in merger cases, and was the window considered by the Department of Transportation when it reviewed airline mergers (Keyes (1987)). We will use three years in our analysis of price and quantity changes below as an additional year provides more precision to our estimates which are based on a small number of routes, with only small effects on the point estimates.

We find that no rivals (no rivals) initiated nonstop service within two (three) years on five routes where the merging parties were nonstop duopolists immediately before the closing of the merger for Delta/Northwest. Rivals did initiate nonstop service on one (two) out of five routes for United/Continental, three (four) out of six routes for American/US Airways and one (one) out of seventeen nonstop duopoly routes for Southwest/Airtran. Therefore, the overall rate of rivals initiating nonstop service was five (seven) out of thirty-three routes, or four (six) out of sixteen if we only consider legacy mergers.[81]

One explanation for a low rate of repositioning is that rivals are ill-suited to provide nonstop service on these routes, so that the merging carriers can exercise market power even if the merger does not generate efficiency advantages (higher quality or lower marginal costs). This will be the explanation that we focus on in our counterfactuals. However, an alternative explanation is that it is efficiencies created through the merger that make it unattractive for rivals to offer nonstop service. An analysis of changes to price and market shares can give some insights into which of these stores are correct.

B.4.2 Changes to the Merging Carriers' Prices and Quantities

We define a treatment group of routes where the merging carriers were nonstop duopolists prior to the merger. We also define a control group of routes where one of the merging carriers is nonstop and the other is either not on the route at all or is at most a quite marginal connecting carrier, with a nondirectional share of traffic of less than 2%. However, we acknowledge that the literature has defined control groups in a number of different ways, with different results (see the literature review

[81] There is no overlap in the routes across these mergers.

in the Introduction), and that to the extent that carriers offer networks, it is implausible that the control routes would be completely unaffected by changes in the treatment routes. We also restrict the control group to only include routes where no carriers initiated new nonstop service after the merger. We define three year pre- and post-merger windows (this provides more power than two year windows, although the pattern of the coefficients are similar using two or three year windows). For Delta/Northwest the windows are Q3 2005-Q2 2008 and Q1 2009-Q4 2011. For United/Continental the windows are Q3 2007-Q2 2010 and Q1 2011-Q4 2013. For American/US Airways the situation is less straightforward as detailed negotiations between the parties, a bankruptcy judge and the Department of Justice were known to be ongoing from at least August 2012. We therefore use windows of Q3 2009-Q2 2012 and Q2 2014-Q1 2017.[82] For Southwest/Airtran we use windows of Q2 2007-Q1 2010 and Q3 2010-Q2 2013.

We use a regression specification

$$y_{imt} = \beta_0 + \beta_1 * \text{Treatment}_{im} * \text{Post-Merger}_{it} + X_{imt}\beta_2 + Q_t\beta_3 + M_{im}\beta_4 + \varepsilon_{imt}$$

where y_{imt} is the outcome variable (the log of the weighted average price or the log of the combined number of local passengers (i.e., passengers just flying the route itself and not making connections to other destinations) on the merging carriers) for merging carrier i in directional airport-pair market m in quarter t, Q_t and M_{im} are quarter and carrier-market dummies and β_1 is the coefficient of interest.[83] m is defined directionally, but we cluster standard errors on the non-directional route. X_{imt} contains dummy controls for the number of competitors (including connecting carriers), distinguishing between legacy and LCC competitors, and one-quarter lagged fuel prices interacted with route nonstop distance and its square. A route is defined to be in the treatment or the control group based on the observed market structure in the last four quarters of the pre-merger window (so to be in the treatment group, for example, both merging carriers must be nonstop in each of these quarters). Note that this means that the treatment samples are different and smaller than those considered for the repositioning analysis above, where we defined duopoly based on the one quarter immediately before the financial closing of the merger. They can also differ from the routes used in our counterfactuals where we will use the market structure from Q2 2006.

The results are presented in Table B.5. We report results for each merger and for the three legacy mergers combined. The upper part of the table presents the results when we only include treatment routes where no rivals launch nonstop service before or during the post-merger window. In the lower panel we only use treatment routes where at least one rival initiated nonstop service after the financial closure of the merger but before or during the post-period window, and, for these routes, we only include post-merger window observations where this rival service was actually

[82] We exclude two American/US Airways routes where rivals began service between the end of the pre-merger window and the financial closing of the merger from the treatment group.

[83] To be clear, in the pre-merger period we combine the number of passengers on the merging carriers and use their weighted average fare, so there is a single observation per market-quarter.

Table B.5: Price and Quantity Changes After Four Mergers

	(1) All Legacy Mergers	(2) Delta/ Northwest	(3) United/ Continental	(4) American/ US Airways	(5) Southwest/ Airtran
	Routes Where No Rivals Initiate Nonstop Service Post-Merger				
Dep. Var.: Log (Average Fare)					
Treatment X Post-Merger	0.111	0.141	0.084	0.108	0.038
	(0.052)	(0.078)	(0.026)	(0.118)	(0.028)
Dep. Var.: Log (Number of Local Passengers)					
Treatment X Post-Merger	-0.295	-0.230	-0.463	-0.323	-0.073
	(0.078)	(0.134)	(0.169)	(0.117)	(0.091)
Number of Non-Directional Routes:					
Treatment Group	9	3	4	2	4
Control Group	298	107	112	79	185
	Routes Where At Least One Rival Initiated Nonstop Service Post-Merger				
Dep. Var.: Log (Average Fare)					
Treatment X Post-Merger	0.032	-	-0.028	-0.047	-0.229
	(0.045)		(0.105)	(0.028)	(0.027)
Dep. Var.: Log (Number of Local Passengers)					
Treatment X Post-Merger	-0.358	-	-0.696	-0.478	0.376
	(0.077)		(0.378)	(0.074)	(0.110)
Number of Non-Directional Routes:					
Treatment Group	4	-	1	3	1
Control Group	298	107	112	79	185

Notes: an observation is a carrier-directional airport pair, and only observations for the merging carrier(s) are included. Dependent variable is the log of the weighted average of fares or the log of the combined number of local passengers (i.e., not including passengers connecting to other destinations) on the merging carriers. The pre- and post-merger windows are defined in the text. For treatment routes where a rival initiated nonstop service we only use post-merger observations after the rival began nonstop service. Standard errors in parentheses are clustered on the non-directional route.

provided.

The results are suggestive, despite the small number of treatment observations. For the legacy mergers the pattern is that prices increase and the number of local passengers falls in the treatment routes when no rivals initiate nonstop service, consistent with an increase in market power and limited synergies from combining service on the treatment routes. The fall in the number of local passengers is large, but this pattern appears to be robust: for example, if we also include a linear time trend for the treatment group markets, to allow for the possibility that demand was falling in the type of markets that are nonstop duopolies, the coefficient is -0.293 with a standard error of 0.092. This is almost identical to the coefficient of -0.295 reported in Table B.5, column (1). On the other hand, on routes where rival nonstop service is initiated there is no clear pattern of price increases. The number of passengers carried by the merging carriers declines in these markets, presumably due to competition from the new nonstop carrier.

The pattern is different for Southwest/Airtran, although we note that we have fewer treatment routes than the sixteen routes that were nonstop duopolies immediately before the merger because, in a number of routes, a legacy carrier stopped its nonstop service during the pre-merger window once both Southwest and Airtran were nonstop. There is no statistically significant price increase on the nonstop duopoly routes when Southwest and Airtran merge and there is no statistically significant decline in the number of passengers. This result suggests that this LCC merger may have generated route-level synergies.

C Estimation and Robustness Checks

This Appendix provides additional detail on how we solve the model, the performance of our estimation algorithm and the robustness of our estimates. Section C.1 explains how we solve the model. Section C.2 explains the choice of the importance densities used in estimation. Sections C.3-C.6 analyze aspects of the performance of the estimation algorithm in more detail, including the fit of the model and the robustness of the results to reducing the number of moments. Sections C.7 presents estimation results using moment inequalities. The reader is referred to Li, Mazur, Park, Roberts, Sweeting, and Zhang (2018) for details of a Monte Carlo procedure that illustrates the good performance of our estimation procedures, under our baseline assumption and using inequalities.

C.1 Solving the Model

Our baseline assumption is that service choices are made sequentially in a known order. For a given set of service choices on a given route, we can solve for a unique Bertrand Nash pricing equilibrium in each direction by solving the system of first-order conditions. One approach for solving the service choice game would be to compute equilibrium variable profits for every possible service choice combination and then apply backwards induction. However, we are able to speed up solving the game, by 80% or more, by selectively *growing the game tree forward*.

To do so, we first calculate whether the first mover would earn positive profits as a nonstop carrier if it were the only carrier in the market, given its fixed cost.[84] If not, then we do not need to consider any of the branches where it provides nonstop service, immediately eliminating half of the game tree from consideration. If it is profitable, then we need to consider both branches. We then turn to the second carrier, and ask the same question, for each of the first carrier branches that remain under consideration, and we only keep the nonstop branch for the second carrier if nonstop service yields it (i.e., the second carrier) positive profits. Once this has been done for all carriers, we can solve backwards to find the unique subgame perfect equilibrium using the resulting tree, which usually has many fewer branches than the full game tree.

C.2 Specification of g, Random Variable Supports and Preliminary Estimation

Choice of g and W. The use of importance sampling assumes that the importance densities $g(\theta_m|X_m)$ and the distributions assumed by the model $f(\theta_m|X_m,\Gamma)$ have the same supports which do not depend on Γ, the parameters to be estimated. As discussed by Geweke (1989), consistency of the importance sampling estimator also requires that g is sufficiently similar to f that the variance of $y(\theta_{ms},X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$ is finite. These considerations lead us to use a multi-round estimation approach, as recommended by Ackerberg (2009), where we specify wide supports for the demand

[84] To be clear, this is not the same as testing whether nonstop service is more profitable than connecting service.

Table C.1: Description of g For the Final Round of Estimation

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

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

and cost draws, including all values that we believe may be relevant.[85]

In the first round we matched a subset of the price, share and service choice moments through straightforward experimentation to provide us with the initial parameterization reported in the notes to Table C.1, and we then ran two further rounds of estimation of the whole model, with the resulting estimates providing the $g(\theta_m|X_m)$ densities (reported in the table) that we use in the final round of estimation that produces the estimates reported in Section 5. The final round uses 2,000 importance draws for each route, with $S = 1,000$ used in estimation and samples from the full pool of 2,000 used when estimating standard errors using a bootstrap where routes are resampled.

The computational burden is reasonable for academic research: solving 2,000 games for 2,028 routes takes less than two days on a medium-sized cluster, and the parameters are estimated in one day on a laptop without any parallelization.[86]

We form the weighting matrix by using the results from the penultimate round of estimation

[85] The one exception to the rule of using wide supports is that we restrict the nesting parameter to lie between 0.5 and 0.9. This range covers most estimates from the existing literature (for example, Berry and Jia (2010) and Ciliberto and Williams (2014)). We experimented using the full range of [0,1], but found that the objective function often had local minima where the estimated nesting parameter was very close to 0 or very close to 1, but the fit of the moments was poor.

[86] In Roberts and Sweeting (2013) we bootstrap the entire multi-round procedure to calculate standard errors. In the current paper we are bootstrap the final stage, while acknowledging that the choice of g was informed by our initial attempts at estimation. See Li, Mazur, Park, Roberts, Sweeting, and Zhang (2018) for Monte Carlo evidence on how varying the gs affects the estimates.

(where we use an identity weighting matrix). As the number of moments (1,384) is large relative to the number of observations (16,130 carrier-route-directions) estimates of the covariances of the moments are likely to be inaccurate, so our final round uses a diagonal weighting matrix, with equal total weight on the groups of moments associated with price, share and service choice outcomes and, within each group, the weight on each moment is proportional to the reciprocal of the variance of that moment from the penultimate round.

C.3 Performance of the Estimation Algorithm For the Baseline Estimates

The use of importance sampling during estimation has two benefits: it greatly reduces the computational burden and it generates a smooth and continuous objective function.

Figure C.1 shows the shape of the objective function when we vary each parameter around its estimated value, holding the other parameters fixed. While these pictures certainly should not be interpreted as strong evidence that there is a global minimum in multiple dimensions, it is comforting that the objective function is convex in almost all dimensions.

C.4 Variance of the Moments

For an importance sample estimate of a moment to be consistent the variance of $y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms}|X_m)}$ must be finite (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 is likely to continue to jump wildly as S rises.

Figure C.2 shows these estimates of the sample variance for the moments associated with three market-level outcomes, namely the weighted nonstop fare, the weighted connecting fare and the quantity-based sum of squared market shares for the carriers in the market, based on the estimated parameters. The number of simulations is on the x-axis (log scale) and the variance of $\frac{1}{M}\sum y(\theta_{ms}, X_m)\frac{f(\theta_{ms}|X_m,\Gamma)}{g(\theta_{ms})}$ across simulations $s = 1, .., S$ is on the y-axis. Relative to examples in Koopman, Shephard, and Creal (2009), the jumps in the estimated sample variance are quite small for $S > 500$. In our application we are using $S = 1,000$.

C.5 Model Fit

Section 5.2 of the text briefly discusses the performance of the model at matching service choices. Table C.2 provides more detail of how well the model predicts service choices for carriers at some of their major hubs. In general, the model matches the fact that hub carriers serve most routes nonstop, although it does underpredict service at both Salt Lake City and Newark.

Table C.3 uses the same draws to show the fit of average prices and shares by type of service and by terciles of the market size distribution. We match average *differences* in market shares

Figure C.1: Shape of the Objective Function Around the Estimated Parameters For the Parameter Estimates in Column (1) of Tables 4 and 5 (black dot marks the estimated coefficient value)

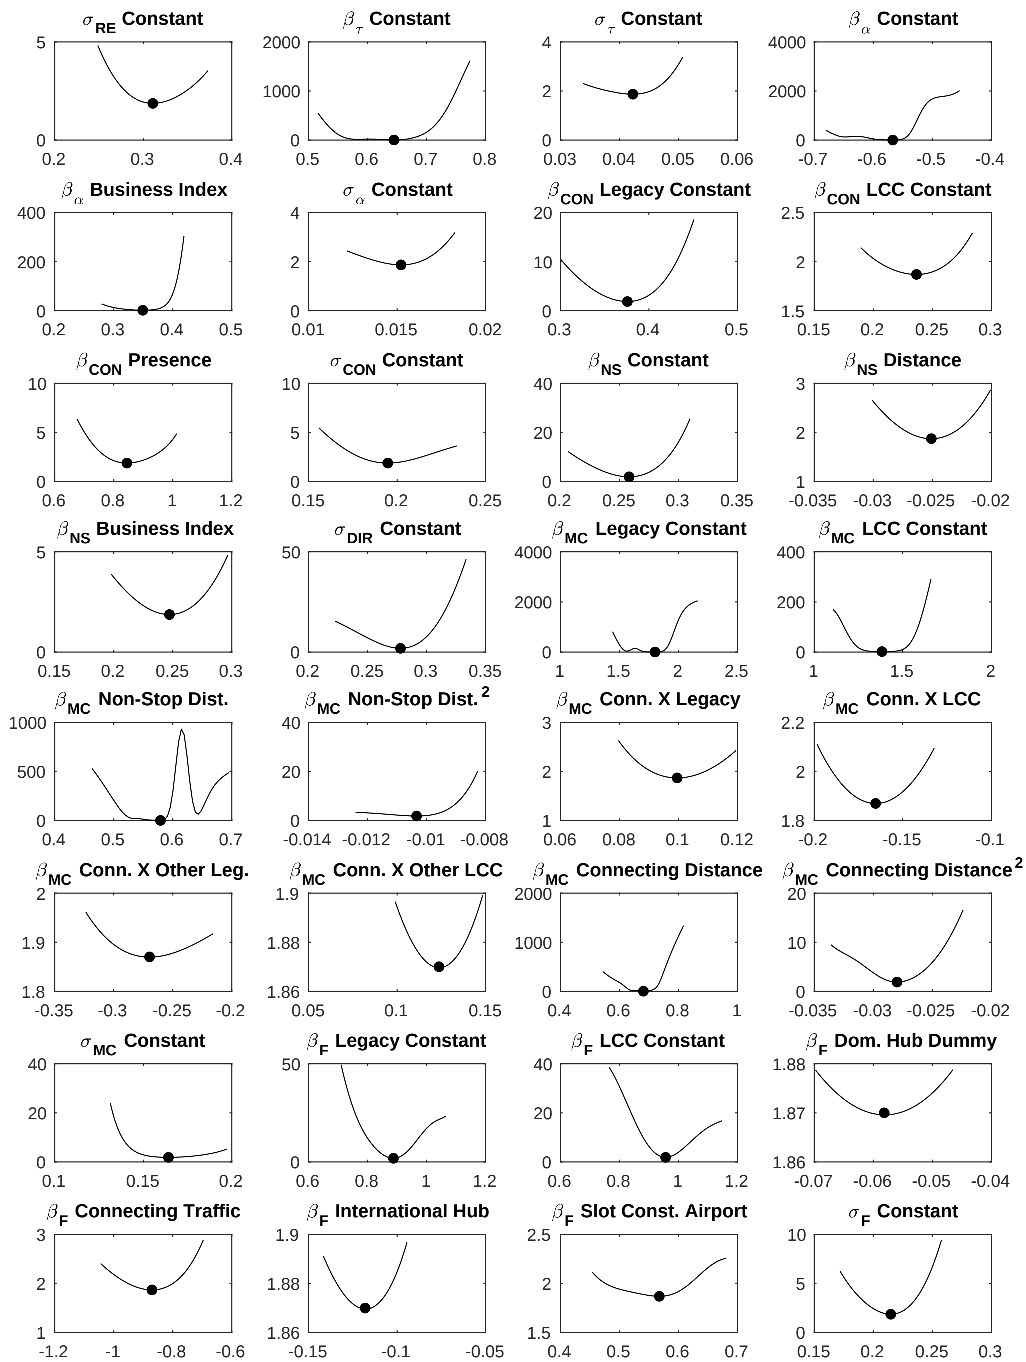

Figure C.2: Sample Variance of Three Moments as the Number of Simulation Draws is Increased (logarithm of the number of draws on the x-axis)

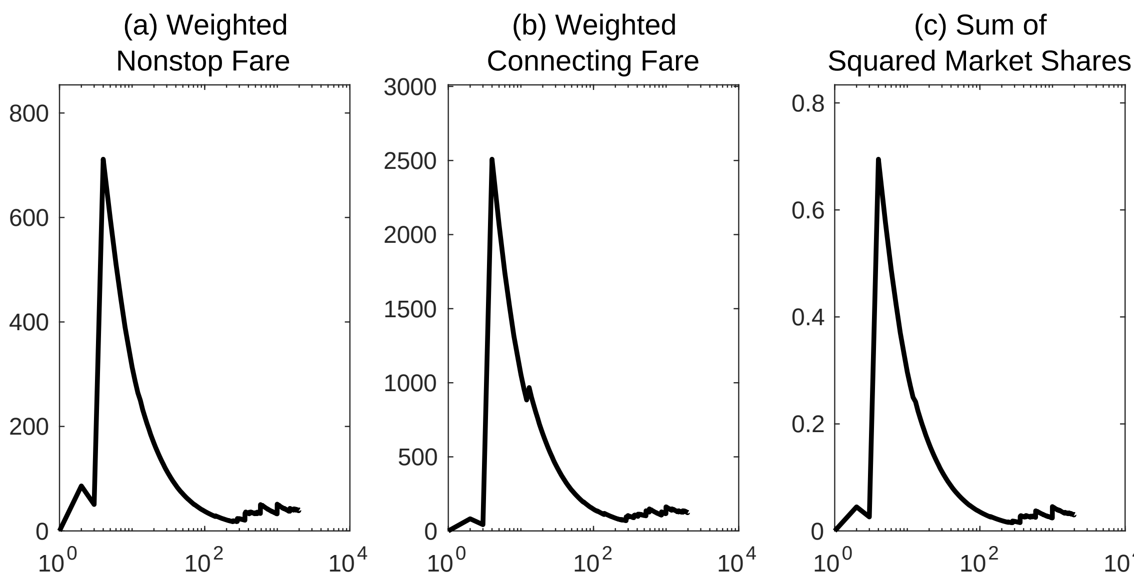

Table C.2: Model Fit: Prediction of Service Choices by Carriers at a Selection of Domestic Hubs

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

Notes: predictions based on the average of 20 simulated draws for each route using the estimated parameters in column (1) of Tables 4 and 5. Standard errors based on additional sets of 20 draws for each of the bootstrap estimates used to calculate standard errors in the same tables.

Table C.3: 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)
		Nonstop	$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 Shares	All Markets	Any Service	7.1%	8.4% (0.3%)
		Nonstop	17.9%	20.5% (0.9%)
		Connecting	4.9%	5.8% (0.3%)
	Market Size Groups			
	1st Tercile	Nonstop	25.6%	29.8% (2.4%)
		Connecting	8.6%	8.0% (0.4%)
	2nd Tercile	Nonstop	23.1%	26.6% (1.5%)
		Connecting	4.3%	5.5% (0.3%)
	3rd Tercile	Nonstop	15.9%	18.7% (0.8%)
		Connecting	1.8%	3.4% (0.3%)

Notes: see the notes to Table C.2.

and prices across service types very accurately, although we overpredict the levels of prices and market shares. This partly reflects our use of new draws to assess fit rather than the draws used in estimation, as the estimation draws provide a closer fit to levels as well.

C.6 Robustness of the Results to Reducing the Number of Moments

As mentioned in the text, we have repeated our estimation using only the 740 moments that are based on carrier-specific outcomes.

Estimates. Table C.4 shows our estimates from the main text and the estimates when we use the reduced number of moments. Most of the coefficients are very similar, and even where individual coefficients are different they have similar implications. For example, even though the individual coefficients measuring the incremental value of nonstop service change significantly, the implied mean value of the increment falls only from 0.299 to 0.268.

Fit. Table C.5 compares model fit for prices and market shares for the two sets of estimates. The predictions are very similar to each other.

Table C.4: Estimates Based on Different Sets of Moments (bootstrapped standard errors in parentheses)

				(1) Text Estimates from (from Table 4 and 5)		(2) Carrier-Specific Moments Only	
Demand: Route-Level Parameters							
Demand RE	Std. D.	σ_{RE}	Constant	0.311	(0.138)	0.377	(0.142)
Nesting Parameter	Mean	β_{τ}	Constant	0.645	(0.012)	0.641	(0.013)
	Std. D.	σ_{τ}	Constant	0.042	(0.010)	0.029	(0.008)
Demand Slope	Mean	β_{α}	Constant	-0.567	(0.040)	-0.591	(0.036)
(price in \$100 units)			Business Index	0.349	(0.110)	0.400	(0.101)
	Std. D.	σ_{α}	Constant	0.015	(0.010)	0.013	(0.008)
Demand: Carrier Qualities							
Carrier Quality for	Mean	β_{CON}	Legacy Constant	0.376	(0.054)	0.332	(0.049)
Connecting Service			LCC Constant	0.237	(0.094)	0.187	(0.094)
			Presence	0.845	(0.130)	0.910	(0.154)
	Std. D.	σ_{CON}	Constant	0.195	(0.025)	0.199	(0.030)
Incremental Quality	Mean	β_{NS}	Constant	0.258	(0.235)	0.000	(0.210)
of Nonstop Service			Distance	-0.025	(0.034)	-0.001	(0.039)
			Business Index	0.247	(0.494)	0.653	(0.483)
	Std. D.	σ_{NS}	Constant	0.278	(0.038)	0.334	(0.051)
Costs							
Carrier Marginal Cost	Mean	β_{MC}	Legacy Constant	1.802	(0.168)	1.713	(0.137)
(units are \$100)			LCC Constant	1.383	(0.194)	1.210	(0.135)
			Conn. X Legacy	0.100	(0.229)	0.107	(0.230)
			Conn. X LCC	-0.165	(0.291)	-0.150	(0.264)
			Conn. X Other Leg.	-0.270	(0.680)	-0.226	(0.147)
			Conn. X Other LCC	0.124	(0.156)	0.217	(0.151)
			Nonstop Dist.	0.579	(0.117)	0.654	(0.096)
			Nonstop Dist.^{2}	-0.010	(0.018)	-0.024	(0.016)
			Conn. Distance	0.681	(0.083)	0.732	(0.099)
			Conn. Distance^2	-0.028	(0.012)	-0.034	(0.012)
	Std. D.	σ_{MC}	Constant	0.164	(0.021)	0.153	(0.015)
Carrier Fixed Cost	Mean	β_F	Legacy Constant	0.887	(0.061)	0.878	(0.062)
(units are \$1 million)			LCC Constant	0.957	(0.109)	0.923	(0.113)
			Slot Const. Airport	0.568	(0.094)	0.530	(0.095)
	Std. Dev.	σ_F	Constant	0.215	(0.035)	0.223	(0.036)
Carrier Network			Dom. Hub Dummy	-0.058	(0.127)	0.000	(0.207)
Variables (offset			$\widehat{\text{Connecting}}$ Traffic	-0.871	(0.227)	-0.761	(0.281)
fixed costs)			Intl. Hub	-0.118	(0.120)	-0.355	(0.142)

Note: standard errors in parentheses based on a bootstrap where routes are re-sampled and simulations are drawn from a pool of 2,000 draws for each selected route.

Table C.5: Model Fit: Average Market Shares and Prices Based on Different Sets of Moments

				Model Predictions	
			Data	Text Estimates (Table C.3)	Carrier Moments
Average Prices (directions weighted by market shares)	All Markets	Any Service	$436	$455	$455
		Nonstop	$415	$436	$442
		Connecting	$440	$458	$459
	Market Size Groups				
	1st Tercile	Any Service	$460	$465	$466
	2nd Tercile	Any Service	$442	$460	$461
	3rd Tercile	Any Service	$412	$441	$442
Average Carrier Market Shares	All Markets	Any Service	7.1%	8.4%	8.5%
		Nonstop	17.9%	20.5%	21.5%
		Connecting	4.9%	5.8%	5.5%
	Market Size Groups				
	1st Tercile	Nonstop	25.6%	29.8%	30.4%
		Connecting	8.6%	8.0%	7.9%
	2nd Tercile	Nonstop	23.1%	26.6%	26.4%
		Connecting	4.3%	5.5%	5.2%
	3rd Tercile	Nonstop	15.9%	18.7%	18.7%
		Connecting	1.8%	3.4%	3.1%

Notes: Predictions from the model calculated based on twenty simulation draws from each route from the relevant estimated distributions.

Table C.6: Predicted Effects of a United/US Airways Merger, under the Baseline Merger Assumption, in Four Nonstop Duopoly Markets Based on Different Sets of Moments and the Conditional Distributions

	United/US Airways Nonstop Duopoly Routes		United/ US Airways Nonstop with Nonstop Rivals	
	Text Estimates (from Table 8)	Carrier Moments	Text Estimates (from Table 11)	Carrier Moments
Mean Pre-Merger United/ US Airways Price	$531.97	$531.97	$350.02	$350.02
Predicted Change in Nonstop Rivals Post-Merger	+0.10	+0.08	+0.05	+0.03
Mean Predicted Post-Merger Newco Price	$573.37	$574.29	$377.24	$377.55

Counterfactuals. Finally, we consider predicted price effects and service changes after a merger between United and US Airways. We compute predictions using the four routes where the United and US Airways were nonstop duopolists and American provided connecting service and the ten routes where United and US Airways were nonstop and there was another nonstop rival. We consider the case where we account for selection by forming conditional distributions, under our baseline merger assumption that the lower presence carrier is removed, so that our results correspond to row 2 of Table 8 and the third row of Table 11. The results from the text and the estimates using the smaller number of moments are almost identical.

C.7 Estimation Using Moment Inequalities

Our baseline estimates assume that carriers make service choices in a known sequential order, so that there is a unique equilibrium. An alternative approach is to allow for simultaneous choices, or an unknown order of moves, and estimate parameters based on moment inequalities. We present results based on this approach here.

The form of the inequalities is

$$h(y, X, Z, \Gamma) = \mathbb{E}\left[\begin{array}{c} y_m^{data} - \underline{\mathbb{E}(\widehat{y_m(X}, \Gamma))} \\ \\ \overline{\mathbb{E}(\widehat{y_m(X}, \Gamma))} - y_m^{data} \end{array} \otimes Z_m\right] \geq 0$$

where y_m^{data} are observed outcomes in the data and Z_m are non-negative instruments. $\underline{\mathbb{E}(\widehat{y_m(X}, \Gamma))}$ and $\overline{\mathbb{E}(\widehat{y_m(X}, \Gamma))}$ are minimum and maximum expected values for y_m given a set of parameters Γ. The minimum and maximum are formed by using the minimum and maximum values of the outcome across different equilibria or across orders for each simulated draw from the importance density. For example, if the outcome is whether firm A is nonstop, the lower bound (minimum) would be formed by assuming that whenever there are equilibrium outcomes where A is **not** nonstop, one of them will be realized, whereas the upper bound (maximum) would be formed by assuming that whenever there are equilibrium outcomes where A is nonstop, one of them is realized. We can also do the same type of calculation of minima and maxima for prices and market shares. If there is a unique outcome the minimum and maximum will be the same. The expected values of the minimum and maximum are calculated by re-weighting the different simulations in the same way that we do when assuming a known sequential order, and we form moments using the same outcomes and interactions that we use for our primary estimates. We note that our use of moment inequalities differs from how it has been used in some entry-type games, such as Eizenberg (2014) and Wollmann (2018), where selection on demand and marginal cost shocks is ruled out by assumption and the moments are based on an equation for fixed costs with an additive structural error.

The objective function that is minimized is

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

where t is a vector equal in length to the vector of moments, which sets equal to zeros the inequalities that are satisfied. W is a weighting matrix, and, as for the baseline estimates, we use a diagonal weighting matrix, dividing the moments into three groups (service choices, shares and prices). The sum of the diagonal components for each group equals one, with each element scaled so that it is proportional to the inverse of the variance of the moment evaluated at an initial set of estimates, which were calculated using the identity matrix.

Estimates. The ideal procedure for presenting the results of an estimation based on inequalities is to present confidence sets for coefficients because the coefficients may not be point identified. The construction of confidence sets is very difficult with large numbers of parameters and moments, and, as we have emphasized in the text, certain features of the data mean that we expect the parameters to be point identified even when we use inequalities in our setting.[87] Therefore in the right-hand column of Table C.7 we simply present the point estimates that we find minimize the objective function. These estimates are very close to the estimates from the text that are also reported in the table, which we view as confirming the result that we would expect given the nature of the game that we are looking at and the data at hand.

[87]Outcomes where no carrier provides nonstop service (the most common outcome in our data) will always be unique, and a necessary condition for there to be multiple equilibria is that at least two carriers do not have a dominant service strategy. In our setting, in the vast majority of routes there is no more than one carrier with intermediate probabilities of nonstop service based on a simple set of observables, which strongly suggests that multiplicity should be rare. See Appendix D.

Table C.7: Coefficient Estimates Based on Inequalities

				(1)		(2)
				Baseline Assumed Seq. Choice		Parameters Minimizing the Moment Ineq. Obj. Fun.
Demand: Route-Level Parameters						
Demand RE	Std. Dev.	σ_{RE}	Constant	0.311	(0.138)	0.350
Nesting Parameter	Mean	β_{τ}	Constant	0.645	(0.012)	0.647
	Std. Dev.	σ_{τ}	Constant	0.042	(0.010)	0.040
Demand Slope	Mean	β_{α}	Constant	-0.567	(0.040)	-0.568
(price in $100 units)			Business Index	0.349	(0.110)	0.345
	Std. Dev.	σ_{α}	Constant	0.015	(0.010)	0.017
Demand: Carrier Qualities						
Carrier Quality for	Mean	β_{CON}	Legacy Constant	0.376	(0.054)	0.368
Connecting Service			LCC Constant	0.237	(0.094)	0.250
			Presence	0.845	(0.130)	0.824
	Std. Dev.	σ_{CON}	Constant	0.195	(0.025)	0.193
Incremental Quality	Mean	β_{NS}	Constant	0.258	(0.235)	0.366
of Nonstop Service			Distance	-0.025	(0.034)	-0.041
			Business Index	0.247	(0.494)	0.227
	Std. Dev.	σ_{NS}	Constant	0.278	(0.038)	0.261
Costs						
Carrier Marginal Cost	Mean	β_{MC}	Legacy Constant	1.802	(0.168)	1.792
(units are $100)			LCC Constant	1.383	(0.194)	1.331
			Conn. X Legacy	0.100	(0.229)	0.134
			Conn. X LCC	-0.165	(0.291)	-0.077
			Conn. X Other Leg.	-0.270	(0.680)	0.197
			Conn. X Other LCC	0.124	(0.156)	0.164
			Nonstop Distance	0.579	(0.117)	0.589
			Nonstop 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 $1 million)			LCC Constant	0.957	(0.109)	1.015
			Slot Const. Airport	0.568	(0.094)	0.602
	Std. Dev.	σ_F	Constant	0.215	(0.035)	0.198
Carrier Network			Dom. Hub Dummy	-0.058	(0.127)	-0.140
Variables (offset			Log($\widehat{\text{Connecting}}$ Traffic)	-0.871	(0.227)	-0.713
fixed costs)			International Hub	-0.118	(0.120)	-0.168

Notes: standard errors for the baseline, in parentheses, are based on 100 bootstrap replications where 2,028 routes are sampled with replacement, and we draw a new set of 1,000 simulation draws (taken from a pool of 2,000 draws) for each selected route. The Log(Predicted Connecting Traffic) variable is re-scaled so that for routes out of domestic hubs its mean is 0.52 and its standard deviation is 0.34. Its value is zero for non-hub routes. Distance is measured in thousands of miles.

D Multiple Equilibria, Identification and the Explanatory Power of Observed Variables for Service and Entry Choices

One of our striking results is that, at the estimated parameters, less than 2% of simulations from our model could support a different equilibrium outcome (i.e., different service choices) if we allowed for simultaneous moves or any alternative sequential order. As a result, it is not surprising that our coefficient estimates are very similar when we allow for these alternative possibilities (Appendix C.7). Several scholars have commented to us that they find this result surprising given earlier work examining airline entry decisions, notably Berry (1992) and Ciliberto and Tamer (2009), that has found that assumptions about the timing of decisions can affect estimates quite dramatically and that it is common for a simultaneous move game to support multiple different outcomes as equilibria (for example, Ciliberto and Tamer find this is true for 95% of their simulations). In this Appendix, we explain why models estimated using service choices and entry decisions, as defined in the existing literature, can differ so much on this dimension.

We define a carrier to be nonstop based on the number of nonstop flights that a carrier has per quarter (at least 64 in each direction to be defined as nonstop) and the proportion of passengers carried that travel direct (at least 50% without a change of planes). Other carriers are connecting. Carriers that provide nonstop service serve many more passengers than connecting carriers: the median nonstop (named) carrier serves over 1,000 round-trip passengers in DB1 (which is a 10% sample), whereas the median connecting carrier serves only 38 round-trip passengers, and, as noted in Appendix B, there are few carriers close to the 64 or 50% thresholds.[88] Our counterfactuals focus on mergers of nonstop carriers, as an analysis with fixed products indicates that these mergers tend to lead to the largest price increases unless rivals reposition.

In contrast, in Berry (1992) and Ciliberto and Tamer (2009), a carrier is defined as an entrant if it carries, by any type of service, a relatively small number of passengers in a quarter (for example, 20 DB1 passengers in Ciliberto and Tamer (2009)). In the data, there are many carriers with passenger counts that are right around these thresholds: the 25th percentile number of connecting passengers is 14 and the median is 38. Given this pattern and the sampling error in the DB1 sample, it is naturally quite difficult to predict which connecting carriers will be counted as entrants on a particular route.

We illustrate how well our data explains service choices and entry by estimating several probit specifications where the dependent variable are indicators for nonstop service or entry and the explanatory variables are the observed characteristics of the carrier and market/route characteristics (such as the average directional market size). The results are reported in Table D.1.

In the first five columns, the dependent variable is equal to one if the carrier is nonstop, and we use the 8,065 carrier-route observations in our data. The regressors in column (1) are the average of our market size measure across directions and the observable carrier and carrier-network variables

[88] The statistics discussed in this paragraph are for the named carriers we use, and not the composite Other Legacy and Other LCC carriers.

that we include in our specification of fixed costs. Despite the simplicity of the specification the pseudo-R^2 is 0.52. Column (2) adds the business index measure which we allow to shift the price coefficient and the preference for nonstop service. It is statistically significant but barely improves the fit. Column (3) replaces our market size measure with the geometric average population measure that is most commonly used in the literature: the pseudo-R2 decreases to 0.45, indicating that this is a poor alternative to our market size measure (a result which is consistent with the results presented in Appendix Table B.1). Column (4) adds measures of the carrier's presence at each endpoint, which we allow to affect demand, to the second specification, and the pseudo-R^2 increases to 0.65. In column (5) we include interactions between a number of the variables in the specification (as noted beneath the table) as well as measures of the number of rival carriers, and we find the pseudo-R^2 increases to 0.73.

In column (6) we consider instead the decision to enter a route (i.e., to provide either type of service) among the carriers that provide service (to any destination) at both airport endpoints and use a specification similar to column (4). This is the type of binary outcome modeled in in Berry (1992), Ciliberto and Tamer (2009) and Ciliberto, Murry, and Tamer (2020). The pseudo-R^2 is *much* lower (0.134).

What is the implication of these results for whether our model should be expected to support multiple equilibrium outcomes? A game with binary discrete choices can only support multiple outcomes if the more profitable option depends on what other players do for at least two of the players (i.e., at least two players do not have a dominant strategy). Intuitively, players are much less likely to be on the margin between different options when observed variables (that do not reflect what their rivals choose) strongly predict what their service choices will be. The service choice and entry models are clearly very different in this regard.

To illustrate, Figure D.1(a), shows the distribution of predicted probabilities for a carrier providing nonstop service using 40 bins based on column (4). We observe that the predicted probabilities are concentrated either very close to zero or very close to one. Defining intermediate as predicted probabilities between 0.05 and 0.95 based on the column (4) estimates, there are 482 routes (less than 24% of the total) where two or more carriers have intermediate nonstop service probabilities (using thresholds of 0.1 and 0.9, 302 routes would have at least two carriers with intermediate probabilities). In contrast, the predicted probabilities for entry choices, shown in Figure D.1(b) (based on column (6)), lie mainly in the range from 0.2 to 0.8, and 96% of routes have two or more carriers with intermediate entry probabilities. When we perform the exercise of counting how many different outcomes our parameter estimates can support under different timing assumptions, discussed in Section 5.3, we can see the connection between the predicted probabilities of nonstop service in these simple regressions and the multiplicity of equilibrium outcomes: the probability of a simulation draw for one of the 482 intermediate probability routes supporting multiple outcomes is two-and-half times higher than for the remaining routes.

The service choice probit results also have implications for the identification of the model. As

Table D.1: Probit Models of Carrier Service Choice and Entry Decisions

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Nonstop	Nonstop	Nonstop	Nonstop	Nonstop	Enter
Low Cost Carrier	0.808	0.808	0.782	0.537	1.681	0.514
	(0.051)	(0.0516)	(0.0476)	(0.0685)	(0.395)	(0.0376)
Slot Constr. Airport	0.559	0.587	0.724	0.541	0.232	-0.207
	(0.095)	(0.0961)	(0.0927)	(0.112)	(0.132)	(0.0650)
Carrier Intl. Hub	0.940	0.946	0.836	0.0385	-0.165	0.158
	(0.074)	(0.0748)	(0.0738)	(0.0894)	(0.113)	(0.0801)
Carrier Dom. Hub	-6.090	-6.161	-6.942	-6.578	-34.24	-3.740
	(0.645)	(0.647)	(0.623)	(0.648)	(47.37)	(0.627)
Carrier Pred. Connecting	1.341	1.355	1.464	1.160	5.701	0.611
Traffic Measure	(0.107)	(0.107)	(0.104)	(0.108)	(7.932)	(0.106)
Route Business Index		-0.663	-1.364	0.198	0.670	-0.126
		(0.293)	(0.268)	(0.348)	(0.387)	(0.142)
Our Market Size	1.614	1.595		2.019	-0.176	-0.0552
/10,000	(0.064)	(0.0649)		(0.0828)	(0.671)	(0.0405)
Geom. Avg. Pop.			0.0122			
/10,000			(0.00112)			
Carrier Max.				3.543	4.334	1.622
Endpoint Presence				(0.144)	(0.626)	(0.109)
Carrier Min.				1.916	6.814	4.424
Endpoint Presence				(0.276)	(2.510)	(0.266)
Number Rival					-0.167	
Carriers in Market					(0.0237)	
Number Rival Low Cost					0.167	
Carriers in Market					(0.0663)	
Constant	-2.335	-2.065	-1.581	-3.930	-4.131	-0.312
	(0.044)	(0.127)	(0.115)	(0.177)	(0.387)	(0.0662)
Variable interactions	N	N	N	N	Y	N
Observations	8,065	8,065	8,065	8,065	8,065	12,550
Pseudo-R2	0.521	0.522	0.450	0.653	0.726	0.134

Notes: standard errors in parentheses. Observations in columns (1)-(5) are the carrier-route observations that are included in our estimation dataset. Our Market Size is the average of our market size estimate across directions. Geom. Avg. Pop. is the geometric average of the MSA endpoint populations, a popular alternative measure of market size. We measure carrier presence (the number of routes served nonstop by the carrier out of the total number of routes served nonstop by any carrier) at the carrier-airport level and include the higher and lower values separately in the regressions. Observations in column (6) include the observations in our estimation dataset plus observations for carrier-routes where the carrier provides some service at both endpoints but does not meet our criteria for being a competitor on the route in question. The interactions that are included in column (6) are between LCC, domestic hub, the predicted connecting traffic, market size and the two presence measures.

Figure D.1: Predicted Probabilities of Carrier Service Choices (based on Table D.1, column (4)) and Entry Decisions (based on Table D.1, column (6))

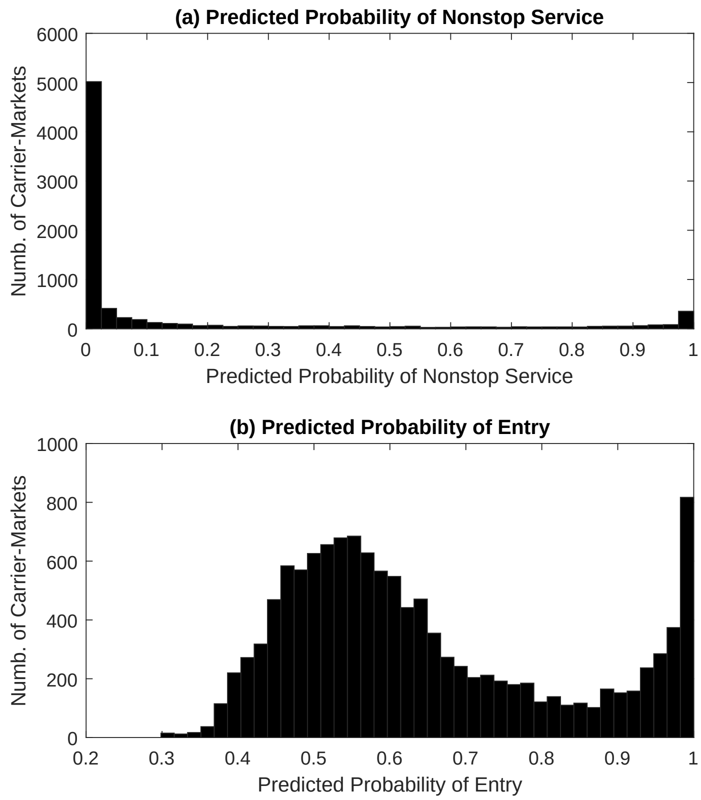

discussed in Section 4, one argument for why the demand and marginal cost parameters are point identified is that there are a large number of routes and carriers for which observed covariates essentially determine their service choices so that there should be (almost) no selection on unobservable demand or marginal cost shocks when they make these choices. Based on the column (4) estimates, 58% of route-carriers predicted nonstop service probabilities are less than 0.01 or more than 0.99, meaning that we have a large number of observations where selection based on demand and marginal cost unobservables is likely to be minimal, and conventional identification arguments should apply.