Simulating the Dynamic Effects of Horizontal Mergers: U.S. Airlines*

C. Lanier Benkard
Stanford University
and NBER

Aaron Bodoh-Creed
U.C. Berkeley

John Lazarev
University of Pennsylvania

This version: August 2020

Abstract

We propose a simple method for studying the medium- and long-run dynamic effects of horizontal mergers that builds on the two-step estimator of Bajari, Benkard, and Levin (2007). Policy functions are estimated on historical pre-merger data, and then future industry outcomes are simulated both with and without the proposed merger. We apply our method to two recent airline mergers as well as one that was proposed but blocked. We find that low-cost carriers play a crucial role in creating offsetting entry. In some cases (United-US Airways), the model predicts substantial scope for offsetting entry, while in others (Delta-Northwest) it does not. Thus, the dynamic analysis is complementary to and yields different conclusions than the static analyses.

*The first draft of this paper was March 2009. We thank Steve Berry, Severin Borenstein, Phil Haile, Darin Lee, and Jon Levin for their useful input.

Correspondence: lanierb@stanford.edu; acreed@haas.berkeley.edu; jlazarev@sas.upenn.edu

1 Introduction

Antitrust enforcement is tasked to protect and promote competition. Despite being a vital part of industry dynamics, mergers—especially those between direct competitors—may present a consequential threat to competition within the industry. The central question of merger review is whether the effect of a merger "may be substantially to lessen competition."[1]

The static effects of a horizontal merger are unambiguous. By definition, horizontal mergers between firms leave the industry with fewer competitors and, therefore, less competition. Whether competition is lessened *substantially* depends on the role that the merging firms play in the competitive process. A merger between large firms producing close substitutes in a concentrated industry is more likely to lessen competition that a merger between smaller firms that produce indirect substitutes in an industry with many active competitors.

The simplest method to evaluate the static effects of a merger on competition is to compute pre- and post-merger concentration measures assuming no post-merger changes in market shares. Large increases in concentration are presumed to lessen competition substantially (Shapiro (1996), US Department of Justice (2010)). More sophisticated methods (Berry and Pakes (1993), Berry, Levinsohn, and Pakes (1995), Nevo (2000)) are available for analyzing mergers in markets with differentiated products, where competition between firms depends critically on the precise characteristics of each firm's array of products. These methods provide richer models of changes in post-merger prices and market shares, but still rely on a static model that holds the set of incumbent firms and products fixed.

Antitrust enforcement, however, recognizes the fact that competition is inherently a dynamic process. For that reason, a static analysis may not be able to capture the full effect of a merger on the competitiveness of the industry. In general, static models do not account for post-merger changes in firms' behavior. By changing firms' incentives, a merger might lead to different levels of entry, exit, investment, and pricing than occurred pre-merger in both merging and non-merging firms (Berry and Pakes (1993), Gowrisankaran (1999)). Lastly, several papers have shown that dynamics can weaken the link between market structure and performance (Berry and Pakes (1993), Pakes and McGuire (1994), Ericson and Pakes (1995), Gowrisankaran (1999), Fershtman and Pakes (2000), Benkard (2004)), making the pre-/post-merger snapshot of market concentration and markups less relevant for understanding the medium- and long-run effects of a merger on competition. A merger that looks "bad" in the short run may, nevertheless, fall short of having a substantial negative impact on competition in the long run if it creates a profitable opportunity

[1] Section 7 of the Clayton Act of 1914.

for new firms to enter the market and reduce market concentration.

When two airlines merge, the industry loses two smaller airlines and gains a new, bigger one. In the short run, markets served by both of the merging airlines are directly affected by the merger as these markets immediately lose a direct competitor. Whether or not this loss attracts a new entrant is ultimately an empirical question. Additionally, since the incentives of the merged airlines change, in the medium run the merger may affect competition in markets without pre-merger overlap. For example, since the merged airlines serve a larger network of markets, previously unprofitable routes may become more attractive for entry. Once passengers are flown in to a hub, they can now reach more destinations without having to change their carrier. At the same time, the combined network of the merged airlines is likely to contain redundant routes. That could incentivize the merged airlines to decrease or even eliminate some duplicative routes in the medium or long run. Again, whether or not these dynamic effects lessen competition substantially, requires an empirical analysis.

The contribution of this paper is to develop a simple set of techniques for analyzing the potential dynamic effects of a merger and to apply these techniques to three recently proposed mergers in the U.S. airline industry. Importantly, since our method does not require estimating the fundamental parameters of an underlying structural model, or solving for a counterfactual equilibrium, it is a practical method for regulators to use given the tight time-frame under which mergers are evaluated.

We start with the standard framework that most existing work on empirical dynamic oligopoly has relied upon – that of Ericson and Pakes (1995) (hereafter EP).[2] This framework models a dynamic industry in Markov perfect equilibrium (MPE). Equilibria of this model cannot be found analytically, so they must instead be computed numerically. In general, inserting mergers into this framework requires a detailed model of how mergers occur (as in Gowrisankaran (1999)), resulting in a complex model that is difficult to compute or apply to data.

To overcome this challenge, we extend the insight of Bajari, Benkard, and Levin (2007) (hereafter BBL). Put in the broadest terms, the main insight of BBL is the following: the equilibrium to a complicated dynamic game can, under certain assumptions, be directly observed in the data. Instead of repeatedly solving the model for different values of unknown parameters trying to find an equilibrium that closely matches the observed data, one can proceed it two steps. The first step is to estimate the equilibrium of the model, i.e. the firms' strategy functions, directly from the data. The second step is to plug these estimated strategies into the equations defined by the equilibrium conditions of the model, which can typically allow one to solve for the unknown parameters much faster.

[2]For an overview, see Ackerberg, Benkard, Berry, and Pakes (2007).

Building on BBL, this paper shows that it is possible to quantify the dynamic effects of a horizontal merger in the EP/BBL framework by adding just one simplifying assumption: we assume that the equilibrium being played does not change after the merger, in the sense that firms' *strategy functions* remain the same. Under this assumption, we can use BBL's forward-simulation procedure to simulate the distribution of future industry outcomes both with and without the merger. We can then compute medium- and long-run concentration measures to evaluate the differential impact of the merger on competition. This allows us to incorporate the effects of offsetting entry into the (static) analysis of concentration commonly used as part of antitrust evaluations.

The proposed assumption is both restrictive and testable. For example, this assumption would hold if mergers are a standard occurrence in equilibrium. Alternatively, it might happen if mergers are very rare, so that equilibrium play is not strongly affected by the likelihood of future mergers, whether or not the merger in question happens. On the other hand, the assumption would not hold in the event that allowing the proposed merger would represent a substantive change in antitrust policy. In that case, the fact that the merger is allowed to go through might change firms' unobserved beliefs about future play, changing their strategy functions. Ultimately, whether this assumption is valid depends both on the counterfactual in question and on the flexibility of the model that defines the data generating process. In more flexible models, the assumption that the counterfactual stays on the equilibrium path can be both internally consistent and supported by the data. In less flexible models, this assumption can fail on both grounds. Thus, our approach requires an extra step to validate this assumption, but the benefit is that our approach is vastly simpler than the alternative of computing a new post-merger equilibrium to the game, an option that, while attractive, could be computationally challenging in most cases.

Note that our methods are not intended to replace static antitrust analyses, described in Shapiro (1996) and Nevo (2000), which seek to measure the short-run effects of a proposed merger on prices, market shares, etc. When used in isolation, our methods generate predictions about the medium and long term effects of a merger on industry structure, which can often be highly informative about the anti-competitive effects of a merger. For example, our dynamic analysis can reveal whether or not off-setting entry is likely to reduce short-run market competition in particular markets, which might determine whether a merger is presumed to be anti-competitive in the long run. However, industry structure may only be a part of the story. Occasionally, a horizontal merger may generate both pro- and anti-competitive effects. To balance them, it may be necessary to quantify the welfare effects of these predictions, which would require an explicit, short-run model of consumer demand and market supply. Thus, in our opinion, merger analyses should include *both* of these tools.

We apply the proposed method to the U.S. airline industry, which has recently gone through a number of large and somewhat controversial mergers. Our focus is narrow. We do not attempt to perform a full-scale prospective analysis of the proposed mergers. Instead, we want to evaluate the differential impact of each of the proposed mergers on the yearly entry/exit decisions of each airline at the level of individual non-stop routes between pairs of metropolitan areas.

Our estimates are based on data from 2003 to 2008. We chose 2003 as the starting period of the analysis to avoid the immediate aftermath of the terrorist attacks on September 11, 2001. We chose 2008 as the end period both because it was the last period prior to the recent wave of mergers, and because it avoids the severe economic downturn experienced during the 2008 financial crisis.

We use these data to estimate the airlines' strategy functions, which are maps from the current state of the market competition game into a decision about which network of nonstop routes to serve in the following year. We validate our simplifying assumption and conclude that the estimated strategy functions do not reject it. These strategy functions exhibit two important features. First, we estimate competition effects that are quite large. For example, we find that, in a market with two incumbent carriers, for a potential entrant indifferent between entering and not, the exit of one incumbent increases the probability of entry by 34% (from 50% to 84%). As a comparative benchmark, the exit of one incumbent competitor is estimated to be roughly half as large as the incumbency effect of airline presence on a given route, an effect typically thought to be large. Estimated competition effects are larger in markets with fewer incumbents, and smaller in markets with more incumbents.

These large competition effects give the model the potential to generate offsetting entry after a merger. However, whether or not there is offsetting entry on a given city pair also depends critically on the availability of potential entrants whose nonstop route network rationalizes entry for that city pair. Similarly to Berry (1992), we find that the size of an airline's network at each end of the route (as measured by how many cities it serves from each end) is an important determinant of entry. For a potential entrant indifferent between entering a given route and not, we find that one additional city served from an endpoint city increases the probability of entry on the route by about 10%.

In addition, competition at the end point cities matters. Airlines with higher market shares at end point cities are more likely to enter a given city pair. Our model does not allow us to determine whether these effects are driven by demand or cost factors. Airlines are also more likely to enter a city pair if an endpoint is an own hub, and less likely to enter if an endpoint is a competitor hub. Airlines are also more likely to enter nonstop routes when there is no convenient alternative one-stop itinerary in their current network.

To summarize, due to the large competition effects, the empirical model is likely to predict offsetting entry after a merger on routes where there is at least one realistic potential entrant with a rich route network in the vicinity of the route in question. However, the existence of realistic potential entrants is far from guaranteed and varies widely in the U.S. airline network.

We consider three mergers: United-US Airways (UA-US), Delta-Northwest (DL-NW), and United-Continental (UA-CO). The UA-US merger was proposed in 2000 and rejected by antitrust authorities. The DL-NW merger was approved in late 2008. The UA-CO merger was approved in late 2010.

According to static measures of concentration, all three of these proposed mergers had great potential to harm consumers. For example, DL-NW and UA-US created seven and six new monopoly non-stop routes respectively. All three mergers also created large increases in concentration for at least a few cities when viewed as a whole, often at existing hubs where concentration was already high. Absent offsetting entry or large cost synergies, these effects point to a high likelihood of price increases and consumer harm.

We use our estimated strategy functions to simulate the 10-year effects of the three mergers. We find that the UA-US merger would have had substantial potential for offsetting entry. From a static perspective, United and US Airways had significant network overlap, particularly in the DC and Philadelphia areas, so short-run reductions in competition would have been significant. However, because the airlines' networks overlapped in areas with heavy low cost carrier presence, our model suggests that the anticompetitive effects likely would have been eliminated within a few years by low cost carrier entry, particularly by Southwest and JetBlue. This merger was blocked based on a static analysis, but our results indicate that a dynamic analysis would have led to a different conclusion.

Like UA-US, the DL-NW merger was predicted to have a strong short-run anticompetitive effect. Unlike UA-US, our simulations show very little scope for offsetting entry in this case. Our model suggests that the markets where the merger leads to reductions in competition are well insulated from entry. The UA-CO merger appears somewhat more benign than the other two mergers in the short run, but our simulations again suggest there was very little scope for offsetting entry. Ironically, it is these last two mergers that were approved and executed.

The rest of the paper is organized as follows. Section 2 discusses the related literature and alternatives it has proposed. We give a detailed description of our method in Section 3. Section 4 shows how to apply our method to the U.S. airline industry. Section 5 describes our data and provides the conclusions of a simple static merger analysis. In section 6, we show how to use econometric and machine-learning techniques to estimate the airlines' strategy functions. We look for but find no empirical evidence that would invalidate the main simplifying assumption that this paper relies upon. In section 7, we evaluate the dynamic effects of the proposed airline mergers. Section 8 concludes.

2 Alternative Approaches and Related Literature

Academic literature and policymakers alike recognize the fact that a horizontal merger may not always substantially lessen competition and that sometimes a static analysis may overestimate its negative effects.

Two recent papers—Li, Mazur, Park, Roberts, Sweeting, and Jun (2018) and Ciliberto, Murry, and Tamer (2018)—seek to address economic questions that are similar to ours. They, however, build on a different set of tools to perform their analysis. Specifically, they use a cross-section of airline markets to estimate a complete information two-period entry/exit model. They use this model to evaluate the likelihood of post-merger entry. A special emphasis is placed on the importance of the unobservable market characteristics that may introduce selection in the airlines' entry and exit decisions.

Both their and our approaches recognize the same challenge: airlines' decisions to enter or exit may be driven by multiple complicated factors. The solutions, however, are somewhat different. The two-period entry/exit papers compress these complicated factors into a single unobservable variable. The distribution of this unobservable then becomes the object of interest. Once this distribution is recovered, the model is solved for a counterfactual industry structure in which the merger is exogenously executed. In contrast, our approach recognizes the fact that panel data may provide superior information on what these underlying unobservable factors may be. Specifically, we employ machine-learning techniques that allow us to make the unobserved part of the entry/exit decisions as much observable as possible. A cross-section analysis, due to data limitations, would not have allowed us to achieve this goal. Second, the two-period entry/exit papers have to assume that the unobserved part would stay the same whether or not the merger goes through. Depending on which factors contribute to this unobservable, their assumption may be problematic. For example, if the main force behind the unobservable part is the overall size of the airline's network, then the merger will necessarily affect it. Our approach controls for these factors explicitly. Third, unlike the two-period papers, our approach recognizes the fact that mergers may be endogenous (Nocke and Whinston (2010)), i.e. a part of the equilibrium industry dynamics. Finally, despite being built on a more complicated dynamic oligopoly model, our analysis is computationally simpler and faster. It is especially suited to data-rich industries that exhibit rich variation in market structures such as the U.S. airline industry.

There are several other related papers in the literature that we have not mentioned yet. Collard-Wexler (2014) uses a Bresnahan and Reiss-style empirical dynamic model to evaluate the hysteresis effects of a merger from duopoly to monopoly in the ready-mix concrete industry. The paper finds that merger to monopoly would generate about 15 years of monopoly. The approach in the paper is similar to ours, but is even simpler than ours as it assumes homogeneous firms.

Three other recent papers (Jeziorski (2014a), Jeziorski (2014b), and Stahl (2009)) use dynamic models similar in spirit to ours to consider recent merger waves in radio and broadcast television respectively. However, the goals of these papers are quite different from ours. They use data on past mergers primarily to evaluate the forces that drove the merger waves, but also to evaluate (ex post) the welfare effects of the merger waves. Our paper instead focuses on the potential future dynamic effects of proposed mergers.

There are also several papers looking at past airline mergers. Most notably, Borenstein (1990) evaluates (ex post) the anticompetitive effects of two airline mergers that occurred in the mid-1980s, each of which led to substantially increased concentration at a major hub. He finds that there is evidence of both price increases and capacity reductions at these hubs after the mergers. Kim and Singal (1993) does a broader ex post evaluation of fourteen airline mergers in the 1980s. Overall they find that after a merger both the merged and unmerged firms substantially increased fares. Peters (2006) also does an ex-post evaluation of static merger simulations (as in Nevo (2000)) using five airline mergers from the mid-1980s. He finds that the standard model does not do very well at predicting the price effects of these mergers, and appears to omit some important supply-side factors (e.g., cost or conduct).

There are also some important results in the literature regarding airline network structure and airline competition that are relevant to our work. Borenstein (1991) finds evidence that a carrier that has a dominant market share of flights out of a given city has increased market power on routes out of that city, even on individual routes where there may be substantial competition. Borenstein (1989) similarly shows that both an airline's market share on an individual route and its share at the endpoint cities influence its ability to mark up price above cost. Our results echo these findings.

Berry (1992) estimates a static model of airline entry with heterogeneous firms and finds, similarly to Borenstein (1989), that an airline's market share of routes out of a given city is an important determinant of entry into other routes from that city. Ciliberto and Tamer (2009) estimates a static entry model that allows for multiple equilibria and for asymmetric strategies. Boguslaski, Ito, and Lee (2004) estimates a static entry model for Southwest that fits the data extremely well and helped inspire some features of our model, such as the way we define entry and exit. Using the rapid expansion of Southwest to generate variation in the level of threat of route-level entry, Goolsbee and Syverson (2008) shows that even the *threat* of entry causes incumbent airlines to drop their fares significantly. Other relevant static airline entry papers include Sinclair (1995) and Reiss and Spiller (1989).

Another recent paper (Aguirregabiria and Ho (2012)) estimates a structural dynamic oligopoly model of airline entry that is similar to our model, and computes equilibrium entry strategies for airlines. Our approach is simpler and less ambitious. However, an advantage of taking a simple approach is that we can include a richer set of airline network state variables in our model, potentially allowing for more robust network-wide route optimization on the part of firms, rather than focusing on one route at a time in isolation from the broader network. In addition, we can avoid some of the simplifications (e.g., the use of inclusive values, the assumption that route entry is determined by route-specific profits) that is required to make the estimation the model of Aguirregabiria and Ho (2012) tractable.

3 Notation and Methodology

We start with a brief characterization of our general approach. Our hope is that the approach is simple enough to be used in a wide variety of settings by practitioners and academics. We apply the approach to airlines in the sections that follow.

3.1 The General Model

The general model closely follows BBL and is a generalization of the EP model. The defining feature of the model is that actions taken in a given period may affect both current profits and, by influencing a set of commonly observed state variables, future strategic interaction. In this way, the model can permit many aspects of dynamic competition, such as entry and exit decisions, mergers, learning, product entry and exit, investment, dynamic pricing, bidding, etc.

There are N firms, denoted $i = 1, ..., N$, that make decisions at times $t = 1, 2, ..., \infty$. Conditions at time t are summarized by a commonly observed vector of state variables $\mathbf{s}_t \in S \subset \mathbb{R}^L$. Depending on the application, relevant state variables might include the firms' production capacities, their technological progress up to time t, the current market shares, stocks of consumer loyalty, or simply the set of incumbent firms.

Given the state $\mathbf{s}_t$, firms choose actions simultaneously. These actions might include decisions about whether to enter or exit the market, investment or advertising levels, or choices about prices and quantities. Let $a_{it} \in A_i$ denote firm i's action at time t, and $\mathbf{a}_t = (a_{1t}, \ldots, a_{Nt}) \in A$ the vector of time t actions. For notational simplicity, we denote a_{it} as a scalar. However, there is no reason that it cannot be vector valued. We will assume that both actions $\mathbf{a}_t$ and states $\mathbf{s}_t$ are observed by the researcher.

We assume that before choosing its action, each firm i receives a private shock ν_{it}, drawn independently across agents and over time from a distribution $G_i(\cdot|\mathbf{s}_t)$ with support $\mathcal{V}_i \subset \mathbb{R}^M$. The private shock might derive from variability in marginal costs of production, profits, or sunk costs of entry or exit. We denote the vector of private shocks across firms as $\nu_t = (\nu_{1t}, ..., \nu_{Nt})$. Again, we have denoted ν_{it} as a scalar, but there is no reason that it cannot be vector valued. We assume that ν_{it} is not known to the researcher.

The assumption of iid private shocks is extremely troublesome in this context. In many empirical applications, some state variables are not observed by the researcher, which could lead to serial correlation in the unobserved shocks. An example would be a serially correlated unobserved demand shifter. In the empirical work we will address this issue by both using machine learning techniques to uncover the relevant observable state and also testing for the presence of serial correlation in the unobserved shock. There is also ongoing research in this area aimed at generalizing these approaches.[3]

To complete the model, BBL and EP outline primitives of the dynamic oligopoly model that determine period profits and the evolution of states. We assume that the state at date $t+1$, denoted $\mathbf{s}_{t+1}$, is drawn from a probability distribution $Q(\mathbf{s}_{t+1}|\mathbf{a}_t, \mathbf{s}_t)$. The dependence of $Q(\cdot|\mathbf{a}_t, \mathbf{s}_t)$ on the firms' actions $\mathbf{a}_t$ means that time t behavior may affect the future strategic environment. This would be the case, for example, for entry/exit decisions or long-term investments. In some applications, some details of the state transition function, such as the investment technology, might also be assumed to have a specific structure. Other aspects of transitions, such as the Markov process determining aggregate demand, might be exogenous and specified quite freely. Others may even be deterministic, as in the case of firm age.

BBL and EP also specify in detail a period profit function, investment process, and entry and exit processes. While these are important fundamentals of the model, we will omit them here for brevity and because, as we will see, in our approach it is possible to proceed without assuming any particular specification. This aspect also makes the approach more general.

To analyze equilibrium behavior, we focus on pure strategy Markov perfect equilibria (MPE). In an MPE, each firm's behavior depends only on the current state and its current private shock. Formally, a Markov strategy for firm i is a function $\sigma_i : S \times \mathcal{V}_i \to A_i$. A profile of Markov strategies is a vector, $\sigma = (\sigma_1, ..., \sigma_n)$, where $\sigma : S \times \mathcal{V}_1 \times ... \times \mathcal{V}_N \to A$. Here, we simply assume that an MPE exists, noting that there could be many such equilibria.[4]

For each agent i the equilibrium generates a distribution over actions a_{it} conditional on states given by the measure of the set of ν_{it} such that action a_{it} is chosen under equilibrium strategy σ_i

$$P_i(a|\mathbf{s}_t) = \int \{\nu_{it}|\sigma_i(\mathbf{s}_t, \nu_{it}) = a\} dG_i(\nu_{it}|\mathbf{s}_t) \tag{3.1}$$

BBL shows that the full model above can be estimated in two steps. In the first step, agents' strategy

[3] See for example Arcidiacono and Miller (2011), Kasahara and Shimotsu (2009), Lazarev (2020).

[4] Doraszelski and Satterthwaite (2010) and Doraszelski and Satterthwaite (2010) provide conditions for equilibrium existence in a closely related model.

functions, σ, and the state transition probability distribution, $Q(\mathbf{s}_{t+1}|\mathbf{a}_t, \mathbf{s}_t)$, are estimated from observations on actions and states. In a second step, remaining profit function parameters are estimated.

3.2 The General Method

Our approach is much simpler than BBL in several respects. Primarily, we will not attempt to estimate the profit function parameters or any of the other dynamic parameters of the model such as entry costs, exit values, or any other investment costs parameters. Releasing ourselves from this burden has the benefit of allowing us to estimate a simpler and more general first stage.

Consider the "reduced form" equilibrium distribution of actions given states, $P_i(a_{it}|\mathbf{s}_t)$, given by (3.1). Since actions and states are observed, it is straightforward to recover these distributions from the data for every agent i. Similarly, we can also recover the transition probability distributions $Q(\mathbf{s}_{t+1}|\mathbf{a}_t, \mathbf{s}_t)$. Under the assumptions of the model, these two sets of distributions completely determine the joint distribution of all future actions and states conditional on any starting state of the world $\mathbf{s}_0$.

$$(3.2) \quad Pr(\mathbf{a}_0, (\mathbf{a}_1, \mathbf{s}_1), ..., (\mathbf{a}_t, \mathbf{s}_t)|\mathbf{s}_0) \quad = \quad P(\mathbf{a}_t|\mathbf{s}_t)Q(\mathbf{s}_t|\mathbf{a}_{t-1}, \mathbf{s}_{t-1}) \ldots P(\mathbf{a}_1|\mathbf{s}_1)Q(\mathbf{s}_1|\mathbf{a}_0, \mathbf{s}_0)P(a_0|\mathbf{s}_0)$$

How can we use these distributions to evaluate the long-run effects of a merger? Assuming that the equilibrium strategy profile is the same both before and after the merger, an assumption we discuss in detail below, a merger is simply a change in the initial state of the industry, $\mathbf{s}_0$. For example, in an industry with three symmetric firms with equal capacities, after a merger the industry has two firms, one with twice the capacity as the other. After a merger between two airlines, we replace the two merging airlines with a single larger airline whose network is the union of the networks of the two merging carriers.

Using equation (3.2), it is straightforward to determine the future distribution of industry outcomes both with and without the merger. In practice, once the first step estimates have been obtained, we use the BBL forward simulation procedure to simulate the distribution of future market outcomes both with and without the merger. These two distributions can then be directly compared. We can even compare industry structures at different times in the future: 5 years, 10 years, or whatever is the period of interest.

3.3 Relation to BBL

The approach of this paper makes much weaker assumptions on the underlying economic model than the BBL estimator does. To estimate the second stage parameters, in the first stage BBL (and similarly all two step approaches in the spirit of Hotz and Miller (1993)) must recover the actual equilibrium strategy functions, σ_i, from the dynamic game. In order to estimate them, the strategy functions must be identified, which places substantial restrictions on the underlying model. For example, identification would typically require the private shock ν_i be single dimensional. This would restrict the researcher to modeling either a single dimensional cost shock or a single dimensional demand shock, but not both. In contrast, in our model the private shocks can be multidimensional, allowing the underlying model to contain many dimensions of both cost and demand shocks (such as cost and demand shocks for each city, city-pair, airline, etc). Identification of the strategy functions would also typically require strong functional form assumptions, including that the private shocks enter the profit function additively. In contrast, our approach places no restrictions on the functional form of the structural model. The private shocks can also enter the underlying model in any way. The reason our approach is so general is that the distribution of actions conditional on states is always identified. It is simply observed directly.

Of course, there are also costs that come with generalizing the model in this fashion. One is that, under our assumptions, we can not compute counterfactuals that would occur if the equilibrium regime changes. Computing such counterfactuals would be desirable, but in the context of industry dynamics it typically comes at great cost, both in modeling assumptions and computational burden. For airline mergers, we believe these requirements would be particularly onerous. Instead, we impose the policy invariance assumption, which is not required in more structural analyses such as BBL. We now discuss this assumption in more detail.

3.4 The Policy Invariance Assumption

3.4.1 Logical Consistency

The central assumption of our approach is that the equilibrium strategy profiles remain the same both in the observed data and in the counterfactual scenario, i.e. both before and after the potential merger. In any model where the merger is part of equilibrium play this assumption would hold. We are therefore implicitly maintaining an assumption that the policy environment is constant in the past data and in the future period of interest, whether or not the merger takes place. If something about the policy environment were to change, either at the point of the merger or any other time, then equilibrium behavior might change, and the past estimates or the future simulations may be invalid. Importantly, as long as the estimated policy functions take into account network effects, hubs, or any other *observable* merger effects, our assumption remains valid.

In the context of mergers, we might particularly worry about evaluating a “game-changing” merger, i.e., one that would never have been approved under the past policy regime. If such a merger were to go through, we might expect that firms would update their beliefs about the future policy regime, and new equilibrium strategies would result. Our method will instead evaluate what would have happened in the industry had the merger taken place with the original equilibrium strategies remaining in place. The only way that we know of to fully evaluate a game changing policy change would be to compute a new MPE strategy profile under the new policy, a much more difficult undertaking than the one we propose. Certainly such an approach would be intractable in the rich airline model that we outline below.

3.4.2 Empirical Validation

The policy invariance assumption can potentially be tested in the data, either directly or indirectly. For a retrospective merger analysis, one can estimate the policy functions both before and after the merger and then formally test whether the estimated functions are statistically different. This direct approach is preferred to an alternative, indirect, approach that estimates only the pre-merger policy functions and then compares the model-predicted evolution of the industry to the actual data. Even though the indirect approach can be useful for predictive purposes, it misses the ultimate goal of the antitrust inquiry. The goal of this inquiry is not to predict what happens after the merger, but to evaluate the effect of the merger by comparing the industry dynamics with and without the merger at stake.

Of course, if a merger analysis is prospective and the merger has not happened yet, a direct test is impossible. However, this assumption can still be validated indirectly. We evaluate its validity by assessing the model's in-sample and out-of-sample fit, by testing the anonymity of the policy functions, and by testing whether the unobserved shock is iid. The data suggest that the policy assumption is likely valid for our empirical specification of the policy functions.

3.5 Estimating Welfare Effects

The procedure described above generates the joint probability distribution of actions and states (3.2) at every point in time for both the merger and no merger cases. In many cases, knowing these distributions may already be enough to shed light on the medium and long-run competitive effects of a merger. For example, in the application to airline mergers, we use the estimates of these distributions to compute measures of market concentration over time and determine when/if off-setting entry is likely to eliminate the short-run increase in market concentration caused by a merger.

A more precise estimate of the welfare implications of a merger would require, in addition, a static model of demand and supply that maps the distribution of observed states $\mathbf{s}_t$ to equilibrium prices, quantities, and consumer welfare. For example, in many markets the model of Berry, Levinsohn, and Pakes (1995) (BLP) would be appropriate for this purpose (though likely not for airlines – see below), and it is straightforward to join a BLP model with the dynamic model above, proving that such an approach would be broadly applicable.

This additionally shows that the dynamic and static approaches are complementary. The dynamic approach alone provides only the future distribution of states $\mathbf{s}_t$ (and dynamic actions). This distribution can be used to compute rudimentary measures of concentration, such as HHIs, but does not yield precise measure of welfare. The static approach alone provides only a mapping from states to welfare. The static approach provides precise measures of welfare, but requires an assumption about what happens to the distribution of states (and typically researchers and policy makers just assume that nothing happens besides the merger). We believe that, put together, the two methods can be made more powerful.

The airline industry has rich data on airline presence and a relatively simple product space, so is well suited to the dynamic analysis. On the other hand, due to dynamic pricing and price discrimination and data limitations (it is not possible to observe consumers' choice sets at the time of purchase), we believe that providing a credible static welfare analysis for airlines would be a highly complex and ambitious undertaking that deserves a separate paper.

3.6 Identification

Under the iid assumption and given that actions and states are observed, theoretical identification is straight forward. However, in practice there could be an issue in the empirical implementation of the approach if there were not enough past data to identify all of the areas of the choice distributions $P(\mathbf{a}_t|\mathbf{s}_t)$ of interest. For example, it would be difficult to estimate the dynamic effects of a merger to monopoly for an industry that had always had at least two firms in the past data. There simply would be no data that would tell us the likelihood of entry when there is a monopolist. We will see below that in our airlines example the data are sufficiently rich that this issue will not arise. Nevertheless, it is something to watch out for in other applications. A separate identification issue, that we discuss further below, is the failure of the iid assumption.

4 Application: A Model of the U.S. Airline Industry

We now outline a model of the US airline industry. In the interest of keeping the model as simple as possible, we will model only airline route presence. It would be possible, computationally tractable even, to also model the extent of entry (e.g., number of seats or flights per day) on each route, but we believe that the marginal benefit of doing so may not be worth the additional complexity. Our hope is that the current approach is both easy to understand and also provides the main insights to be gleaned from the dynamic analysis.

Consider an air transportation network connecting a finite number, K, of cities. A nonstop flight between any pair of cities is called a *segment*. We index segments by $j \in \{1, ..., J\}$ and note that $J = K*(K-1)/2$, though of course not all possible segments may be serviced at any given time.

There are a fixed number, A, of airlines. As entry of new airline carriers is very rare, it would not be possible to estimate the likelihood of new entry occurring using past data, so we will rule it out in the analysis. Each airline i has a network of segments defined by a J dimensional vector, n_i. The jth element of n_i equals one if airline i currently flies segment j and is zero otherwise. Let the $J \times A$ matrix N be the matrix obtained by setting the network variables for each airline next to each other. We call N the *route network*.

In order to travel between two cities, consumers are not required to take a nonstop flight, but might instead travel via one or more other cities along the way. Thus, we define the market for travel between two cities broadly to include any itinerary connecting the two cities. Below we will argue that itineraries involving more than one stop are rarely flown in practice, and will restrict the relevant market to include only nonstop and one-stop flights. Markets are indexed by $m \in \{1, ..., J\}$.

Airline j maximizes the total profit from all markets it serves. Profits depend on city pair characteristics, z_m, as well as the strength of competition in the market as described by the airline route network, N_t. We will not model demand in detail, but we imagine that there are likely to be unobserved profit shifters at the city pair and perhaps airline levels.

We will assume that decisions are made in discrete time at yearly intervals. Each year, t, an airline can make entry and exit decisions at the route segment level that will be reflected in the network the next year, N_{t+1}. Changing the firm's network may also involve costs. Though we will not model them explicitly, we imagine there are three potential sources of costs, in order from largest to smallest: (a) costs of opening or closing a new airline, (b) costs of opening or closing operations at a given airport, (c) costs of opening or closing operations on a given route segment (in which both endpoints already have service). Below we will find that (a) and (b) are large enough to make these events rare in practice.

Each period, each airline chooses its next period's network so as to maximize the expected discounted value of profits. Let Z_t be a vector consisting of the profit shifters z_m for all markets m in period t, and assume that Z_t is Markov. Note that the notation allows Z_t to contain aggregate variables that are relevant to all markets. A Markov perfect equilibrium in this model is characterized by a set of strategy functions of the form:

$$n_i^{t+1}(N_t, Z_t, \nu_{it}),$$

where ν_{it} represents the vector of all of the unobserved profit and cost shifters for airline i in all markets.

Assuming symmetry, these functions would have the property that permuting the order of airlines in N_t (and correctly updating the index i) would not change the value of the function. However, while symmetry is commonly assumed in many applications of dynamic games, here complete symmetry may not be a good assumption as there are at least two kinds of airlines: hub-and-spoke and point-to-point (or "low cost") carriers. This is something that we will explore empirically.

The model above results in a set of behavioral probability distributions for each airline:

$$Pr(n_i^{t+1}|N_t, Z_t) \tag{4.1}$$

that correspond to the equilibrium distribution of actions conditional on states in the general model above. If we knew the underlying primitives of the model, these probabilities could be obtained by computing an equilibrium. However, in our context computing an equilibrium is out of the question, and furthermore there are almost surely going to be many equilbria (with associated strategy functions and behavioral probability distributions). Alternatively, we will follow the general method described above and begin by attempting to recover these distributions empirically.

4.1 Abstractions

In trying to keep the model simple, we have necessarily omitted some important features of the airline industry. Most notably, in modeling the airline network and entry and exit, we have modeled presence only and have not accounted for the extent of entry (e.g., the number and size of flights). As mentioned above, there is plenty of available data so it would be possible to model the extent of entry. However, it would make the model and estimation more complex, surely beyond what would be desirable in a typical merger analysis by antitrust authorities. Additionally, it is not obvious to us that the benefit justifies the cost of such an analysis, which would primarily be a slightly more precise measure of post-merger concentration.

Finally, we will not explicitly allow for hub formation and destruction. Our set of city characteristics variables, Z_t, will include whether or not a city is a hub for a given airline, but this will be treated as exogenous and fixed. Airlines can grow and shrink their networks in each city (hubs and non-hubs), but they cannot form new hubs or dissolve old ones. While it would be relatively straightforward to relax this assumption in theory, forming new hubs or dissolving old hubs is also quite rare in the data, making it difficult to model empirically.[5]

4.2 Policy Invariance: Discussion

For the policy invariance assumption (Section 3.4) to hold, mergers should only affect the observed state (N_t, Z_t), while keeping the policy functions $Pr(n_i^{t+1}|N_t, Z_t)$ intact. The more flexible the definition of the *observed* state is, the more likely this assumption is to be satisfied.

We empirically validate this assumption in Section 6. But first we informally discuss some potential scenarios for which the policy invariance assumption might fail. First, one might worry that the scale of the newly merged airlines are "out of sample." However, entry decisions are made at the route level in our empirical model, and the incentives driving decisions are network and competition features that are local to the route and the city-pair. Therefore, while the post-merger airline may be larger than any existing airlines, the incentives faced on each route are of a similar scale to those faced by the airlines in our sample.

Second, perhaps cost efficiencies unique to the merged airline might make a new entry strategy optimal. The costs typically cited by merging airlines are either fixed costs (e.g., integrating information systems) or more efficient usage of city-specific capital (e.g., hangar space). The former are irrelevant for entry decisions, and the latter are captured by our city service and concentration measures.

Finally, de-hubbing and slot constraints might alter the incentives of the post-merger airline. As discussed in Section 7, we do see some mild de-hubbing of the post-merger airline, but the effects are not strong. Alleviating slot constraints is often cited as a pro-competitive, merger-specific efficiency, and many recent merger have been approved only after the merging airlines agreed to transfer some of their slots to their competitors. It is outside the scope of this paper to determine to what extent these divestitures helped. However, the fact that only 3 out of 60 cities in our data had a slot constrained airport makes us believe that the impact of slot constraints is unlikely to be a first-order issue in our analysis.

5 Data

The principle data source is the Bureau of Transportation Statistics (BTS) T-100 Domestic Segment Data set for the years 2003-2008. More historical data is readily available. However, due to the large impact of the events of 9/11 on the airline industry, we view 2001 and 2002 as not representative of the current industry, so we dropped those from our sample. We did not use data from years prior either because our model requires

[5] The only hubbing or dehubbing event in the period covered by our data is Delta dissolving their Dallas-Fort Worth hub in 2005.

us to use a period where airlines' entry/exit strategy functions are relatively stationary, and we felt that this was not likely to be true over longer time horizons due to changes in policy, technology, etc. However, we note that we have tried extending all of our estimations back all the way to 1993 and achieved very similar results.

The T-100 segment data set presents quarterly data on enplaned passengers for each segment flown by each airline in the U.S. The data defines a segment to be an airport to airport flight by an airline. A one-stop passenger ticket would therefore involve two flight segments. We use data for the segments connecting the 75 largest airports, where size is defined by enplaned passenger traffic. The data was then aggregated to the Composite Statistical Area (CSA) where possible and to the metropolitan statistical area when this was not possible. The end result was segment data connecting 60 demographic areas (CSAs). Note that this means we are treating multiple airports at the same city as one. We feel that this is more appropriate for our purposes than treating them as separate destinations. Appendix A contains the list of airports included in each demographic area.

Although the airline strategy function is defined over the route segment entry decisions, we also allow airlines to carry passengers between a pair of CSAs using one-stop itineraries. The combination of non-stop and one-stop service between two CSAs is denoted the "market" between the CSAs. Using the data on itineraries actually travelled as a guide (DB1B), we define an airline as present in a market if either (1) the airline provides service on the route segment connecting the two CSAs OR (2) the airline provides service on two route segments that connect the CSAs and the flight distance of the two segments is less than or equal to 1.6 times the geodesic distance between the CSAs. Itineraries that use two or more stops are extremely rare in the airline ticket database so we exclude this possibility entirely. Note that in certain places we supplement the T100S data with data from the T100M "market" database, the DB1B ticket database, and the Household Transportation Survey (tourism data).

There are many flights that show up in our data as flown by regional carriers (e.g., Mesa Air) that are in fact flown under contract with a major carrier. On these flights, the major carrier sells the tickets and, typically, the plane would have the major carrier's name on the outside and would generally appear to passengers to be owned by the major carrier (though in many cases it is not). Major carriers can contract with different regional airlines in different parts of the country and contracts change over time in terms of what routes are covered. Regional carriers may also fly some routes under their own name, selling tickets themselves. Flights flown by regional carriers represent about 25-30% of the flights in the major carrier's networks in our data (see Appendix A.3), so ignoring them could potentially be very problematic. In our analysis, we attribute flights flown by regional carriers to the major carrier that they are contracted to. That is, if Mesa flies a plane under contract for Delta, we will call that a Delta flight for the purposes of the analysis and treat it identically to a flight that Delta flies itself.

The T100 data set we use to describe the route networks of the airlines contains numerous flights that are not regularly scheduled, such as charter flights, and even flights diverted due to weather or equipment problems. As a result, if we were to define airline market presence by the existence of a small number of flights on a given market, we would pick up a very large number of phantom entries that did not represent regularly scheduled service. Our goal is to describe stable features of the airline networks rather than idiosyncratic flights flown. We therefore define an airline as having "entered" a segment if at least 9000 passengers are carried on a segment, roughly coinciding with a single daily nonstop flight, in each of four consecutive quarters. Symmetrically, an airline has "exited" a segment if it has not carried 9000 passengers on a segment in each of four consecutive quarters. Our entry definition is explained more thoroughly in Appendix A.

5.1 Data Summary

Table 1 lists summary statistics for segment and market presence by airline. Southwest has the most nonstop routes, followed by the three major carriers: American, United, and Delta. Because the majors have hub-and-spoke networks, as compared with Southwest's point-to-point network, they are present in as many or more one-stop markets as Southwest despite flying fewer nonstop routes. A striking feature of the data is the rapid expansion of Southwest and JetBlue. The other major airlines are growing much more slowly.[6] On average airlines enter and exit between five and ten percent of their routes per year.

Table 2 lists summary statistics for the airline's networks, concentrating on the variables that we will use in the estimations. An observation in the data is an airline-year-city pair; there are ten airlines (not counting America West before it was merged into US Airways) and 1770 city pairs.

5.1.1 City Pair Characteristics

In the past literature, the most commonly used measure of the underlying demand for air travel between two cities is the interaction of the populations of the cities. This population variable is intended to measure the total possible number of visits between residents of the two cities, but it also has the disadvantage that it ignores many other important features of demand such as city proximity, availability of alternative methods of transport, industrial activity, etc. We instead use the variable "Log(2002 Passenger Density)," which measures the log actual passenger density (enplanements) for each market in the year 2002. The density

[6] Growth in US Airways' network is largely due to the merger with America West.

variable helps capture many of the unobservable aspects of market demand that are peculiar to a given city pair. Boguslaski, Ito, and Lee (2004) have shown that passenger density does a very good job in predicting Southwest's entry behavior. Note that in cases such as unserved markets, where the density variable equals zero (over 25% of cases – see Table 2), we set Log(Density) equal to zero. A potential problem with using the density variable is that, because density depends somewhat on the airline networks, it would be endogenous. To mitigate this issue, rather than measuring density lagged one period, which would be valid under the iid assumption but invalid otherwise, we measure density in the period just prior to our estimation sample. As a robustness check we have also tried using density lagged one period, with similar results.

To capture underlying demand in unserved markets, where passenger density is zero, we also include the product of the population at the route's endpoint cities, interacted with a dummy for whether the route is unserved.

We also construct a second density measure that we call "Log(Passenger Density in New Markets)" that reflects a particular route segment's importance in each airline's overall network of markets (nonstop and one-stop flights). Specifically, this variable equals the log difference in total passenger density on the network (in 2002) on the nonstop and one-stop markets served with and without the route segment under consideration. It is meant to capture total potential revenue gain/loss across the entire network from adding/subtracting each route segment individually. This variable was inspired by anecdotal evidence suggesting that American Airlines uses a similar measure in making its entry decisions.[7] Note that this variable is zero more than 50% of the time, reflecting both the presence of unserved markets as above, and also the fact that some routes in an airline's network are extraneous, in the sense that they do not add any new markets to the network but merely duplicate existing service in a more convenient way.

A fourth demand variable, "percent tourist," measures the percentage of passengers travelling in each market who reported that their travel was for the purpose of tourism in the Household Transportation Survey. We found that other city characteristics such as household income had no explanatory power in our data so we excluded them from the analysis.

5.1.2 Competition Variables

In our estimations we use a large number of variables that attempt to characterize competition on each route segment. First we divide competitors into non-stop and one-stop to help pick up consumers' preference for non-stop travel, as well as any cost considerations. The average city-pair has slightly less than one non-stop competitor and 3.5 one-stop competitors. Of course both of these variables have very skewed distributions

[7]This anecdote has been relayed by Steve Berry in several talks but not, to our knowledge, in print.

with many zeros and a few city-pairs that have many carriers. We also measure the number of code-share agreements that each airline has on each route segment.[8] Code shares are fairly rare.

We have also computed a large number of concentration measures for each market. The variable "HHI Among Others (Market)" directly measures the concentration among rival carriers on the city pair in question, including both non-stop and one-stop competitors. The HHI among competitors averages about 5000 in our sample (where HHI ranges from 0 to 10,000).

There is also substantial evidence (Borenstein (1989), Borenstein (1990), Borenstein (1991), Berry (1992)) that the size of a carrier's operations at the endpoint cities influences a carrier's market power on travel between those cities independently of concentration on the market itself. Thus, we also include variables measuring a carrier's market share at each endpoint city ("Own Share (City) Large/Small"). The use of "Large" and "Small" refer to the largest and smallest value out of the city-pair connected by the route segment. For similar reasons we also include measures of concentration at each endpoint city ("HHI Among Others (City) Large/Small"). Note that these variables might also influence entry for cost reasons.

If we measured the market share and HHI variables in the natural way, using the number of enplaned passengers, then it would not be possible to simulate future values of the competition variables under a merger without also estimating a demand system that predicted enplaned passengers at future dates. Thus, we instead measure all of the HHI variables using the number of routes served out of each city. It turns out that this yields essentially identical estimates empirically.

Our final measure of competition is whether or not a competitor has a hub on the route. Own hubs are treated separately below.

5.1.3 Network Characteristics

For each city-pair route segment we also have a large number of measures of local network characteristics. We measure segment (non-stop) presence and market (feasible one-stop) presence separately, as well as endpoint presence ("Present at Both Airports (not Market)"). These variables are non-nested in the sense that an airline can either serve a route segment or be "Present at Both Airports (not Market)", but not both. All of these should have large effects on market presence through the cost side.

We measure how many endpoint cities are a hub for each airline. We also measure how convenient the most convenient hub is to the route segment by taking the non-stop distance and dividing by the one-stop distance for the closest hub. If a hub is very convenient, nearly on a straight line between the two cities, one might expect that the airline could very easily serve the route via one-stop travel. We also measure

[8] This variable is compiled from data that is separately measured for each airline pair-route segment using the ticket data.

the distance to the nearest hub for each end, ranked (Large/Small), which is meant to be a measure of how central to the network the two endpoints are.

Finally, we measure the size of each airlines' network at the endpoint cities using the number of non-stop destinations served at each endpoint city, ranked (Large/Small). This variable could influence market presence through both the demand and the supply sides. Note that it is different than the share variables above because it measures network size rather than network share.

5.2 A Simple Static Merger Analysis

With the notable exception of Southwest, U.S. airlines have hub-and-spoke networks. As a consequence, out of 1770 possible city pairs in our data, the typical major airline flies only 150-220 nonstop routes (in 2008), while still covering 1100-1500 city pairs with reasonable one-stop connections. For comparison, Southwest flies 323 nonstop routes and covers only 1042 markets. (See the first column of both panels of Table 3.) In this section, we perform an analysis of the immediate effects of mergers on market concentration, which we contrast with the medium- and long-run effects in later sections.

Table 3 summarizes the level of competition faced by each airline across its nonstop routes (panel A) and feasible one-stop markets (panel B). Southwest, Delta and Northwest are the airlines most isolated from competition in the sense that they have the most monopoly and duopoly nonstop routes. As a result, the overnight effect of a DL-NW merger is to create an airline with 108 monopoly nonstops, far and away the most of any airline. The story is slightly less stark when we include feasible one-stops. However, even then, DL-NW has 31 monopoly one-stop markets and an additional 97 duopoly one-stop markets. The other two potential mergers we consider, UA-US, and UA-CO, create airlines with only about 35 monopoly nonstops, and only 13 monopoly one-stop markets.

Table 3 also shows the extent to which the data cover many different market structures. Almost every airline flies multiple routes with all combinations of 0 to 6 nonstop competitors and 0 to 10 one stop competitors. We also observe many different pairs and triplets of airlines across markets. This rich variation in competition is what allows us to empirically identify the relationship between competition and route entry.

To measure the short-run anticompetitive effects of each merger, we compute HHIs before and after the merger, in terms of 2008 passengers enplaned.

Table 4 shows the ten worst affected nonstop routes (city-pairs) for the three mergers in terms of increase in the HHI. For DL-NW, there are seven nonstop routes – out of Cincinnati, Minneapolis, Salt Lake City, and Memphis, all hubs – where the merger essentially creates a monopoly carrier. Three other routes between Atlanta (also a hub) and these cities move essentially to duopoly. The HHI changes range from 1500-5000 points. All of these markets violate the merger guidelines by a very large margin. Absent offsetting entry, prices would be expected to rise substantially on these routes after the merger. For UA-US, we see almost the same pattern, with six monopoly routes created, mostly out of DC, Philadelphia, and Charlotte. For UA-CO, the story is not quite as bad – there is only one monopoly route created – but the patterns still would violate the merger guidelines by a large margin.

There is also substantial evidence (Borenstein (1989), Berry (1992)) that, due to frequent flyer programs, market concentration out of a city as a whole is also an important determinant of market power. Table 5 shows the five worst affected cities in terms of HHI increase across all flights from the city. Again, we see large HHI increases in markets that were already very concentrated, clearly in violation of the merger guidelines. The merger that looks worst by this measure is again DL-NW, while UA-CO appears the least bad. For UA-US, the worst case cities are Charlotte, Philadelphia, and DC. Concentration at these cities was cited as the main reason that the merger was blocked.

The HHI results provide a short-run snapshot of the increase in concentration that would result from the three proposed mergers. By these short-run measures, all the mergers look pretty bad, in the sense of increasing concentration and leading to upward price pressure. Of course, arguments can still be made in favor of the mergers. Large cost savings could offset the harm from decreased competition. However, cost savings would have to be large and system wide to justify the increases in concentration seen in many markets. We will not explore this avenue in this paper. Alternatively, it may be that entry costs are low and that the cities discussed above are likely to experience offsetting entry in a short period of time. Below, we use our dynamic model to explore this possibility by simulating medium and longer term market outcomes.

6 Empirical Implementation

6.1 Overview

Our goal is to find an empirical specification for the policy functions that would be consistent with the policy invariance assumption. The primary difficulty with estimating the airline model is that, in their general form, the choice probabilities in (4.1) are high dimensional and would be identified only by variation in the data over time. Variation across airlines could also be used if we were to assume some symmetry across carriers. However, given that there are at least, in principle, two types of carriers, hub-and-spoke carriers and low cost carriers, we do not necessarily want to impose symmetry across all carriers. At the very least we should explore this empirically. Furthermore, given that we have only ten carriers and six years of data, that still only leaves 60 observations to determine a very high dimensional set of probabilities.

Therefore, to estimate these probabilities we will need to use at least some of the variation in the data within an airline's network (across city pairs) to identify the strategy functions. In principle, all segments in the whole system are chosen jointly, and we would like the model to reflect that. However, it also seems unlikely that the entry decisions are very closely related for segments that are geographically distant and not connected in the network. Thus, our empirical approach will be to start with a fairly simple model, and then add complexity until we exhaust the information in the data. As a robustness check, we will test for any remaining complementarity of entry decisions.

The simplest model we can think of would allow the entry decisions across segments to be correlated only through observable features of the market, so this will be our base model. In the base model, we assume that there are only nonstop segment level shocks and that these shocks are independent across nonstop segments. We model segment presence (entry/stay in=1, exit/stay out=0) using a probit model.

Tables 6 and 7 show the baseline probit estimates for route presence using data pooled for all airlines. Table 6 includes year dummies, city dummies, and all of the route demand variables, while Table 7 adds route fixed effects and drops all variables that have no variation at the route level. To help interpret the coefficient magnitudes, the third column of each table reports the marginal effects of the estimated coefficients for an airline that is indifferent (in expectation) between entering/staying in and exiting/staying out, while the fourth column reports the marginal effects of a one SD change in each variable.

We first show that the data reject the specification in Table 6. We then show that specifications more flexible than the one in Table 7 fail to deliver significant improvements. We conclude that the data do not invalidate the policy invariance assumption for the policy functions presented in Table 7, our base specification.

6.2 Serially Correlated Unobserved Shocks

Before discussing the estimates, we first consider the important issue of serially correlated unobserved shocks that can invalidate the two-step approach. Recall that the primary effect of serially correlated market level demand or cost shocks would be to bias the competition coefficients upward (more positive/less negative). This is because a higher value of the shock would also lead to persistently more entry and less exit, and make it appear as if competition was less unfavorable than it really was. The first specification (Table 6) controls for unobserved shocks using city fixed effects and detailed demand variables. The idea is to make demand shocks as observable as possible, leaving little unexplained. The second specification uses more detailed route fixed effects (Table 7).

While we found that the demand variables did help remove some of the bias in the competition coefficients, it is clear from the estimates in Table 6 that this strategy has not entirely worked. The coefficient on nonstop competitors is very small (-0.14). Most importantly, it is 23 times smaller than the coefficient on route presence (3.25), suggesting that it would take 23 additional nonstop competitors to offset the effect of being present on a route. To add another point of comparison, the estimated coefficient on nonstop competitors is of similar magnitude to the coefficient for adding a single connecting flight out of the largest endpoint city on the route. The estimated competition effects in this model seem implausibly small.

Moving to the estimates from the model with route fixed effects in Table 7, we can see that the estimated competition coefficients are now an order of magnitude larger. According to these estimates, it takes only two competitors (-2.93) to fully offset the effect of route presence (2.48). This magnitude seems much more plausible.

To further investigate this issue, we also implement a statistical test. Lazarev (2020) suggests testing the Markov property of the observed distribution of actions given states against a general alternative. In our setting the enormous size of the state space would make such a test impractical. Furthermore, with such a large state space, that test is likely to have very low power, and thus failure to reject it would not be very convincing.

Instead, to achieve a balance between power and generality of the alternative, to test for a violation of the Markov property we re-estimated the same models including the variable "presence two periods ago." Under the null hypothesis that the observed distribution of actions given states is Markov, the coefficient on this variable should be zero, i.e., presence two periods ago should not have any remaining explanatory power in the probit model. On the other hand, if there is a serially correlated demand or cost shock, then presence two periods ago would be a function of this shock and its coefficient should be nonzero.

For the model with the demand variables and city fixed effects only, we strongly reject the null hypothesis (p=0.004). This finding supports our conclusion that the magnitude of the competition effects in the model with city fixed effects is implausibly low. For the model with route fixed effects, we fail to reject the null (p=0.883). Moreover, the coefficient on "presence two periods ago" is estimated to be very close to zero (0.015), suggesting that the impact of any remaining unobserved heterogeneity on the coefficient estimates is likely to be minimal. We therefore proceed by using the specification in Table 7 as our base model.[9]

[9]This empirical specification of policy functions can be derived from a model with fixed route-specific profit shifters that take additive form.

6.3 Base Model Estimates

We now discuss the base model estimates in Table 7. First, consider the competition variables. The effect of nonstop competitors is estimated to be strong and nonlinear. The marginal effect of going from zero to one nonstop competitor (-1.93) is strongly negative, almost as large in magnitude as the effect of market presence (2.48). These are the two largest coefficients. Adding a second, third, or fourth competitor also has strong negative effects, starting about half the size of the first competitor and trending smaller. Adding competitors above four was not estimated to have a further effect, perhaps in part because there are a limited number of routes with more than four nonstop competitors so these effects are estimated with less precision. The effect of one-stop competition was consistently found to be small and insignificant so that variable is omitted from our base model.

We also found that the market structure at the two endpoint cities plays a significant role, particularly the larger endpoint city. For an airline indifferent to market presence, a ten percent increase in market share at the larger endpoint city leads to a 11.7% increase in the probability of entry/presence. For the smaller endpoint city this effect is 7.2%. Similarly, the more concentrated competition is at the larger endpoint city, the more likely that an airline enters the route. Finally, if a competitor has hub at one endpoint, entry/presence is significantly less likely—the marginal effect for an indifferent airline is -23%. These findings are consistent with prior literature that finds that market power at the endpoint cities increases market power on all the routes out of those cities. Note that concentration at the smaller endpoint is estimated to have a negative effect, though that coefficient is insignificant.

The other important factor determining entry/presence is the thickness of the airline's route network in the vicinity of the endpoint cities. As mentioned above, prior route presence is the single most important variable. This represents the "stickiness" in airline entry decisions and is likely induced by the effect of any sunk costs of entry such as the costs of setting up operations at an airport and of advertising the new route. Other than prior route presence, the most important network variables are the number of nonstop cities served at each endpoint. Adding a single nonstop route out of the larger endpoint city increases the probability of entry/presence by 11%, while adding a single nonstop route out of the smaller endpoint increases the probability of entry/presence by 9%. Note that these variables vary widely in the data, so this means that they have a large amount of explanatory power in the model. They are therefore among the most important variables in the model (along with the competition variables) because they play a large role in determining what airlines are potential entrants for a particular route.

We also found that endpoint hubs have a large significant positive effect on entry/presence, and hub convenience has a negative effect, presumably because the route is easily served by stopping at the hub instead. Distance to nearest hub does not add much once the units are accounted for.

6.4 Model Fit

To evaluate in sample fit, Table 8 computes psuedo R-squared statistics for different subsets of the data. The first two columns of the table compute a psuedo R-squared by airline, for routes where the airlines stays in, and routes where it stays out. These are, of course, the easiest routes to predict, as airline route presence is static for most routes from one year to the next. Looking first at the stayers, we can see that routes where an airline stays out have psuedo R-squareds that are typically greater than 0.99, so the model is matching these routes almost perfectly. Routes where airlines stay in are also matched very closely, with the possible exception of JetBlue, which is a young, rapidly expanding airline in our data period.

Columns three and four compute psuedo R-squareds only for routes with new entry and exit. These psuedo R-squareds are much lower, averaging around 0.2, with exit predicted slightly better than entry. This is a very strict test of the model as it asks whether the model is capable of predicting exactly which new routes each airline enters/exits each year. We therefore view these psuedo R-squareds as quite high, which adds support to our policy invariance assumption.

In order to see more clearly which aspects of the data the model is fitting well, the final two columns of the table show the fit of the model in 2008, when it is simulated for the entire data period 2003-2008 using 2002 as a starting point (and never updated using the actual outcomes). Again we consider only routes with new entry and exit over that period. These psuedo R-squareds are extremely high, averaging around 0.65, with exit fitting better than entry.

Overall, we conclude that the model fits new entry and new exit very well by airline-route, but is less good at predicting the exact year of new entry/exit. In other words, the model predicts the marginal routes very well but is not as good at predicting the timing of new entry/exit on these routes. The fact that the model does not capture timing well is perhaps not surprising as our data do not have good measures of year to year demand or cost variation at the route level. (The main source of yearly demand/cost variation is the year dummies.)

6.5 Validating Policy Invariance

In general, there always exists a model flexible enough for which the policy invariance assumption holds. One just needs to redefine the state space appropriately. For example, the state space can be augmented to account for the potential endogeneity of merger decisions and/or of the merger enforcement policy. Of course, without appropriate variation in the data, a more flexible model will not be identified, which creates a trade-off for the researcher. In this section we describe several attempts at generalizing the base model. These attempts all failed. A recurring theme is that the increase in flexibility improves the in-sample fit of the model, but fail tests of overfitting using cross validation.

6.5.1 Asymmetric Strategy Functions

One generalization we explored was to allow the probit entry functions to differ across airlines. The most obvious approach here is to group carriers into two groups: traditional "hub-based" carriers and point-to-point "low cost" carriers. We also explored estimating different strategy functions for each airline. Note that such an approach is somewhat complicated by the fact that the airlines' networks are not entirely overlapping.

In each case, we found similar results: allowing for separate probit entry functions increased the in-sample fit. However, the coefficient estimates were in general less precise, and more frequently showed unrealistic signs or magnitudes than the base model estimates above, giving us less confidence in them. Moreover, cross validation rejected the asymmetric strategy functions in favor of the symmetric base model.

In addition, we explored adding airline dummies to the base model. For the model without route fixed effects, the only airline dummy that showed up as mattering at all in terms of magnitude and statistical significance was JetBlue, which is included in the model in table 6. For the model with route fixed effects, all of the airline dummies were insignificant and close to zero, so they were omitted.

6.5.2 Nonparametric Estimation

Since the primary use for the base model is to predict airline entry behavior with and without mergers, it seems potentially fruitful to consider machine learning approaches developed for prediction models. Moreover, in our context, since the functional form of the strategy functions is unknown, BBL suggests estimating them nonparametrically. In this section we explore these possibilities by reestimating the base model using an artificial neural network (ANN) of varying dimensions.

The ANN has the ability to match arbitrary nonlinearities, including interactions between variables. The implementation of the ANN used here had a linear input layer and logit middle layers, and was estimated using maximum likelihood with a normal error. This makes it a nonlinear probit and a direct generalization of the linear probit model above. The ANN is computationally intensive to estimate, so it was estimated on the version of the base model above without fixed effects. The ANN was estimated with 10-fold cross validation.

Table 9 lists the value of the cross-validation likelihood for the linear probit, as well as the ANN with successively higher dimension. As we make the ANN richer, the in-sample fit improves substantially (not reported), but the cross-validation results show that this improvement is due to overfitting. Cross validation picks the linear probit as the preferred model, rejecting all of the nonparametric models.

We believe this result is primarily caused by the fact that the linear probit model is already explaining much of the variation in the data (see discussion of table 8 above), leaving little room for nonlinearities to improve the fit of the model. Additionally, many of the explanatory variables in the model take on only a small number of values, so nonlinearities are not important for these variables.

6.5.3 LASSO and Fixed Effect Selection

An important issue with prediction models is overfitting from including too many variables in the model, which produces unbiased but noisy predictions. We found above that the route fixed effects were effective at controlling for unobserved serially correlated shocks, but the drawback of having so many (1770) of them is that there is the potential for many of them to be poorly estimated, resulting in overfitting.

To address this issue, we attempted to use LASSO to reduce the set of fixed effects and reduce out of sample prediction error, as measured by cross validation likelihood. Unfortunately we were unsuccessful in this endeavor. When LASSO was applied to the full base model, it did significantly improve the cross validation likelihood. However, the resulting model estimates were similar to those above for the base model without fixed effects: the competition coefficients were unrealistically small and the estimated model failed the test for serially correlated unobserved errors. We also made a second attempt at using LASSO, where we only applied LASSO to the fixed effects coefficients, holding the rest of the coefficients fixed at their values estimated in the base model with all the fixed effects included. Perhaps not surprisingly, this process yielded best estimates that included almost all of the fixed effects. It did slightly reduce the cross validation likelihood, but this improvement was achieved not by reducing the number of fixed effects, but by shrinking all of the fixed effect estimates toward zero. We therefore proceed below with the base model including all the fixed effects.

6.5.4 City Specific Shocks

The final generalization we explored was to generalize the variance structure of the probit model. The idea here is to try to capture the fact that there might be city specific cost and demand shocks that affect entry decisions for all routes out of a given city. Anecdotally, it often seems that airlines expand and contract their networks on a city wide basis rather than route by route, and city-specific shocks may help explain this behavior. The base model assumes that shocks are independent across routes.

Let n_{ij}^{t+1} indicate presence for airline i on route j in period $t+1$, and let $x_{ij,t}$ represent all of the explanatory variables in the base model above. Then in this new version of the model, airline i is present on route j, ($n_{ij}^{t+1} = 1$), if

$$x'_{ij,t}\beta + \gamma n_{ij}^{t} + \xi_{j_1,t+1} + \xi_{j_2,t+1} + \epsilon_{ij,t+1} > 0 \tag{6.1}$$

where $\xi_{j_1,t+1}$ and $\xi_{j_2,t+1}$ are drawn from a $N(0,\tau^2)$ and are city specific shocks for cities j_1 and j_2, the endpoint cities for route j; $\epsilon_{ij,t+1}$ are i.i.d. market shocks drawn from a $N(0,\sigma^2)$; and γ is an entry threshold. Note that this formulation simply adds $\xi_{j_1,t+1}$ and $\xi_{j_2,t+1}$ to the base probit model above.

The city specific shocks generate a correlation structure in the route presence variables such that the shocks for the route between CSA k and l and the route between CSA m and n have a normal distribution with means given by the left hand side of (6.1) and variance matrix Σ where

$$\mathbf{\Sigma}_{kl,mn} = \begin{cases} 2\tau^2 + \sigma^2, \text{ if } k = m \text{ and } l = n, \\ \tau^2, \text{ if } k = m \text{ or } l = n \text{ but not both}, \\ 0, \text{ otherwise.} \end{cases} .$$

Estimating this model is similar to estimating a 1770-dimensional probit model. It is quite computationally intensive, and was only made feasible using a Gibbs sampling routine that takes full advantage of the special structure of the covariance matrix.[10] To reduce computational burden, it was again necessary to use the base model without fixed effects for this exercise.

We found no support for this model in the data. In the Gibbs sampler, the estimate of τ, the variance of the city shocks, collapses to zero. The Gibbs sampling distributions become degenerate when this happens, which invalidates the approach, so we do not report the estimates here.

7 Results: Simulating the Long-Run Effects of Airline Mergers

Tables 10-14 show simulation results for the hub/low cost pooled model above over the 10 years following our data set (i.e., 2009 – 2018). We run four simulations: no mergers, DL-NW only, UA-US only, and UA-CO only.

The estimations include year dummies that absorb aggregate shocks to the industry. For the simulations,

[10] Details of the procedure are available from the authors upon request.

we have to choose forecast values for these shocks. Since we are not so much interested in forecasting aggregate demand, and are instead mainly interested in the differences between the no merger and merger cases, how we set them does not seem too important to the results as long as they are set to be reasonably stable. In order to mimic what a pre-merger antitrust analysis might entail, in the simulations below we chose to set them such that the total number of routes served by all airlines was roughly constant over the simulation period for the base (no merger) case.

7.1 Effects on Overall Airline Networks

We first assess the effect of the mergers on the total (national) size of the airlines' networks of routes and markets. Looking at Tables 10 and 11, we see that the model predicts that the DL-NW and UA-CO mergers reduce the total number of routes served by competitor airlines (relative to no merger), while the denied UA-US merger modestly increases the number of routes served by competitors. That is, at the national level there is more offsetting entry in the latter merger. The effect on competitor low cost airlines is generally larger than the effect on competitor hub-and-spoke carriers.

Looking more closely, the DL-NW merger causes a modest decrease in the number of routes served by American, United, Continental, and JetBlue ten years after the merger. The largest negative effect of the merger is imposed on Southwest, whose network is 10% smaller ten years after the merger than it would have been without the merger. This effect is reflected in slower network growth rather than network shrinkage. Surprisingly, US Airways's network is 10% larger ten years after the merger, although this is the result of a slower rate of shrinkage. The net effect is that competing airlines serve 33 fewer routes in total if the merger occurs. By this aggregate measure, the dynamic effect of the merger is actually worse than the static one.

The UA-CO merger has similar effects to the DL-NW merger. American serves modestly fewer routes under the merger, but Southwest is greatly affected with a 10% smaller network as a result of the merger. US Airways, which was predicted to shrink by almost 10% in the absence of a merger, instead retains the same route network size. Delta, Northwest, JetBlue, and Alaska are almost completely unaffected at this aggregate level. The net effect is that competitors serve 4 fewer routes if the merger occurs.

The UA-US merger has almost no effect on the total number of routes served by the hub-and-spoke carriers. The primary effects are that Southwest serves 18 fewer routes than would be the case without the merger, but this is offset by a 21 route increase in the size of JetBlue's network. The net effect is that competitor airlines serve 7 more routes if the merger occurs.

For most of the airlines, qualitatively similar effects are seen in the one-stop markets served by the airlines (right hand panel of Table 11). But, there are some exceptions. For example, American and United see their network of markets grow slightly despite serving fewer nonstop routes. In contrast the network of US Airways markets grows by less than 2% despite experiencing a 10% growth in its nonstop route network.

The merger simulations also make predictions about the time required for these network changes to occur (Table 11). The merged airline aggressively enters routes in the year following the merger because the radically changed network of the merged airline provides strong incentives for expansion in certain markets. Meanwhile, for the airlines that are not part of the merger, the changes in the network occur more gradually over time.

7.2 Effects on Competition within cities (CSAs)

Next we assess the effects of the three mergers on the worst affected cities (Table 12). The main question we are interested in is to what extent the large static anti-competitive effects shown earlier (Table 5) are relieved by entry over the ten year simulation horizon.[11]

We find that the dynamic effects on airline competition at the city level are different for each of the three mergers. According to the simulations, in the case of DL-NW, there is essentially no offsetting entry in the worst affected cities. In fact, there are some markets (MEM, BDL) where we see the opposite. In these cities the merger stifles entry, typically by Southwest, and also leads to increased entry by the merged carrier, increasing its overall market share. These effects strongly suggest that prices would rise in these markets, at least in the absence of very large cost synergies. However, we also note that there is a small potential offsetting effect that consumers may also benefit from the convenience from being able to fly one carrier to more destinations.

In the other two mergers, the simulations do generate offsetting entry in the worse affected cities. As a result, in both cases the concentrating effects of the merger are reduced over time. For UA-US, most of this offsetting entry typically comes from JetBlue. For UA-CO, there is also offsetting entry by Southwest. For UA-CO, after 10 years new entry has offset a large fraction of the initial rise in concentration.

7.3 Effects on Competition on Routes

We now look at the simulated effects of the three mergers on individual routes (Tables 13 and 14). Starting again with DL-NW, we see that the immediate effect of the merger is to move six routes and three markets from duopoly to monopoly. After ten years, the simulations suggest that this number will increase such that there will be 14 more monopoly routes than if there were no merger. The number of monopoly routes

[11] Here we compute HHIs by # of routes served rather than passengers enplaned, since that is what we model. Computing predicted HHI by passengers enplaned would require a model of passenger demand.

increases over time for two reasons. First, there is little scope for offsetting entry on the monopolized routes. On these existing routes, the merger seems likely to cause permanent price increases. However, in addition, the merged carrier enters new routes that were previously unserved, and these routes become monopolies. The reason DL-NW enters these markets is that it has increased levels of service at the endpoints, which causes these markets to become economically viable. (Our model cannot distinguish whether this is for demand or cost reasons.) This new entry effect seems very likely to be welfare enhancing, potentially offsetting some of the harm caused by the merger from higher prices.

The UA-US merger has the greatest immediate effect on individual routes, with nine routes and three markets moving from duopoly to monopoly. However, in this merger there is substantial offsetting entry. This merger also leads to substantial new entry on previously unserved routes. Thus, at the end of the ten year simulation period, there are twelve routes with new service, and yet a net of only three additional monopoly routes. In fact, in the simulations, after ten years, competition at the route level does not seem to be substantially reduced from what it would have been without the merger, while there is more service. It seems plausible that this merger increased welfare after ten years. Ironically, it was the only merger of the three to be blocked.

Of the three mergers, the UA-CO merger has the smallest immediate impact on competition at the route level, with only one additional monopoly route and actually two fewer monopoly markets (because the merger immediately creates viable one-stop service where it did not previously exist). After ten years there is a net effect of five additional monopoly routes, but there is new service on nine routes. Thus, after ten years this merger also seems to have fairly benign overall effects at the route level.

8 Conclusions

We draw two sets of conclusions from this research. The first is that our method provides a simple yet effective way to provide some empirical insight and rigor to questions of how a particular merger will affect the evolution of an industry over time. While we have applied the method to airlines, it could equally well be applied to many industries, so long as there is rich enough past data available. A great advantage of the method is that it requires minimal econometric tools and computational power.

Our analysis of the two completed (DL-NW and UA-CO) mergers and one contested (UA-US) merger suggest that the contested merger would have had less of an anticompetitive effect than the two approved ones. While this is somewhat ironic, it is rather easy to explain. From a static perspective, the high degree of overlap in the UA-US merger makes it appear that on many routes two direct competitors will be merged.

However, from a dynamic perspective, the network overlap for United and US Airways was primarily in areas where JetBlue and Southwest were expanding, and the decrease in competition would increase the incentives for these low-cost carriers to enter the less competitive routes dominated by the merged firm.

On the other hand, the two approved mergers reduced the incentives for the low-cost carriers to expand, which helped perpetuate the anticompetitive effects of the mergers. The larger hub-and-spoke airlines appear to have an at best weak incentive to enter routes even after a merger reduces competition, and part of the blame for this lack of potential entrants is the strong influence of hubs on the network expansion strategies of these airlines.

Finally, an interesting and unforeseen finding in our analysis is the scope for new service on marginal routes. The effect of having a richer network in the merged carrier both increases the potential demand for each additional nonstop route, and also potentially decreases the cost of serving such routes. According to our model, some combination of these effects can lead to nontrivial new service as a direct result of the merger, potentially offsetting some of the harm from higher prices on routes with reduced competition.

Our methodology is meant to assess whether or not the medium- and long-run evolution of the industry will be substantially similar with and without the proposed merger, not make precise predictions about future outcomes. Importantly, the antitrust case law has recognized this nuance explicitly. In *Brown Shoe,* the Supreme Court noted: "Congress used the words 'may be substantially to lessen competition' to indicate that its concern was with probabilities, not certainties."[12] Comparing our simulations with the realized outcome is complicated by a range of factors including the impact of the great recession, the three airline mergers that occurred in the 2008-2014 time-frame, and the choice by some airlines to effectively de-hub from some cities following the respective mergers. We can use our model to simulate the outcome of the successive mergers, and we do not find general trends in the differences between our simulations and reality.[13] For the five worst affected cities of the 2008 DL-NW and 2010 UA-CO mergers (see Table 11), we find that our predictions of whether off-setting entry ameliorated the static increase in market concentration by 2013 are correct in 7 of the 10 cases despite the relatively short simulation horizon.[14]

[12]Brown Shoe v. United States, 370 U.S. 294, 321-22 (1962). See also, US Department of Justice (2010): "Most merger analysis is necessarily predictive, requiring an assessment of what will likely happen if a merger proceeds as compared to what will likely happen if it does not."

[13]In order to help make the simulations match the large economic shocks from the financial crisis, we set aggregate year dummies so that they give the best match between the simulations and the actual data, but these dummies do not help in matching city- or airline-specific shocks due to the financial crisis.

[14]Recall that our analysis suggests the model is effective at identifying marginal routes that are targets for entry, but the model is less accurate at predicting precise years of entry than longer-term entry patterns.

References

Ackerberg, D., C. L. Benkard, S. Berry, and A. Pakes (2007). Econometric tools for analyzing market outcomes. Handbook of econometrics 6, 4171–4276.

Aguirregabiria, V. and C. Ho (2012). A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. Journal of Econometrics 168, 156–173.

Arcidiacono, P. and R. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. Econometrica 79(6), 1823–1867.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. Econometrica 75(5), 1331 – 1370.

Benkard, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. Review of Economic Studies 71(3), 581 – 611.

Berry, S. (1992). Estimation of a model of entry in the airline industry. Econometrica 60(4), 889–917.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. Econometrica 60, 889–917.

Berry, S. and A. Pakes (1993). Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. American Economic Review 83(2), 247 – 252.

Boguslaski, C., H. Ito, and D. Lee (2004). Entry patterns in the southwest airlines route system. Review of Industrial Organization 25, 317–350.

Borenstein, S. (1989). Hubs and high fares: Dominance and market power in the us airline industry. RAND Journal of Economics 20(3), 344–365.

Borenstein, S. (1990). Airline mergers, airport dominance, and market power. American Economics Review 80(2), 400–404.

Borenstein, S. (1991). The dominant-firm advantage in multiproduct industries: Evidence from the us airlines. Quarterly Journal of Economics 106(4), 1237–66.

Ciliberto, F., C. Murry, and E. Tamer (2018). Market structure and competition in airline markets.

Ciliberto, F. and E. Tamer (2009). Market structure and multiple equilibria in airline markets. Econometrica 77(6), 1791–1828.

Collard-Wexler, A. (2014). Mergers and sunk costs: An application to the ready-mix concrete industry. American Economic Journal: Microeconomics 6(4), 407–47.

Doraszelski, U. and M. Satterthwaite (2010). Computable markov-perfect industry dynamics. RAND Journal of Economics 41(2), 215–243.

Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. Review of Economic Studies 62(1), 53 – 82.

Fershtman, C. and A. Pakes (2000). A dynamic oligopoly with collusion and price wars. RAND Journal of Economics 31(2), 207 – 236.

Goolsbee, A. and C. Syverson (2008). How do incumbents respond to the threat of entry? evidence from the major airlines. The Quarterly Journal of Economics 123(4), 1611–1633.

Gowrisankaran, G. (1999). A dynamic model of endogenous horizontal mergers. RAND Journal of Economics 30(1), 56 – 83.

Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. The Review of Economic Studies 60(3), 497–529.

Jeziorski, P. (2014a). Effects of mergers in two-sided markets: The u.s. radio industry. American Economic Journal: Microeconomics 6(4), 35–73.

Jeziorski, P. (2014b). Estimation of cost synergies from mergers: Application to the u.s. radio. RAND Journal of Economics 45(4), 816–846.

Kasahara, H. and K. Shimotsu (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. Econometrica 77(1), 135 – 175.

Kim, E. and V. Singal (1993). Mergers and market power: Evidence from the airline industry. American Economic Review 83(3), 549–569.

Lazarev, J. (2020). The identification power of the markov assumption in dynamic discrete choice models. Working paper, New York University.

Li, S., J. Mazur, Y. Park, J. Roberts, A. Sweeting, and J. Jun (2018). Repositioning and market power after airline mergers. Mimeo, University of Maryland.

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. Rand Journal of Economics 31(3), 395–421.

Nocke, V. and M. D. Whinston (2010). Dynamic merger review. Journal of Political Economy 118(6), 1200–1251.

Pakes, A. and P. McGuire (1994). Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. RAND Journal of Economics 25(4), 555 – 589.

Peters, C. (2006). Evaluating the performance of merger simulation: Evidence from the u.s. airline industry. Journal of Law and Economics 44, 627–649.

Reiss, P. and P. Spiller (1989). Competition and entry in small airline markets. Journal of Law and Economics 32, 179–202.

Shapiro, C. (1996). Mergers with differentiated products. Antitrust 10, 23–30.

Sinclair, R. (1995). An empirical model of entry and exit in airline markets. Review of Industrial Organization 10, 541–557.

Stahl, J. (2009). A dynamic analysis of consolidation in the broadcast television industry. Working Paper, Board of Governors of the Federal Reserve System.

US Department of Justice (2010). Horizontal Merger Guidelines. U.S. Department of Justice.

A Data Appendix

As an example of the CSA aggregation, the CSA containing San Francisco contains the Oakland International Airport (OAK), the San Francisco International Airport (SFO), and the Mineta San Jose International Airport (SJC). Once the data was aggregated, passengers from all three airports in the San Francisco Bay Area CSA were treated as originating from the CSA as opposed to the individual airports within the CSA. This aggregation captures the fact that these airports are substitutes both for passenger traffic and for airline entry decisions.

The portion of the T100 data set that we use contains quarterly data on passenger enplanements for each airline on segments connecting between the 60 demographic areas of interest for our study. The segment data is in principle so accurate that if a NY-LA flight is diverted to San Diego due to weather, then it shows up in the data as having flown to San Diego. This leads to there being a fair amount of "phantom" entry occurrences in the raw data. To weed out these one-off flights, an airline is defined to have entered a segment that it had not previously served if it sends 9000 or more enplaned passengers on the segment per quarter for four successive quarters. The level chosen is roughly equivalent to running one daily nonstop flight on the segment, a very low level of service for a regularly scheduled flight. For example, if airline X sends at least 9000 passengers per quarter along segment Y from the third quarter of 2005 through the second quarter of 2006 (inclusively), then it is defined to have entered segment Y in the third quarter of 2005. If an airline entered a segment in any quarter of a given year, then it is said to have entered during that year. Once an airline has entered a segment, it is considered present on that segment until an exit even has occurred. We define exit event symmetrically with our entry definition. If an airline is defined to be "In" on a segment, four successive quarters with fewer than 9000 passengers enplaned on the segment defines an exit event. Therefore, if airline X had been in on segment Y in quarter 2 of 2005, but from quarter 3 of 2005 through quarter 2 of 2006 the airline had fewer than 9000 enplanned passengers, the airline is noted as having exited segment Y in quarter 3 of 2005. Once an airline has entered a segment, it is defined as present on that segment until an exit even occurs for that airline on that segment. Similarly, once an airline has exited a segment, it is defined as not present on the segment until an entry event occurs. The data on segment presence is initialized by defining an airline as present if it had 9000 or more enplaned passengers on a segment in quarter 1 of 2003 and not present otherwise.

A.1 Hub Definitions by CSA

American: Los Angeles, CA; Ft. Lauderdale, FL; Chicago, IL; St. Louis, MS; New York, NY; Dallas, TX

United: Los Angeles, CA; San Francisco, CA; Denver, CO; Washington, D.C.; Chicago, IL

Southwest: Phoenix, AZ; Los Angeles, CA; Oakland, CA; Denver, CO; Chicago, IL; Baltimore, MD; Las Vegas, NV

Delta: Atlanta, GA; New York, NY; Cincinnati, OH; Salt Lake City, UT

Continental: New York, NY; Cleveland, OH; Houston, TX

Northwest: Detroit, MI; Minneapolis/St. Paul, MN; Memphis, TN

US Airways: Washington, D.C.; Charlotte, NC; Philadelphia, PA

JetBlue: Long Beach, CA; Fort Lauderdale, FL; New York, NY

American West: Phoenix, AZ; Las Vegas, NV

Alaska: Anchorage, AK; Los Angeles, CA; Portland, OR; Seattle, WA

A.2 CSA Airport Correspondences

CSA code	CSA Airports	Pop 2000	Median Inc.	# pass (mark, 2000)	# deps 2000
12	BUR, LAX, ONT, SNA	16373645	52069	63366291	651974
32	MDW, ORD	9312255	54421	62343200	699212
22	EWR, JFK, LGA	21361797	56978	58882013	689529
4	ATL	4548344	52957	55337406	499976
37	OAK, SFO, SJC	7092596	66657	51131131	503844
18	DAL, DFW	5346119	49146	49770836	580463
13	BWI, DCA, IAD	7538385	67752	42311686	514799
45	PHX	3251876	48124	33102813	367510
26	HOU, IAH	4815122	46480	31547559	388080
19	DEN	2449054	55149	31311309	300264
29	LAS	1408250	49171	31081307	299968
10	BOS, MHT, PVD	1582997	51310	29349066	360982
23	FLL, MIA	5007564	43091	29309146	275868
57	STL	2698687	48361	25674940	303880
31	MCO	1697906	43952	25459140	236478
20	DTW	5357538	50471	25396816	280110
35	MSP	3271888	58459	25124724	267797
53	SEA	3604165	53900	22497342	238320
44	PHL	5833585	53266	18812458	241778
55	SLC	1454259	50357	16205369	148173
15	CLT	1897034	44402	16052317	198542
17	CVG	2050175	48022	15283486	197718
50	SAN	2813833	56335	15118565	163921
58	TPA	2395997	41852	14373207	144221
46	PIT	2525730	41648	13979823	182791
43	PDX	1927881	49227	12134527	150319
30	MCI	1901070	50179	11320857	151568
14	CLE	2945831	44049	10842047	192681
25	HNL	876156	60485	10320878	71179
36	MSY	1360436	39479	9497691	108138
47	RDU	1314589	49449	9221253	137888
33	MEM	1205204	41065	8651773	118131
8	BNA	1381287	45194	8552027	120258
56	SMF	1930149	54071	7728952	80867
54	SJU	2509007	19403	7067099	51241
6	BDL	1257709	59912	6963738	84986
5	AUS	1249763	50484	6950039	82864
27	IND	1843588	48399	6885666	93134
51	SAT	1711703	43263	6624018	77632
16	CMH	1835189	47075	6163317	89701
1	ABQ	729649	43070	5871686	71116
34	MKE	1689572	47799	5445851	90630
42	PBI	5007564	43091	5376385	51452
48	RNO	342885	48974	5294211	61475
28	JAX	1122750	47323	4955361	60860
38	OGG	128094	57573	4840509	49519
49	RSW	2395997	41852	4629297	42883
11	BUF	1170111	41947	3770970	54207
52	SDF	1292482	42943	3702821	57119
40	OMA	803201	48826	3585827	49920
60	TUS	843746	41521	3500323	39440
39	OKC	1160942	39743	3367555	53260
59	TUL	908528	40512	3253687	53582
21	ELP	679622	30968	3142143	47032
24	GEG	417939	41667	2933340	42947
7	BHM	1129721	43290	2884829	43839
9	BOI	464840	46960	2667242	41537
41	ORF	234403	31815	2577507	39326
2	ALB	825875	50828	2438339	37108
3	ANC	319605	60180	2293263	21837

A.3 Regional Carriers

To account for flights operated by regional carriers, we performed the following steps. First, using a publicly available representative sample of all airline tickets sold in the United States (”DB1B”), we identified routes operated by regional carriers. For each route and regional carrier, we identified how many tickets each mainline carrier sold on each regional partner within a given quarter and the corresponding shares. We then took these shares, multiplied by the total traffic of the regional carrier and added it directly to total traffic of the corresponding mainland carrier. We used this combined traffic to determine whether or not a given mainline airline is currently operating in a given route.

The table below illustrates why accounting for regional carriers may make a significant difference in our network analysis. In particular, it shows that failing to account for regional carriers will systematically distort the full scope of all mainline carriers' networks. Any model that tries to rationalize such distorted data will inevitably struggle, especially given the fact that mainline carriers regularly change the set of regional carriers they contract with and do sometimes alternate between regional service with mainline service based on both demand and cost factors.

Table A.1: Fraction of Routes Served Without Regional Affiliates

Fraction of Routes Served Without Regional Affiliates

Year	2002	2003	2004	2005	2006	2007
American	0.832	0.816	0.767	0.742	0.698	0.700
United	0.783	0.681	0.642	0.662	0.670	0.660
Southwest	1.000	0.993	0.942	0.948	0.966	0.991
Delta	0.795	0.714	0.685	0.63	0.620	0.596
Continental	0.923	0.877	0.892	0.888	0.640	0.630
Northwest	0.893	0.879	0.804	0.793	0.814	0.796
US Airways	0.849	0.758	0.885	0.889	0.859	0.873
JetBlue	1.000	0.941	0.654	0.813	0.608	0.880
AmericaWest	0.970	0.910	0.880	N/A	N/A	N/A
Alaska	0.595	0.692	0.775	0.756	0.767	0.721

B Tables and Figures

Table 1: Airline Route and Market Statistics, 2003-2008

	Nonstop Routes					Markets		
Carrier	**Avg**	**Min**	**Max**	**Entry/yr**	**Exit/yr**	**Avg**	**Min**	**Max**
American	224	219	232	7	8	1260	1237	1296
United	182	166	193	6	2	1331	1237	1372
Delta	230	220	241	14	14	1453	1400	1504
Continental	121	103	147	10	2	920	772	1126
Northwest	155	136	169	6	2	1173	1145	1215
US Airways	158	146	190	14	6	730	665	982
Southwest	298	269	323	15	4	937	824	1042
JetBlue	32	16	51	8	1	128	61	226
Alaska	41	37	43	2	1	115	94	123
DL + NW	373	349	386	18	14	1566	1550	1579
UA + US	309	292	341	16	7	1455	1379	1494
UA + CO	286	254	321	15	3	1485	1396	1523

Table 2: Airline Route and Market Statistics, 2003-2008

Regressor	**Avg**	**SD**	**Min**	**25%**	**50%**	**75%**	**Max**
City Pair Characteristics:							
Log(2002 Pass Dens)	7.6	5.6	0	0	11	13	16
Pop1*Pop2(*1e-12)*Dens=0	0.82	3.2	0	0	0	0.34	82
Log(Pass. Den. New Markets)	2.7	4.6	0	0	0	5.5	16
% Tourist	0.37	0.35	0	0	0.33	0.67	1
Competition Variables:							
Number Non-Stop Comps.	0.76	0.99	0	0	0	1	6
Number One-Stop Comps.	3.5	2	0	2	4	5	9
Number CS Agreements	0.051	0.23	0	0	0	0	3
HHI Among Others (Market)	0.49	0.44	0	0	0.51	1	1
HHI Among Others Large (City)	0.32	0.15	0.012	0.2	0.28	0.42	0.72
HHI Among Others Small (City)	0.17	0.079	0.0054	0.13	0.16	0.2	0.68
Own Share Large (City)	0.16	0.16	0	0.051	0.11	0.21	0.85
Own Share Small (City)	0.056	0.068	0	0	0.042	0.074	0.77
Competitor Hub on Route	0.68	0.47	0	0	1	1	1
Network Characteristics:							
Present in Segment	0.09	0.29	0	0	0	0	1
Present in Market (not Seg)	0.41	0.49	0	0	0	1	1
Present Both Apts (not Market)	0.18	0.38	0	0	0	0	1
Number of Hubs	0.15	0.37	0	0	0	0	2
Hub Conv (NS dist/OS dist)	0.76	0.28	0.011	0.57	0.89	0.99	1
Dist Nearest Hub Large (100s)	12	9.3	0	5	8.6	18	48
Dist Nearest Hub Small (100s)	4.4	4.9	0	1.2	2.9	5.5	47
# Nonstops Large (City)	8.4	12	0	2	4	8	56
# Nonstops Small (City)	2.3	3.1	0	0	2	3	53
Distance Variables:							
Distance > 250	0.95	0.21	0	1	1	1	1
Distance > 500	0.84	0.37	0	1	1	1	1
Distance > 1000	0.58	0.49	0	0	1	1	1
Distance > 1500	0.37	0.48	0	0	0	1	1
Distance > 2000	0.22	0.42	0	0	0	0	1
Distance > 2500	0.11	0.32	0	0	0	0	1
Distance > 3000	0.075	0.26	0	0	0	0	1

Table 3: U.S. Airline Route Network Competition

This table lists the total number of segments/markets flown by each airline, followed by the number of segments/markets where they are the only carrier, where there is one additional carrier, etc.

		with number of competitors equal to											
2008: segments	Total	0	1	2	3	4	5	6	7	8	9	10	Avg
American (AA)	223	21	49	66	41	31	11	4	0	0	0	0	2.27
United (UA)	190	4	31	71	49	22	9	4	0	0	0	0	2.51
Southwest (WN)	323	51	94	92	64	14	7	1	0	0	0	0	1.76
Delta (DL)	220	64	66	35	17	21	13	4	0	0	0	0	1.64
Continental (CO)	146	30	45	28	13	18	9	3	0	0	0	0	1.88
Northwest (NW)	157	42	60	33	15	5	1	1	0	0	0	0	1.29
US Airways (US)	190	30	46	54	38	13	8	1	0	0	0	0	1.93
JetBlue (B6)	50	0	4	8	10	14	11	3	0	0	0	0	3.58
Alaska (AS)	43	6	17	11	3	3	3	0	0	0	0	0	1.74
DL + NW	366	108	125	63	33	21	13	3	0	0	0	0	1.41
UA + US	341	35	85	121	61	28	8	3	0	0	0	0	1.99
UA + CO	320	34	78	99	57	38	13	1	0	0	0	0	2.09

		with number of competitors equal to											
2008: markets	Total	0	1	2	3	4	5	6	7	8	9	10	Avg
American (AA)	1272	13	29	58	105	174	237	261	219	120	43	13	5.44
United (UA)	1366	6	21	87	113	209	271	265	218	120	43	13	5.36
Southwest (WN)	1042	11	49	64	83	136	169	197	168	114	38	13	5.33
Delta (DL)	1489	13	50	99	143	238	274	276	220	120	43	13	5.15
Continental (CO)	1125	7	14	33	67	152	217	242	217	120	43	13	5.71
Northwest (NW)	1145	15	19	59	80	153	204	234	205	120	43	13	5.52
US Airways (US)	982	5	21	42	55	107	152	221	203	120	43	13	5.79
JetBlue (B6)	226	0	0	1	3	7	21	29	50	59	43	13	7.33
Alaska (AS)	123	2	11	12	12	17	14	14	1	13	14	13	5.37
DL + NW	1580	31	97	150	249	303	312	247	135	43	13	0	4.31
UA + US	1483	13	57	121	204	286	342	265	139	43	13	0	4.58
UA + CO	1526	13	38	144	250	329	311	260	125	43	13	0	4.48

Note: the 13 markets that are served by ALL 11 carriers are as follows:
Boston - Los Angeles, Boston - Las Vegas, Boston - San Francisco, Boston - Phoenix, Boston - San Diego, Los Angeles - Washington, Los Angeles - Miami, Los Angeles - Orlando, Washington - Las Vegas, Washington - San Francisco, Washington - San Diego, Miami - San Francisco, Orlando - San Francisco

Table 4: Top 10 Nonstop Routes by HHI Increase, Passengers Enplaned, 2008

DL-NW				
		HHI Passengers		
CSA1	CSA2	Pre	Post	Chng
CVG	MSP	5056	9982	4926
CVG	DTW	4954	9875	4921
BHM	MSP	5156	10000	4844
MSP	SLC	5237	9902	4665
DTW	SLC	5263	9885	4622
ATL	DTW	3606	6622	3016
ATL	MSP	3494	5812	2318
MEM	SAN	7678	9468	1790
BDL	MEM	7632	9255	1623
ATL	MEM	4016	5512	1496
UA-US				
		HHI Passengers		
CSA1	CSA2	Pre	Post	Chng
OAK, SFO, SJC	PHL	5467	9981	4514
CLT	DEN	5931	10000	4069
CLT	MDW, ORD	4342	8064	3722
BUR, LAX, ONT, SNA	PHL	6442	9978	3536
OAK, SFO, SJC	PIT	6845	9975	3130
BWI, DCA, IAD	MSY	3622	6720	3098
BWI, DCA, IAD	PHL	6616	9487	2871
CLT	OAK, SFO, SJC	7235	10000	2765
DEN	PHL	2785	4894	2109
BWI, DCA, IAD	PIT	3337	5255	1918
UA-CO				
		HHI Passengers		
CSA1	CSA2	Pre	Post	Chng
CLE	DEN	5276	9804	4528
DEN	HOU, IAH	3227	5451	2224
DEN	EWR, JFK, LGA	3385	5174	1789
BWI, DCA, IAD	CLE	3808	5185	1377
HOU, IAH	MDW, ORD	3031	4305	1274
CLE	MDW, ORD	2892	3901	1009
BWI, DCA, IAD	HOU, IAH	5896	6847	951
HOU, IAH	OAK, SFO, SJC	6963	7906	943
EWR, JFK, LGA	OAK, SFO, SJC	1783	2629	846
EWR, JFK, LGA	MDW, ORD	2989	3790	801

Table 5: Top 5 Cities by HHI Increase, Passengers Enplaned, 2008

DL-NW			
	HHI Passengers		
CSA	Pre	Post	Chng
MEM	5284	6591	1307
MSP	5372	6013	641
CVG	7558	8131	573
DTW	4912	5458	546
BDL	1679	2106	427
UA-US			
	HHI Passengers		
CSA	Pre	Post	Chng
CLT	6533	7477	944
PHL	3357	4104	747
BWI, DCA, IAD	1559	2288	729
PIT	1772	2442	670
ALB	2154	2712	558
UA-CO			
	HHI Passengers		
CSA	Pre	Post	Chng
CLE	4108	4778	670
EWR, JFK, LGA	1631	1943	312
OMA	1501	1787	286
HOU, IAH	4738	4988	250
MSY	1626	1869	243

Table 6: Probit for Entry/Exit/Stay, All Carriers

Variable	Units/ Range	Beta	SE	Marg Eff	1SD Marg Eff
Demand Vars:					
Log(2002 Pass Dens)	[0,16]	0.043	0.012	0.02	0.10
Pop1*Pop2(*1e-12)*Dens=0	[0,82]	0.014	0.010	0.01	0.02
Log Pass. Den. New Markets	[0,0.02]	11.8	4.4	4.7	0.02
% Tourist	[0,1]	0.14	0.07	0.06	0.02
Competition Vars:					
Number NonStop Comps.	{0,...,6}	-0.14	0.03	-0.06	-0.06
Number One-Stop Comps.	{0,...,9}	-0.01	0.02	-0.00	-0.01
Number CS Agreements	{0,...,3}	0.22	0.06	0.09	0.02
Competitor Hub on Route	{0,1}	-0.17	0.07	-0.07	-0.03
HHI Among Others (Market)	[0,1]	-0.37	0.06	-0.15	-0.06
HHI Among Others Large (City)	[0,1]	2.08	0.41	0.83	0.12
HHI Among Others Small (City)	[0,1]	1.19	0.74	0.48	0.04
Own Share Large (City)	[0,1]	3.81	0.46	1.52	0.25
Own Share Small (City)	[0,1]	2.53	0.55	1.01	0.07
Network Vars:					
Present in Route	{0,1}	3.25	0.09	1.30	0.38
Present in Market (not Route)	{0,1}	0.34	0.08	0.13	0.07
Present Both Apts (not Market)	{0,1}	0.26	0.08	0.10	0.04
Number of Hubs	{0,1,2}	0.58	0.07	0.23	0.09
Hub Conv (NS dist/OS dist)	[0,1]	-0.36	0.15	-0.15	-0.04
Dist Nearest Hub Large (100s)	1000mi	0.06	0.06	0.03	0.02
Dist Nearest Hub Small (100s)	1000mi	-0.20	0.11	-0.08	-0.04
# Nonstops Large (City)	{0,...,57}	0.11	0.03	0.04	0.05
# Nonstops Small (City)	{0,...,54}	-0.07	0.10	-0.03	-0.01
Distance > 250	{0,1}	0.24	0.09	0.10	0.02
Distance > 500	{0,1}	-0.19	0.07	-0.08	-0.03
Distance > 1000	{0,1}	-0.24	0.07	-0.10	-0.05
Distance > 1500	{0,1}	-0.18	0.08	-0.07	-0.03
Distance > 2000	{0,1}	-0.03	0.09	-0.01	-0.01
Distance > 2500	{0,1}	-0.14	0.13	-0.06	-0.02
Distance > 3000	{0,1}	-0.80	0.27	-0.32	-0.08
JetBlue dummy	{0,1}	0.58	0.09	0.23	0.07
N		79650			
Likelihood		-2639			
Fixed Effects		Year, City			
Test for Markov unobservables:		Coeff	SE	p-value	
		0.239	0.083	0.004	

Table 7: Probit for Entry/Exit/Stay, All Carriers, Route Fixed Effects

Variable	Units/ Range	Beta	SE	Marg Eff	1SD Marg Eff
Demand Vars:					
Log Pass Dens New Mkts	[0,0.02]	10.12	5.73	4.04	0.02
Direct Competitors:					
1 Nonstop Comp	{0,1}	-1.93	0.13	-0.77	-0.35
2 Nonstop Comps	{0,1}	-2.93	0.17	-1.17	-0.39
3 Nonstop Comps	{0,1}	-3.79	0.20	-1.51	-0.32
4 Nonstop Comps	{0,1}	-4.79	0.27	-1.91	-0.24
>4 Nonstop Comps	{0,1}	-5.27	0.32	-2.10	-0.15
Other Comp Vars:					
Number CS Agreements	{0,...,3}	0.21	0.08	0.08	0.02
Competitor Hub on Route	{0,1}	-0.57	0.16	-0.23	-0.11
HHI Among Others Large (City)	[0,1]	1.98	0.58	0.79	0.12
HHI Among Others Small (City)	[0,1]	-1.18	1.13	-0.47	-0.04
Own Share Large (City)	[0,1]	2.93	0.62	1.17	0.19
Own Share Small (City)	[0,1]	1.79	0.77	0.72	0.05
Network Vars:					
Present in Route	{0,1}	2.48	0.08	0.99	0.29
Present in Mkt (not Route)	{0,1}	0.18	0.08	0.07	0.04
Number of Hubs	{0,...,2}	0.51	0.09	0.21	0.08
Hub Conv (NS dist/OS dist)	[0,1]	-0.42	0.21	-0.17	-0.05
Dist Nearest Hub Large (100s)	1000mi	0.12	0.08	0.05	0.04
Dist Nearest Hub Small (100s)	1000mi	-0.27	0.15	-0.11	-0.05
# Nonstops Large (City)	{0,...,57}	0.27	0.05	0.11	0.13
# Nonstops Small (City)	{0,...,54}	0.24	0.14	0.09	0.03

N	79650			
Likelihood	-2114			
Fixed Effects	Year, Route			
Test for Markov unobservables:	Coeff	SE	p-value	
	0.015	0.099	0.883	

Table 8: Measures of Fit by Airline: All Airlines Pooled, Route FE's

	Actual Last Period Status				Full Sample Simulated	
Airline	Stay In	Stay Out	New Entry	New Exit	New Entry	New Exit
American (25,27)	0.979	0.995	0.177	0.174	0.523	0.578
United (24,17)	0.978	0.996	0.211	0.263	0.653	0.748
Southwest (66,12)	0.973	0.989	0.234	0.173	0.658	0.604
Delta (31,52)	0.972	0.995	0.191	0.237	0.569	0.845
Continental (39,7)	0.978	0.996	0.248	0.169	0.703	0.791
Northwest (19,11)	0.985	0.998	0.085	0.215	0.523	0.784
US Airways (86,29)	0.973	0.996	0.671	0.231	0.757	0.687
JetBlue (33,0)	0.926	0.996	0.125	0.396	0.446	NaN
Alaska (5,1)	0.971	0.999	0.138	0.417	0.490	1.000

Table 9: Model Selection: Probit and ANN dimension

Model	CV Likelihood
Probit	-618
ANN (dim=0)	-630
ANN (dim=1)	-641
ANN (dim=2)	-686
ANN (dim=3)	-697
ANN (dim=4)	-730
ANN (dim=5)	-757

Table 10: 10 year simulations, Median Nonstop Routes Served

Median number of routes served, by year

Year	**0**	1	2	3	4	5	6	7	8	9	10
	No merger										
American	227	225	224	223	223	222	221	221	220	218	217
United	198	199	199	199	198	198	197	196	196	195	194
Southwest	325	332	337	343	349	354	360	365	370	375	379
Delta	220	215	211	207	203	199	196	193	190	187	184
Continental	146	145	145	145	144	144	143	143	142	142	141
Northwest	157	157	156	155	154	153	151	150	149	148	147
USAirways	227	224	220	217	214	211	209	206	203	201	198
JetBlue	50	53	57	61	65	69	74	79	83	88	93
Alaska	43	42	42	42	42	42	42	42	42	41	41
	DL-NW merger										
American	227	225	223	222	220	219	217	216	214	213	211
United	198	199	198	198	197	196	195	194	193	192	191
Southwest	325	330	333	336	338	340	342	344	345	346	348
DL + NW	366	371	373	375	376	377	378	379	380	381	381
Continental	146	146	145	144	143	141	140	139	138	137	136
-merged-	0	0	0	0	0	0	0	0	0	0	0
USAirways	227	227	226	225	224	223	222	221	221	220	219
JetBlue	50	54	58	61	64	67	71	74	78	81	85
Alaska	43	43	43	43	42	42	42	42	42	42	42
	UA-US merger										
American	227	229	228	227	226	225	223	222	220	219	217
UA + US	358	366	370	374	378	381	384	387	389	391	393
Southwest	325	337	343	347	351	354	356	358	359	361	361
Delta	220	219	216	212	207	203	199	195	192	189	186
Continental	146	147	147	146	145	144	143	142	141	141	140
Northwest	157	160	159	158	157	156	154	153	151	150	149
-merged-	0	0	0	0	0	0	0	0	0	0	0
JetBlue	50	59	65	71	77	83	89	95	101	108	114
Alaska	43	44	44	44	44	44	44	44	44	44	44
	UA-CO merger										
American	227	226	224	223	221	220	219	217	216	215	213
UA + CO	326	331	335	337	340	342	343	345	346	347	348
Southwest	325	331	334	338	341	343	346	348	349	351	352
Delta	220	216	211	207	202	198	194	191	188	185	183
-merged-	0	0	0	0	0	0	0	0	0	0	0
Northwest	157	158	157	155	154	152	151	150	148	147	146
USAirways	227	228	228	228	228	228	228	228	228	228	228
JetBlue	50	54	58	62	66	71	75	80	84	89	93
Alaska	43	43	43	43	43	43	43	43	43	43	43

Table 11: Nonstop Routes Served: Simulated Distribution in Year 10

	Number of Nonstop Routes Served								Number of Markets Served							
Carrier	**base**	mean	std	min	max	q0.25	med	q0.75	**base**	mean	std	min	max	q0.25	med	q0.75
	No merger															
American	227	217	8	189	254	212	217	223	1284	1290	30	1195	1421	1268	1286	1312
United	198	193	8	165	221	188	194	199	1370	1364	34	1209	1459	1343	1367	1388
Southwest	325	379	12	336	420	372	379	387	1074	1459	48	1254	1611	1428	1462	1493
Delta	220	185	8	158	221	179	184	190	1489	1414	28	1300	1498	1396	1417	1435
Continental	146	141	4	121	158	138	141	144	1125	1133	26	1031	1201	1112	1134	1155
Northwest	157	147	4	132	165	144	147	150	1145	1136	11	1059	1195	1131	1135	1141
USAirways	227	198	9	165	235	192	198	204	1235	1116	50	913	1252	1084	1120	1153
JetBlue	50	93	10	57	134	87	93	100	226	430	55	195	625	393	430	468
Alaska	43	41	4	24	57	38	41	44	123	167	23	82	246	152	168	183
	DL-NW merger															
American	227	211	7	186	242	206	211	216	1284	1292	31	1181	1424	1268	1290	1314
United	198	190	7	166	218	185	191	195	1370	1377	31	1244	1452	1357	1379	1399
Southwest	325	348	12	299	389	339	348	356	1074	1358	55	1140	1517	1321	1360	1397
DL + NW	366	381	11	345	426	374	381	389	1580	1585	10	1521	1619	1579	1586	1591
Continental	146	136	5	116	149	133	136	139	1125	1123	26	1028	1189	1102	1124	1145
-merged-	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
USAirways	227	219	8	187	251	214	219	225	1235	1136	47	920	1287	1107	1141	1170
JetBlue	50	85	10	48	121	78	85	91	226	427	55	220	645	388	426	464
Alaska	43	42	4	27	58	39	42	45	123	161	25	74	240	144	162	180
	UA-US merger															
American	227	217	8	193	249	212	217	222	1284	1297	29	1173	1415	1275	1292	1316
UA + US	358	393	9	358	428	387	393	399	1512	1540	13	1466	1580	1532	1541	1550
Southwest	325	361	12	320	405	354	361	369	1074	1417	46	1222	1555	1386	1419	1450
Delta	220	186	8	160	214	181	186	191	1489	1460	17	1371	1512	1450	1461	1472
Continental	146	140	5	120	159	137	140	143	1125	1130	26	1030	1195	1110	1133	1151
Northwest	157	149	5	133	170	145	149	152	1145	1146	15	1089	1222	1136	1144	1154
-merged-	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
JetBlue	50	114	11	77	174	106	114	122	226	503	56	297	721	466	503	540
Alaska	43	44	4	29	61	41	44	47	123	169	23	75	253	154	170	185
	UA-CO merger															
American	227	213	7	188	244	208	213	218	1284	1289	29	1168	1421	1267	1285	1311
UA + CO	326	348	9	316	385	342	348	354	1530	1565	13	1497	1602	1557	1566	1574
Southwest	325	352	12	309	399	344	352	360	1074	1362	54	1140	1539	1326	1363	1400
Delta	220	183	6	158	207	179	183	187	1489	1457	17	1378	1504	1447	1458	1469
-merged-	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Northwest	157	146	4	132	165	143	146	149	1145	1144	15	1073	1220	1134	1141	1151
USAirways	227	228	8	196	258	222	228	233	1235	1142	47	925	1277	1112	1146	1176
JetBlue	50	94	10	62	131	87	93	100	226	444	54	266	666	406	443	480
Alaska	43	43	4	28	59	41	43	46	123	169	23	79	239	154	170	186

Table 12: Top 5 Cities, Static vs Simulated Dynamic Effect

DL-NW					
		Static		Dynamic	
CSA	# Carriers (Pre)	Pre	Post (Yr 0)	No Merger (Yr 10)	Merger (Yr 10)
MEM	6	5904	6451	5679	6571
MSP	6	5861	6379	5676	6215
CVG	6	5977	6358	6130	6332
DTW	7	4374	4918	4460	4898
BDL	7	1787	2202	1559	2272
UA-US					
		Static		Dynamic	
CSA	# Carriers (Pre)	Pre	Post (Yr 0)	No Merger (Yr 10)	Merger (Yr 10)
CLT	7	4452	5592	4482	4512
PHL	7	3363	3954	2966	3146
BWI, DCA, IAD	9	2098	2726	2313	2542
PIT	8	2580	2996	1898	3272
ALB	7	1882	2188	1655	1882
UA-CO					
		Static		Dynamic	
CSA	# Carriers (Pre)	Pre	Post (Yr 0)	No Merger (Yr 10)	Merger (Yr 10)
CLE	7	4271	4683	4206	4355
EWR, JFK, LGA	9	1946	2119	1961	2078
OMA	7	1524	1745	1524	1745
HOU, IAH	8	3921	4261	3866	4160
MSY	8	1678	1886	1706	1853

Note: HHIs in this table are by # routes served and therefore differ from those in Table 5.

Table 13: Simulated Nonstop Route-Level Market Structures in Year 10

	Year 0	Year 10						
Number of ...		mean	std	min	max	q0.25	med	q0.75
		No merger						
markets with 0 carriers	846	846	4	832	869	843	846	849
markets with 1 carrier	511	512	7	488	538	507	512	517
markets with 2 carriers	250	243	7	214	271	238	243	248
markets with 3 carriers	98	102	6	80	125	98	102	106
markets with >=4 carriers	65	66	4	51	82	64	66	69
		DL-NW merger						
markets with 0 carriers	846	834	4	820	848	831	833	836
markets with 1 carrier	517	526	7	501	552	522	526	531
markets with 2 carriers	244	240	8	210	265	235	240	245
markets with 3 carriers	100	101	6	77	127	97	101	105
markets with >=4 carriers	63	69	4	55	86	67	69	72
		UA-US merger						
markets with 0 carriers	846	834	4	820	852	831	834	836
markets with 1 carrier	520	515	7	487	540	511	515	520
markets with 2 carriers	265	256	8	227	285	250	256	261
markets with 3 carriers	98	106	6	84	130	102	106	110
markets with >=4 carriers	41	59	4	44	77	57	59	62
		UA-CO merger						
markets with 0 carriers	846	837	4	822	852	834	837	839
markets with 1 carrier	512	517	7	491	543	512	517	521
markets with 2 carriers	254	247	8	220	276	242	247	252
markets with 3 carriers	95	104	6	81	129	100	104	108
markets with >=4 carriers	63	66	4	50	81	63	66	68

Table 14: Simulated Market-Level Market Structures in Year 10

	Year 0	Year 10						
Number of ...		mean	std	min	max	q0.25	med	q0.75
		No merger						
markets with 0 carriers	25	19	2	12	28	17	19	20
markets with 1 carrier	77	67	5	52	93	64	67	70
markets with 2 carriers	113	115	8	91	153	109	115	120
markets with 3 carriers	179	161	12	121	227	152	160	169
markets with >=4 carriers	1376	1408	16	1331	1453	1398	1409	1419
		DL-NW merger						
markets with 0 carriers	25	17	2	11	26	16	17	19
markets with 1 carrier	80	69	5	55	90	66	69	72
markets with 2 carriers	140	135	8	103	187	129	134	139
markets with 3 carriers	207	187	18	136	263	173	185	198
markets with >=4 carriers	1318	1362	22	1262	1426	1348	1364	1378
		UA-US merger						
markets with 0 carriers	25	17	2	11	27	16	17	19
markets with 1 carrier	80	67	4	54	85	64	67	70
markets with 2 carriers	149	119	9	90	161	113	119	125
markets with 3 carriers	248	207	13	163	257	197	206	215
markets with >=4 carriers	1268	1360	18	1289	1412	1348	1361	1373
		UA-CO merger						
markets with 0 carriers	25	17	2	10	27	16	17	19
markets with 1 carrier	75	67	4	53	86	64	66	69
markets with 2 carriers	118	113	7	90	148	108	112	117
markets with 3 carriers	203	179	12	139	231	171	178	186
markets with >=4 carriers	1349	1395	14	1317	1435	1386	1396	1405