Price Competition and Endogenous Product Choice in Networks: Evidence from the US Airline Industry[*]

Christian Bontemps[†] Cristina Gualdani[‡] Kevin Remmy[§]

March 2025

Abstract

We develop a two-stage game in which competing airlines first choose the networks of markets to serve in the first stage before competing in price in the second stage. Spillovers in entry decisions across markets are allowed, which accrue on the demand, marginal cost, and fixed cost sides. We show that the second-stage parameters are point identified, and we design a tractable procedure to set identify the first-stage parameters and to conduct inference. Further, we estimate the model using data from the domestic US airline market and find significant spillovers in entry. In a counterfactual exercise, we evaluate the 2013 merger between American Airlines and US Airways. Our results highlight that spillovers in entry and post-merger network readjustments play an important role in shaping post-merger outcomes.

Keywords: endogenous market structure, networks, airlines, oligopoly, product repositioning, mergers, remedies.

[*]We would like to thank Manuel Arellano, Steven Berry, Jason Blevins, Andrew Chesher, Pierre Dubois, Isis Durrmeyer, Natalia Fabra, Silke Forbes, Alessandro Gavazza, Rachel Griffith, Alessandro Iaria, Panle Jia, Robert de Jong, Myrto Kalouptsidi, Estelle Malavolti, David Martimort, Thierry Magnac, Pedro Mira, Francesca Molinari, Massimo Motta, Ariel Pakes, Aureo de Paula, Mathias Reynaert, Jörg Stoye, Miguel Urdanoz, Elie Tamer, Frank Verboven, and Matthew Weinberg, as well as participants at various conferences and seminars for helpful comments. We acknowledge funding from the French National Research Agency (ANR) under the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010), the German Research Foundation (DFG) through CRC TR 224 (Project A02), and the Economic and Social Research Council (grant ES/X011186/1). An earlier version of the paper was presented at the EARIE Annual Conference in 2018 under the title "A Two-Stage Model of Airline Competition".

[†]Email: christian.bontemps@tse-fr.eu, Ecole Nationale de l'Aviation Civile & Toulouse School of Economics, University of Toulouse, Toulouse, France.

[‡]Email: c.gualdani@qmul.ac.uk, Queen Mary University of London, London, United Kingdom.

[§]Email: remmy@uni-mannheim.de, University of Mannheim, Mannheim, Germany and CEPR.

1 Introduction

When evaluating a merger, antitrust authorities must trade off the costs of an increase in market power and the benefits of merger-induced efficiencies, putting upward and downward pressure on prices, respectively (Williamson, 1968). In network industries, such as the airline industry studied in this paper, post-merger market repositioning can amplify or nullify these effects. The hub-and-spoke system of this industry creates opportunities for the newly merged entity to offer passengers more destination choices, hence increasing the willingness-to-pay on the demand side and reducing costs on the supply side. Driven by such incentives, the merged entity may decide to re-optimize its network of routes.Moreover, rivals may react to the merger by exiting some markets in which the merged entity has become powerful and entering others where more competitors can remain profitable.

Despite these arguments, traditional merger analyses do not formally incorporate post-merger entry-exit patterns. To address this gap, we build and estimate a structural model of the airline market in which airlines choose their network of markets to transport passengers from one city to another. Our model allows us to conduct an exhaustive evaluation of airline mergers, thus considering the possibility for each airline to redefine the set of destinations offered to passengers. In particular, we quantify the importance of post-merger network re-optimization in shaping final outcomes using the 2013 merger between American Airlines and US Airways. We also evaluate the global effect of the remedies imposed by antitrust authorities on the merging parties in order to constrain post-merger network readjustments and protect consumer surplus.

Endogenizing entry decisions in a model for the airline industry is challenging. The presence of an airline in a given market affects the demand, marginal costs, and fixed costs of the itineraries offered by the same airline in neighboring markets and, hence, spills over into the airline's decision to operate in those neighboring markets. As a result, an airline does not take its entry decisions market-by-market but rather builds the network of served markets on a global basis so as to internalize spillovers in entry. These spillovers arise from the hub-and-spoke system operated by airlines. In addition to flights transporting passengers directly from one city to another, an airline can offer flights connecting cities via a common hub, which acts as a stopover point towards many final destinations. Such connecting flights can lead to marginal cost savings by activating economies of density and increase demand by boosting the value of loyalty programs in all the markets linked to the same hub. At the same time, they may increase fixed costs due to the risk of congestion at hubs.

We model airlines' decisions as a two-stage game. In the first stage, each airline forms its network by weighing changes in the fixed costs against changes in the expected variable profits. In the second stage, conditional on all the networks, airlines face demand for the offered itineraries, pay the variable costs, and choose the prices to charge while competing in a classic Bertrand-Nash pricing game. The two-stage structure of our model is similar to Eizenberg (2014) and Wollmann (2018), but we face additional methodological and computational challenges: the spillovers in entry break the usual separability across markets of the characteristic space where products are defined.

The timing of the game permits us to identify the supply and demand parameters as in traditional supply-demand models for differentiated products (Berry and Haile, 2014). Identification of the fixed cost parameters is hampered by the possibility of multiple Nash equilibrium networks, which prevents us from writing down a well-defined likelihood function. Further, constructing the set of Nash equilibrium networks for a given value of the fixed cost parameters is computationally burdensome, due to the large number of markets and the presence of spillovers in entry. We circumvent these issues by following the literature on revealed preferences (Pakes, 2010; Pakes et al., 2015), and derive moment inequalities from best-response implications. Namely, if the networks chosen by the airlines constitute a pure strategy Nash equilibrium, then the airlines' profits in these networks are higher than those in counterfactual networks. These moment inequalities are easy to evaluate because they require us to neither impose any ad-hoc equilibrium selection assumption, nor to construct the set of equilibrium networks for each possible value of the fixed cost parameters.

Due to our fixed cost specification, the identified set defined by the moment inequalities is a convex polytope. Convexity has long been recognized as a desirable feature in the set identification literature (Beresteanu and Molinari, 2008; Bontemps et al., 2012; Kaido and Santos, 2014). Indeed, convexity often reduces the computational burden of estimation because analysts can directly estimate the frontier points of the identified set via the support function. This allows us to estimate the identified set of the fixed cost parameters using an easy-to-implement procedure based on solving linear programs. However, constructing a confidence region depends on nuisance parameters that cannot be uniformly estimated, as is typical in the set identification literature. By exploiting the linearity of the moment inequalities, we design a method to appropriately smooth the identified set. We show that this smoothing step yields a strictly convex outer set and, hence, guarantees the asymptotic normality of the estimated support function with a variance that can be easily computed from the data. In turn, we can construct confidence intervals for each component (or linear combinations of components) of the vector of fixed cost parameters by solving linear programs with linear and exponential cone constraints.

We estimate our model using US domestic tickets data from the Airline Origin and Destination Survey during the third quarter of 2011. We consider the flights operated by the main airlines between the top 85 US cities. Our empirical findings reveal significant spillovers in entry on the demand, marginal cost, and fixed cost sides. Specifically, on the demand side, consumers benefit from flying with airlines offering many connections out of the itinerary's endpoints due to an increase in the value of loyalty programs. Hence, dense networks increase consumers' willingness-to-pay for an airline's flight. On the supply side, the marginal costs of an itinerary decrease when an airline allows passengers to reach many cities from the itinerary's endpoints and intermediate stops, due to economies of density. Hence, dense networks generate marginal cost savings. At the same time, we find that the denser the networks, the higher the fixed costs of offering direct flights out of hubs due to congestion effects.

Our empirical findings have significant consequences on analyzing environmental changes, such as mergers. In particular, we use our estimates to study the merger between American Airlines and US Airways. These two firms merged in 2013, subject to a series of remedies imposed by the Department of Justice (DoJ) to restrict post-merger network readjustments and protect consumer surplus. We highlight three main counterfactual results.

First, without the remedies, the merger leads to a decrease in consumer surplus by around 0.95%. With the remedies, consumer surplus increases by about 3.09%, suggesting that these measures contribute significantly to preventing losses in consumer surplus.

Second, the impact of the merger differs between the markets that the merging parties served pre-merger ("old markets"), on which antitrust authorities typically focus, and the markets where the merged entity enters post-merger ("new markets"), which are usually ignored by antitrust authorities. On the one hand, old markets undergo consumer surplus losses of around 7.52%. If the merger's effect on consumer surplus in the old markets was the relevant criterion, then the merger should have been blocked. This is in line with the DoJ's initial attempt to stop the merger. On the other hand, new markets experience an increase in consumer surplus by around 68.14%, driven by the high willingness-to-pay for direct flights. It reveals substantial positive effects of the merger and can be used to legitimize its implementation. Further, the DoJ's remedies, which were tailored for old markets, reduce consumer surplus losses in old markets to 2.74% but, at the same time, weaken consumer surplus gains in new markets to 55.17%. This highlights the need for antitrust authorities to carefully balance these two effects when designing post-merger interventions. To the best of our knowledge, this tension between consumer surplus losses in old markets and consumer surplus gains in new markets is a novel empirical finding that has major implications for policymakers, and clearly shows the inadequacy of the fixed network approach in evaluating mergers.

Third, spillovers in entry substantially shape post-merger outcomes. In particular, the differences between new and old markets are driven by the expansion of the American Airlines network in an attempt to leverage spillovers on the demand and marginal cost sides, and the reduction of competitors' networks, which are unable to compete with the new powerful player. Importantly, such network changes align with the real entry-exit patterns observed after 2013. We further quantify the impact of spillovers in entry by running our merger simulation without them. In their absence, the overall impact of the merger changes significantly, leading to a drastic reduction in the merged entity's re-optimized network and leaving consumers substantially worse off.

The rest of the paper is organized as follows. Section 2 summarizes the literature. Section 3 presents the model. Sections 4 and 5 discuss identification and inference. Section 6 presents the data of our empirical application on the US domestic market. Section 7 displays our estimates of the structural model and Section 8 studies the merger between US Airways and American Airlines. Section 9 concludes. Further details are available in the Online Appendix.

2 Literature Review

This paper contributes to a flourishing literature on entry, exit, and product positioning (Mazzeo, 2002; Seim, 2006; Ho, 2009; Holmes, 2011; Eizenberg, 2014; Houde et al., 2023; Kuehn, 2018; Wollmann, 2018; Crawford et al., 2019; Aguirregabiria et al., 2020; Fan and Yang, 2024). Although the two-stage structure of our model is similar to that of Eizenberg (2014) and Wollmann (2018), the spillovers in entry create additional methodological and computational challenges by preventing us from applying a market-by-market analysis. Further, we develop and implement a formal inference procedure for the fixed cost parameters that leverages the convexity of their identified set.

This paper also relates to the recent advances in the econometrics of network formation games (Graham and de Paula, 2020). The methods proposed in this literature typically require the analyst to construct the set of equilibria for each candidate parameter value. However, this is unfeasible in our setting due to the large dimension of the airlines' networks. Our approach demonstrates that exploiting the necessary conditions for equilibrium, along the lines of Pakes (2010) and Pakes et al. (2015), represents an alternative route for constructing identified sets in large network formation games that offers notable computational advantages.

Furthermore, and from a more conceptual standpoint, while the network literature typically focuses on "decentralized" network formation, where each node is controlled by a different agent, this paper is among the first to examine a "centralized" network formation process, in which a single airline ("general manager") determines the link decisions for *all* its nodes (cities). This distinction affects how we characterize the model's equilibrium. Specifically, in the context of our "centralized" model, when evaluating changes in an airline's profits due to a deviation from equilibrium in one market, we must account for any ripple effects that this deviation has on neighboring markets where the airline operates via the spillover variables. These effects must be fully integrated into the airline's profit changes. Conversely, within a "decentralized" network formation process, the neighboring markets are overseen by distinct "local managers", each with its own profit function and potentially competing objectives. As a result, the equilibrium implications differ significantly between the two approaches.

More broadly, this paper combines two strands of the literature. The first strand is on structural models of demand and supply. This literature estimates demand and supply equations, taking entry decisions as exogenously given (Berry, 1994; Berry et al., 1995). For applications of these models to airlines, see Berry et al. (1996), Berry and Jia (2010), Ciliberto and Williams (2014). For applications to merger analysis, see Nevo (2000), Bjornerstedt and Verboven (2016), and Miller and Weinberg (2017). The second strand is the literature on entry models. This literature estimates the payoffs from entering markets, assuming that entry decisions are independent across markets, without considering demand and supply equations. See de Paula (2013) for a review. Applications of entry models to airlines can be found, for instance, in Reiss and Spiller (1989), Berry (1992), Goolsbee and Syverson (2008), and Ciliberto and Tamer (2009).

Last, this paper is among the first to model the airlines' entry decisions by taking endogenous spillovers in entry into account and combining this step with a supply-demand framework. Ciliberto et al. (2021) and Li et al. (2022) develop methods to estimate models of entry and price decisions. Differently from us, they assume that entry decisions are independent across markets and do not endogenize spillovers in entry. However, these papers allow entry decisions to be correlated with both fixed cost shocks and supply-demand shocks, whereas we only allow the first channel of correlation.

Aguirregabiria and Ho (2012), Benkard et al. (2020), Park (2020), and Jia and Yuan (2024) introduce endogenous spillovers in entry in their models for the airline industry. However, differently from our paper, Aguirregabiria and Ho (2012) assume that entry decisions are made by independent local managers, each concerned with maximizing profits for a particular local sub-network of their employing airline, in order to reduce computational complexity. Benkard et al. (2020) investigate the dynamic effects of mergers on the airlines' networks. They focus on medium- to long-run transitions towards the post-merger equilibrium and do not model the demand side. In contrast, our framework includes the demand side and we are interested in the welfare effects emanating from the change between pre- and post-merger equilibrium. Park (2020) endogenizes entry and slot choices at the Ronald Reagan Washington National Airport only. Jia and Yuan (2024) model entry, capacity, and price decisions, but do not consider spillovers in entry on the fixed cost side. Finally, our paper is the first to disentangle and estimate the impact of hub-and-spoke structures, both on the variable profits and fixed costs.

3 The Model

There are N airlines, labeled by $f \in \mathcal{N} := \{1, \ldots, \mathrm{N}\}$, which play a two-stage game. The timing of the game is represented in Figure 1. In the first stage, the airlines design their networks to transfer passengers from one city to another and pay the fixed costs. Cities are connected directly and/or via hubs. Hubs are exogenously pre-determined. In the second stage, given the networks, the airlines face the demand for their products, pay the variable costs, and compete in prices. In the first stage, the airlines observe their own and their competitors' fixed cost shocks. However, they do not observe their own and their competitors' demand and supply shocks, which are discovered before the second stage. Therefore, in the first stage, the airlines choose the networks maximizing their expected second-stage profits minus the fixed costs. In what follows, we describe the game starting from the second stage.

Figure 1: Timing of the game

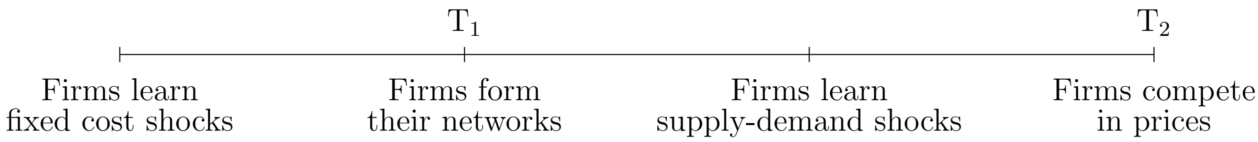

3.1 The Second Stage: Demand and Supply

In the second stage, the airlines take as given the networks and consequent product choices. Markets are non-directional city-pairs, such as Los Angeles-Boston, which allows the possibility to fly from Boston to Los Angeles or from Los Angeles to Boston. Each market is indexed by $m \in \mathcal{M}$, where $\mathcal{M}$ is the set of markets. Alternatively, when we need to keep a record of the endpoint cities, the market whose endpoints are cities a and b is denoted by $\{a, b\} \in \mathcal{M}$. Each product offered in market m is indexed by $j \in \mathcal{J}_m$, where $\mathcal{J}_m$ is the set of products offered in market m.

Products are airline-itinerary combinations. For example, in the Los Angeles-Boston market, American Airlines offers a direct flight between Boston and Los Angeles, a one-stop flight between Los Angeles and Boston with an intermediate stop at the Dallas hub, and a one-stop flight between Los Angeles and Boston with an intermediate stop at the Chicago hub.

In every market, the airlines face the demand for their products, pay the variable costs, and simultaneously choose the prices to maximize the variable profits, under complete information. We now present the demand and supply equations.

Demand We consider the nested logit demand with two nests; one for the airline products, and the other for the outside option of not travelling or travelling by other means (Berry, 1994). The utility that individual i receives from buying product j in market m is specified as:

$$U_{i,j,m} = X_{j,m}^{\top}\beta - \alpha p_{j,m} + \xi_{j,m} + \nu_{i,m}(\lambda) + \lambda\epsilon_{i,j,m}. \tag{1}$$

The outside option is denoted by 0 and its utility normalized to $\epsilon_{i,0,m}$. In (1), $X_{j,m}$ is a vector of product characteristics and $p_{j,m}$ is the product price, both observed by the researcher. $\xi_{j,m}$ represents the product characteristics that are unobserved by the researcher and can be arbitrarily correlated with prices. $(\nu_{i,m}(\lambda), \epsilon_{i,j,m}, \epsilon_{i,0,m})$ denote the consumer tastes, unobserved by the researcher, i.i.d. across i, j, m, and independent of all the other variables. The probability distribution of $(\nu_{i,m}(\lambda), \epsilon_{i,j,m}, \epsilon_{i,0,m})$ is chosen to yield the familiar nested logit market share function, with $\lambda \in (0, 1]$.

We include in $X_{j,m}$ various product characteristics, such as the number of stops and the distance flown, along with carrier and city fixed effects. Further, $X_{j,m}$ contains the number of direct flights offered out of market m's endpoints by the same carrier offering itinerary j (hereafter, "*Nonstop Origin*"). The variable "*Nonstop Origin*" captures the value of frequent flier programs (Berry and Jia, 2010). In fact, the larger the number of destinations for which consumers can redeem frequent flier miles, the higher the value of such loyalty programs, and so the higher the utility consumers can get from flying with a given carrier. Moreover, an airline that flies to many cities is likely to have more convenient parking and gate access and provide better services. See Figure 2 for an example. Note that, due to the variable "*Nonstop Origin*", the demand for product j in market m depends on the entry decisions of an airline in neighboring markets. This gives rise to spillovers in entry across markets on the demand side.

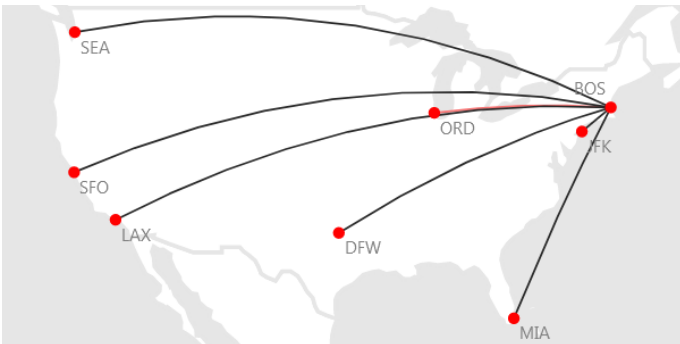

Figure 2: Let market m be Boston-Chicago and product j be a direct flight between Boston and Chicago offered by American Airlines. The larger the number of direct flights offered by American Airlines to passengers from Boston (for instance, direct flights to Chicago as well as New York, Dallas, Los Angeles, Seattle, Miami, and San Francisco), the higher the value of American Airlines' frequent flier programs, and the more facilities American Airlines will provide to customers at Boston's airport. These mechanisms are expected to increase the utility of buying product j.

From utility-maximizing behavior, we obtain the predicted demand in market m.

$$\begin{aligned} &\text{For product } j, && s_{j,m}(X_m, p_m, \xi_m; \theta_d) \times \mathrm{MS}_m = \frac{\exp(\delta_{j,m}/\lambda)}{D_m} \frac{D_m^\lambda}{1 + D_m^\lambda} \times \mathrm{MS}_m. \\ &\text{For the outside option } 0, && s_{0,t}(X_m, p_m, \xi_m; \theta_d) \times \mathrm{MS}_m = \frac{1}{1 + D_m^\lambda} \times \mathrm{MS}_m, \end{aligned} \tag{2}$$

where $X_m := (X_{j,m} : j \in \mathcal{J}_m)$, $p_m := (p_{j,m} : j \in \mathcal{J}_m)$, $\xi_m := (\xi_{j,m} : j \in \mathcal{J}_m)$, $\theta_d := (\beta, \alpha, \lambda)$, $s_{j,m}(X_m, p_m, \xi_m; \theta_d)$ is the product share of product j in market m, MS_m is the market size, $\delta_{j,m} := X_{j,m}^\top \beta - \alpha p_{j,m} + \xi_{j,m}$, and $D_m := \sum_{j=1}^{J_m} \exp(\delta_{j,m}/\lambda)$. The researcher observes the product shares without errors, as is standard in the literature.

Supply The airlines simultaneously set the prices in each market m to maximize the variable profits, under complete information:

$$\sum_{m \in \mathcal{M}} \sum_{j \in \mathcal{J}_{f,m}} (p_{j,m} - \mathrm{MC}_{j,m}) \times s_{j,m}(X_m, p_m, \xi_m; \theta_d) \times \mathrm{MS}_m, \tag{3}$$

where $\mathcal{J}_{f,m}$ is the set of products offered by airline f in market m and $\mathrm{MC}_{j,m}$ is product j's marginal cost. For each airline f and market m, we obtain the Bertrand-Nash F.O.C.s in the usual way:

$$\mathrm{MC}_{f,m} = p_{f,m} + \Big(\frac{\partial s_{f,m}(X_m, p_m, \xi_m; \theta_d)}{\partial p_{f,m}}\Big)^{-1} s_{f,m}(X_m, p_m, \xi_m; \theta_d), \tag{4}$$

where $\mathrm{MC}_{f,m}$, $p_{f,m}$, and $s_{f,m}(X_m, p_m, \xi_m; \theta_d)$ are the vectors stacking $\mathrm{MC}_{j,m}$, $p_{j,m}$, and $s_{j,m}(X_m, p_m, \xi_m; \theta_d)$, respectively, for each product $j \in \mathcal{J}_{f,m}$. $\frac{\partial s_{f,m}(X_m, p_m, \xi_m; \theta_d)}{\partial p_{f,m}}$ is the matrix collecting the partial derivatives of product shares with respect to prices.

As standard in the literature, we express product j's marginal cost as a function of observed and unobserved cost shifters:

$$\mathrm{MC}_{j,m} = W_{j,m}^{\top} \theta_s + \omega_{j,m}, \tag{5}$$

where $W_{j,m}$ is a vector of marginal cost shifters that are observed by the researcher and $\omega_{j,m}$ represents the marginal cost shifters that are unobserved by the researcher.

As for the demand side, we include in $W_{j,m}$ various product characteristics, such as the number of stops, the distance flown, and whether the itinerary is short-haul or long-haul, along with carrier fixed effects. Further, $W_{j,m}$ contains the number of cities that are reachable from the endpoints and intermediate stops of itinerary j with the same carrier offering itinerary j (variable "*Connections*"). This variable captures economies of density; that is, the fact that more densely traveled markets tend to generate marginal cost savings due to engineering reasons (Berry and Jia, 2010). In particular, the larger the number of final destinations consumers can reach via connecting flights, the more the opportunities for an airline to pool passengers from several itineraries into the same planes, and so the more an airline can efficiently use large aircrafts, which tend to have lower unit costs. Note that, due to the variable "*Connections*", the marginal cost of product j in market m depends on the entry decisions of an airline in neighboring markets. This gives rise to spillovers in entry across markets on the marginal cost side.

3.2 The First Stage: Entry

In the first stage, the airlines design the networks to transfer passengers from one city to another and pay the fixed costs. Cities are connected directly and/or via hubs with, at most, one intermediate stop. Our data contain very few observations of flights with more than one intermediate stop and flights connecting via non-hubs (see Section 6 for more details). We assume that hub locations are exogenously determined before game begins. This is because the transition from point-to-point to hub-and-spoke operations was a historical process that was started by the airlines after the US Airline Deregulation Act of 1978 and was quickly completed by the 1990s, many years before the period considered in our empirical application. Once decided upon, hub locations were not altered in any major way by the airlines, even after mergers and other restructuring events. This is also highlighted in Figures 3 to 6 of Section 6.[1]

We formalize the network formation process as follows. Given market $\{a, b\}$, let:

$$G_{ab,f} = \begin{cases} 1 & \text{if airline } f \text{ offers direct flights between cities } a \text{ and } b, \\ 0 & \text{otherwise.} \end{cases}$$

Let $G_f \coloneqq (G_{ab,f} : \{a, b\} \in \mathcal{M})$ be the network of airline f, where the nodes of the network are the cities and the links of the network are the markets served by airline f with direct flights. In the first stage, each airline f chooses its network G_f, that is, the set of markets served by direct flights. This choice automatically determines which markets are served by airline f with one-stop flights. In particular, if $G_{ah,f} = G_{hc,f} = 1$ and city h is one of airline f's hubs, then airline f also competes in market $\{a, c\}$ by offering one-stop flights between cities a and c via h.

When entering markets, the airlines pay the fixed costs of building and maintaining the physical, technological, and human infrastructures. Examples are the costs of salaries, insurance, scheduling coordination, computer reservation and revenue management systems, advertising, and aircraft financing. The fixed costs also include the fees for ticket offices, baggage conveyors, gates, lounges, parking, and hangars at the airports. Further, the literature suggests that hub-and-spoke operations can increase the fixed costs due to the risk of operational congestion at hubs where many connections must be carefully coordinated. For instance, consider the extra resources that need to be invested in order to harmonize flight schedules, the leasing of contiguous gates, and the management of passenger traffic along different parts of the airport, in the case of closely scheduled flights.

[1] See Holmes (2011) for an empirical framework that incorporates supercenter (hub) location choices in a study of the supermarket industry.

Taking these points into account, we specify the fixed costs sustained by airline f as:

$$\mathrm{FC}_f(G_f, \eta_f; \gamma) = \sum_{\{a,b\}\in\mathcal{M}} G_{ab,f}(\gamma_{1,f} + \eta_{ab,f}) + \sum_{h\in\mathcal{H}_f} \gamma_{2,f}\Big(\sum_{\substack{a\,\in\,\mathcal{C}\\ a\,\neq\,h}} G_{ha,f}\Big)^2, \tag{6}$$

where $\mathcal{H}_f$ is the set of airline f's hubs, $\mathcal{C}$ is the set of cities, $\eta_f := (\eta_{ab,f} : \{a,b\} \in \mathcal{M})$ is a vector of mean-zero shocks observed by the airlines but unobserved by the researcher, and $\gamma := (\gamma_{1,f}, \gamma_{2,f} : f \in \mathcal{N})$ collects the parameters to be identified. The fixed cost equation consists of two parts. First, there are market-specific contributions, $\gamma_{1,f}+\eta_{ab,f}$, for each market $\{a,b\}$ served by direct flights. Second, there are quadratic terms, $\gamma_{2,f}(\sum_{\substack{a\,\in\,\mathcal{C}\\ a\,\neq\,h}} G_{ha,f})^2$, for each hub h, that account for the risk of congestion at hubs, as discussed in the previous paragraph. In particular, $\sum_{\substack{a\,\in\,\mathcal{C}\\ a\,\neq\,h}} G_{ah,f}$ is the degree of hub h, that is, the number of markets served out of hub h with direct flights by airline f (also called "spokes"). Due to such quadratic terms, the increase in the fixed costs sustained by airline f when adding a spoke to hub h depends on the number of spokes that hub h already has. This gives rise to spillovers in entry decisions across markets on the fixed cost side.

In Equation (6), we interpret the quadratic term for hub congestion as an operational and internal cost of airline f, unrelated to competition with other firms, which we instead capture in the second stage of the model. Excluding competitive entry effects from the fixed cost equation, such as those arising from slot constraints, is consistent with our approach of not differentiating between airports within the same city. In fact, in our data, there are no hub cities where all airports are slot-constrained at the time of our empirical application. Consequently, the impact of competitors' entries at hubs on airline f's fixed costs is likely diminished. However, our econometric methodology is flexible enough to incorporate such factors if needed. We refrain from doing so to maintain a parsimonious and computationally manageable framework, and more importantly, because we lack the necessary data on slot and gate allocations, capacity, and frequency required to explicitly model slot constraints.

We assume that the fixed cost shocks, $\eta := (\eta_f : f \in \mathcal{N})$, are common knowledge among the airlines. In fact, in the airline industry, the fixed costs capture fairly standard balance sheet entries that pertain to the long-term side of the business and do not typically involve any industrial or technological "secrets". Hence, it is plausible to suppose that the airlines can predict the competitors' fixed cost shocks reasonably well.

We also assume that, in the first stage, the airlines know everything about the second stage, except their own and their competitors' demand and marginal costs shocks. This is a natural assumption, because the legacy carriers (which, together with Southwest Airlines, are the main players of our empirical application) typically operate with a time lag between the entry decisions and the sale of flight tickets.

Therefore, the airlines simultaneously choose the networks $G := (G_f : f \in \mathcal{N})$ that maximize the expected second-stage profits minus the fixed costs:

$$\mathbb{E}[\Pi_f(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \xi^{\oplus}, \omega^{\oplus}, G; \theta) | X^{\oplus}, W^{\oplus}, \mathrm{MS}, \eta] - \mathrm{FC}_f(G_f, \eta_f; \gamma), \tag{7}$$

where $\Pi_f(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \xi^{\oplus}, \omega^{\oplus}, G; \theta)$ is the second-stage profit of airline f. Hereafter, we denote by $\mathcal{J}_m^{\oplus}$ the set of *all potential* products in market m, including the products not chosen for production. In turn, $X^{\oplus} \coloneqq (X_{j,m} : j \in \mathcal{J}_m^{\oplus}, m \in \mathcal{M})$, $W^{\oplus} \coloneqq (W_{j,m} : j \in \mathcal{J}_m^{\oplus}, m \in \mathcal{M})$, $\xi^{\oplus} \coloneqq (\xi_{j,m} : j \in \mathcal{J}_m^{\oplus}, m \in \mathcal{M})$, and $\omega^{\oplus} \coloneqq (\omega_{j,m} : j \in \mathcal{J}_m^{\oplus}, m \in \mathcal{M})$ are the vectors of observed demand shifters, observed marginal cost shifters, demand shocks, and marginal cost shocks of all potential products across all markets.[2] $\mathrm{MS} \coloneqq (\mathrm{MS}_m : m \in \mathcal{M})$ is the vector of market sizes. $\theta \coloneqq (\theta_d, \theta_s)$ is the vector of second-stage parameters. Note that the expectation of the second-stage profits is computed by integrating over the demand and supply shocks, $(\xi^{\oplus}, \omega^{\oplus})$, conditional on the variables observed by the airlines in the first stage, $(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \eta)$. Note also that the second-stage profits depend on the networks formed by the airlines. In fact, these networks determine the products offered and their characteristics, and, given the supply-demand shocks, the equilibrium prices in each market.

3.3 Equilibrium

The airlines solve the game by working backwards from the second stage. First, they calculate the equilibrium profits under any possible networks, demand shocks, and marginal cost shocks. Then, they choose the networks that are maximizing the expected value of those profits. A subgame perfect pure strategy Nash equilibrium consists of networks and price functions, $\{G^*, (P_m^*(\xi_m^{\oplus}, \omega_m^{\oplus}, G) : m \in \mathcal{M})\}$, constituting a pure strategy Nash equilibrium in every subgame.

The existence and uniqueness of $(P_m^*(\xi_m^{\oplus}, \omega_m^{\oplus}, G) : m \in \mathcal{M})$ is established by Nocke and Schutz (2018) for the case of multi-product nested logit, which is what we consider here. We allow for multiple G^*. Multiple G^* are possible because the airlines compete at the entry stage through the second-stage pricing game. As explained in Appendix B, it is difficult to show that at least one G^* exists, due to the presence of spillovers in entry on demand, marginal cost, and fixed cost sides. In what follows, we assume that G^* exists. In Appendix C, we show that the moment inequalities derived in Section 4.2 are robust to the possibility that G^* does not exist for some parameter values and variable realizations.

4 Identification

This section discusses the identification of the vector of parameters, $(\theta, \gamma) \in \Theta \times \Gamma \subseteq \mathbb{R}^{\mathrm{K}} \times \mathbb{R}^{\mathrm{P}}$, where K is the dimension of θ and P is the dimension of γ.

[2]Analogously, we define the market-specific vectors $X_m^{\oplus} \coloneqq (X_{j,m} : j \in \mathcal{J}_m^{\oplus})$, $W_m^{\oplus} \coloneqq (W_{j,m} : j \in \mathcal{J}_m^{\oplus})$, $s_m^{\oplus} \coloneqq (s_{j,m} : j \in \mathcal{J}_m^{\oplus})$, and $P_m^{\oplus} \coloneqq (p_{j,m} : j \in \mathcal{J}_m^{\oplus})$. We will also use this notation in Section 4.1.

4.1 Identification of the Demand and Supply Parameters

To identify $\theta := (\theta_d, \theta_s) \in \Theta$, we follow the identification arguments for standard supply-demand models with differentiated products (Berry and Haile, 2014). Intuitively, the vector of demand parameters, θ_d, is identified from the distribution of prices, sales, and product covariates. Once θ_d is identified, the markups are also identified from the F.O.C.s in (4). In turn, the marginal costs are identified as the difference between the prices and the markups. Last, the variation in the identified marginal costs and product covariates identifies the vector of marginal cost parameters, θ_s.

More precisely, there are two potential sources of endogeneity to be considered here. First, the list of products offered in the second stage is selected by the airlines in the first stage and may be correlated with the supply-demand shocks ("selection on supply-demand shocks"). Second, the prices and within-group market shares are correlated with the supply-demand shocks because the latter are observed by the airlines when playing the second stage (standard endogeneity issue in supply-demand models). We address these two sources of endogeneity by leveraging the timing of information release about the shocks in our model (first, the fixed cost shocks, then the supply-demand shocks) and ruling out any correlation between the fixed cost shocks and the supply-demand shocks. Formally, we assume that

$$\mathbb{E}(\xi_{j,m}, \omega_{j,m} | X^{\oplus}, W^{\oplus}, \mathrm{MS}, \eta) = 0 \quad a.s., \tag{8}$$

for every product $j \in \mathcal{J}_m^{\oplus}$. That is, the information owned by the airlines in the first stage (in particular, knowing the realization of fixed cost shocks) does not help them to predict the supply-demand shocks better.

The mean independence (8) eliminates "selection on supply-demand shocks" because it implies that $\mathbb{E}(\xi_{j,m}, \omega_{j,m}|G) = 0$ for every product $j \in \mathcal{J}_m^{\oplus}$, that is, the supply-demand shocks are mean independent of the list of products offered in the second stage. Further, the mean independence assumption allows us to instrument the prices and within-group market shares in the usual way of supply-demand models, without requiring parametric distributional assumptions on the supply-demand shocks. Specifically, let $z_{j,m}(X_m^{\oplus}, W_m^{\oplus})$ be an $\mathrm{L} \times 1$ vector of instruments pertaining to product $j \in \mathcal{J}_m^{\oplus}$, where $\mathrm{L} \geq \mathrm{K}$. Given $\rho_{j,m} := (\xi_{j,m}, \omega_{j,m})$, the mean independence assumption implies:

$$\mathbb{E}(\rho_{j,m} \times z_{j,m,l}(X_m^{\oplus}, W_m^{\oplus})|G) = 0 \quad \forall l = 1, \dots, \mathrm{L}, \tag{9}$$

for every product $j \in \mathcal{J}_m^{\oplus}$. Berry et al. (1995) show that we can uniquely express $\rho_{j,m}$ as a function of the product covariates and θ ("BLP inversion"):

$$\rho_{j,m} = \tau_{j,m}(X_m^{\oplus}, W_m^{\oplus}, \mathrm{MS}_m, s_m^{\oplus}, P_m^{\oplus}, G; \theta).$$

Therefore, we obtain:

$$\mathbb{E}(\tau_{j,m}(X_m^{\oplus}, W_m^{\oplus}, \mathrm{MS}_m, s_m^{\oplus}, P_m^{\oplus}, G; \theta) \times z_{j,m,l}(X_m^{\oplus}, W_m^{\oplus})|G) = 0 \quad \forall l = 1, \ldots, \mathrm{L}, \tag{10}$$

for every product $j \in \mathcal{J}_m^{\oplus}$. The above moment equalities depend only on variables that are observed by the researcher and guarantee point identification of θ under an appropriate rank condition. Following Berry et al. (1995), we use functions of the observed demand shifters as instruments for the price and the within-group market share. See Footnote 12 in Section 6 for the list of second-stage instruments.

The mean independence (8) is standard in empirical two-stage games (Eizenberg, 2014; Holmes, 2011; Houde et al., 2023; Kuehn, 2018; Wollmann, 2018). For an alternative approach handling "selection on supply-demand shocks", see Ciliberto et al. (2021), which considers a one-stage game where the realization of all shocks is revealed at the start of the game and allows for correlation between the fixed cost shocks and supply-demand shocks, under a parametric specification of the fixed cost shock distribution.

4.2 Identification of the Fixed Cost Parameters

This section shows how $\gamma \in \Gamma$ is set-identified from moment inequalities. We follow the literature on revealed preferences (Pakes, 2010; Pakes et al., 2015) and derive these moment inequalities from best-response implications. Namely, if the networks chosen by the airlines constitute a pure strategy Nash equilibrium (PSNE), then they should lead to higher profits than if the airlines were to deviate from those networks.

This strategy allows us to circumvent two potential difficulties. First, there may be multiple PSNE networks. Second, constructing the set of PSNE networks is computationally burdensome due to the large number of markets and the presence of spillovers in entry. Our moment inequalities do not require us to impose any ad-hoc equilibrium selection assumption and construct the set of PSNE networks for each candidate parameter value.

Best-response Implications

Let G_{-f} denote the networks formed by firm f's competitors and $G \coloneqq (G_f, G_{-f})$. If G_f is part of a PSNE, as we assume, then it should lie on airline f's best-response curve. In turn, the increase in profits that airline f would receive if it deviated from G_f and the other firms were mandated to keep G_{-f} is negative. To keep the computational burden manageable, we consider one-link deviations only; that is, each airline f can add/delete direct flights in one market at a time. Restricting our attention to one-link deviations does not introduce any misspecification, as one-link deviations are a subset of all possible deviations airlines guard against when choosing their best responses. If anything, this may result in an identified set that is theoretically not sharp, and we discuss this further at the end of this section.

For each airline f, take market $\{a, b\}$ that airline f does not serve with direct flights in equilibrium; that is, for which $G_{ab,f} = 0$. Let airline f's counterfactual network be the network in which airline f operates in all markets served under G_f and, additionally, it offers direct flights between cities a and b. We denote it by $G_{(+ab),f}$. From the revealed preference principle, the difference between the expected variable profits earned by airline f under its counterfactual network $G_{(+ab),f}$ and the expected variable profits earned by airline f under its factual network G_f, while the competitors maintain G_{-f}, must be less than or equal to the extra fixed-cost that airline f pays for the added direct flight:

$$\Delta\Pi_{(+ab),f} \leq \Delta\mathrm{FC}_{(+ab),f}, \tag{11}$$

where $\Delta\Pi_{(+ab),f}$ is the difference between the expected variable profits at $G_{(+ab),f}$ and G_f and $\Delta\mathrm{FC}_{(+ab),f}$ is the difference between the fixed costs at $G_{(+ab),f}$ and G_f.[3] Following Equation (6), the right-hand side of (11) is a linear expression in $\gamma_{1,f}$, $\gamma_{2,f}$ and $\eta_{ab,f}$. Denoting by $\Delta\mathrm{Q}_{(+ab),f}$ the difference between the quadratic terms in Equation (6) at $G_{(+ab),f}$ and G_f, we have:

$$\Delta\mathrm{FC}_{(+ab),f} = \gamma_{1,f} + \gamma_{2,f}\Delta\mathrm{Q}_{(+ab),f} + \eta_{ab,f}.$$

When neither a nor b is a hub of airline f, $\Delta\mathrm{FC}_{(+ab),f} = \gamma_{1,f} + \eta_{ab,f}$. When only a is a hub, $\Delta\mathrm{FC}_{(+ab),f} = \gamma_{1,f} + \gamma_{2,f}(2N_{a,f} + 1) + \eta_{ab,f}$, where $N_{a,f}$ is the number of spokes of hub a in G_f. When only b is a hub, $\Delta\mathrm{FC}_{(+ab),f} = \gamma_{1,f} + \gamma_{2,f}(2N_{b,f} + 1) + \eta_{ab,f}$. As we expect $\gamma_{1,f}$ and $\gamma_{2,f}$ to be positive, (11) provides a "lower" bound for $\gamma_{1,f}$ and $\gamma_{2,f}$.

Similarly, for each airline f, take market $\{a, b\}$ that airline f serves with direct flights in equilibrium; that is, for which $G_{ab,f} = 1$. Let airline f's counterfactual network be the network in which airline f operates in all markets served under G_f, but does no longer offer direct flights between cities a an b. We denote it by $G_{(-ab),f}$. Following the revealed preference principle, we have:

$$\Delta\Pi_{(-ab),f} \geq \Delta\mathrm{FC}_{(-ab),f}, \tag{12}$$

where $\Delta\Pi_{(-ab),f}$ is the difference between the expected variable profits at G_f and $G_{(-ab),f}$ and $\Delta\mathrm{FC}_{(-ab),f}$ is the difference between the fixed costs at G_f and $G_{(-ab),f}$.[4] As above, the

[3]More precisely, based on the notation used in Equations (6) and (7),

$$\begin{aligned}
\Delta\Pi_{(+ab),f} &\coloneqq \mathbb{E}[\Pi_f(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \xi^{\oplus}, \omega^{\oplus}, G_{(+ab),f}, G_{-f}; \theta)|X^{\oplus}, W^{\oplus}, \mathrm{MS}] \\
&\quad - \mathbb{E}[\Pi_f(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \xi^{\oplus}, \omega^{\oplus}, G_f, G_{-f}; \theta)|X^{\oplus}, W^{\oplus}, \mathrm{MS}], \\
\Delta\mathrm{FC}_{(+ab),f} &\coloneqq \mathrm{FC}_f(G_{(+ab),f}, \eta_f; \gamma) - \mathrm{FC}_f(G_f, \eta_f; \gamma).
\end{aligned}$$

[4]More precisely, based on the notation used in Equations (6) and (7),

$$\begin{aligned}
\Delta\Pi_{(-ab),f} &\coloneqq \mathbb{E}[\Pi_f(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \xi^{\oplus}, \omega^{\oplus}, G_f, G_{-f}; \theta)|X^{\oplus}, W^{\oplus}, \mathrm{MS}] \\
&\quad - \mathbb{E}[\Pi_f(X^{\oplus}, W^{\oplus}, \mathrm{MS}, \xi^{\oplus}, \omega^{\oplus}, G_{(-ab),f}, G_{-f}; \theta)|X^{\oplus}, W^{\oplus}, \mathrm{MS}], \\
\Delta\mathrm{FC}_{(-ab),f} &\coloneqq \mathrm{FC}_f(G_f, \eta_f; \gamma) - \mathrm{FC}_f(G_{(-ab),f}, \eta_f; \gamma).
\end{aligned}$$

right-hand side of (12) can be written as:

$$\Delta \mathrm{FC}_{(-ab),f} = \gamma_{1,f} + \gamma_{2,f}\Delta \mathrm{Q}_{(-ab),f} + \eta_{ab,f},$$

where $\Delta \mathrm{Q}_{(-ab),f}$ is the difference between the quadratic terms in (6) at G_f and $G_{(-ab),f}$. Again, when neither a nor b is a hub of airline f, $\Delta \mathrm{FC}_{(-ab),f} = \gamma_{1,f} + \eta_{ab,f}$. When only a is a hub, $\Delta \mathrm{FC}_{(-ab),f} = \gamma_{1,f} + \gamma_{2,f}(2N_{a,f} - 1) + \eta_{ab,f}$. When only b is a hub, $\mathrm{FC}_{(+ab),f} = \gamma_{1,f} + \gamma_{2,f}(2N_{b,f} - 1) + \eta_{ab,f}$. As we expect $\gamma_{1,f}$ and $\gamma_{2,f}$ to be positive, (12) provides an "upper" bound for $\gamma_{1,f}$ and $\gamma_{2,f}$.

Before proceeding, we remark that the inequalities in (11) and (12) are also compatible with equilibrium notions weaker than PSNE. For instance, they resemble the notion of pairwise stability used in network theory, as discussed in Appendix B. Further, Appendix D outlines the steps to calculate $\Delta\Pi_{(+ab),f}$ and $\Delta\Pi_{(+ab),f}$ in (11) and (12). Note here that, despite considering one-link deviations, these quantities are not computed as if entry decisions were independent across markets or as if the network formation process were decentralized at the level of city local managers. In fact, a one-link deviation generates payoff-relevant ripple effects in neighboring markets due to the creation of new products and changes in the characteristics of existing products of airline f. For example, suppose a is a hub and airline f introduces a previously non-existent direct flight between cities a and b, that is, it adds a new product in market $\{a, b\}$. This new route creates a one-stop flight via a between b and any city c directly connected to a, effectively adding a new product in all markets $\{b, c\}$ such that $G_{ac,f} = 1$. Additionally, the characteristics of existing products offered by airline f in neighboring markets change due to the spillover variables. In particular, the fixed costs of all direct flights to/from hub a (which are products of markets $\{a, c\}$ for each c such that $G_{ac,f} = 1$) increase due to the congestion effects at hub a that are potentially exacerbated by the new direct flight between a and b. We must also update the demand attractiveness of these flights due to the "*Nonstop Origin*" spillover variable on the demand side. Lastly, the marginal costs of all flights originating from or destined to a or b change due to the "*Connections*" spillover variable on the marginal cost side. This makes our method different from the approaches that assume that entry decisions are independent across markets or that the network formation process is decentralised.

From Best-response Implications to Moment Inequalities

In this section, we use the inequalities in (11) and (12) to derive moment inequalities. As a first step, we take the expectation of (11) and (12) over *all* markets $\{a, b\}$ and obtain the following moment inequalities for each airline f:

$$\mathbb{E}\left[\Delta\Pi_{(+ab),f}\right] \le \mathbb{E}\left[\Delta\text{FC}_{(+ab),f}\right],$$
$$\mathbb{E}\left[\Delta\Pi_{(-ab),f}\right] \ge \mathbb{E}\left[\Delta\text{FC}_{(-ab),f}\right].$$

Using $\mathbb{E}(\eta_{ab,f}) = 0$, these simplify to

$$\begin{aligned}\mathbb{E}\left[\Delta\Pi_{(+ab),f}\right] &\le \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta\text{Q}_{(+ab),f}\right],\\ \mathbb{E}\left[\Delta\Pi_{(-ab),f}\right] &\ge \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta\text{Q}_{(-ab),f}\right].\end{aligned} \tag{13}$$

Next, we think about how to practically use (13) to partially identify γ. Calculating $\mathbb{E}\left[\Delta\text{Q}_{(+ab),f}\right]$ and $\mathbb{E}\left[\Delta\text{Q}_{(-ab),f}\right]$ is straightforward based on simple algebra. However, when we attempt to compute $\mathbb{E}\left[\Delta\Pi_{(+ab),f}\right]$ and $\mathbb{E}\left[\Delta\Pi_{(-ab),f}\right]$, we face a fundamental challenge. To see why, we apply the law of iterated expectations to $\mathbb{E}\left[\Delta\Pi_{(+ab),f}\right]$ and $\mathbb{E}\left[\Delta\Pi_{(-ab),f}\right]$:

$$\begin{aligned}\mathbb{E}\left[\Delta\Pi_{(+ab),f}\right] &= \mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=0\right]\Pr(G_{ab,f}=0) + \mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=1\right]\Pr(G_{ab,f}=1),\\ \mathbb{E}\left[\Delta\Pi_{(-ab),f}\right] &= \mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=0\right]\Pr(G_{ab,f}=0) + \mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=1\right]\Pr(G_{ab,f}=1).\end{aligned} \tag{14}$$

In (14), $\mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=0\right]$ and $\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=1\right]$ *cannot* be computed. This issue mirrors the *selection problem* in the treatment effect literature (Manski, 2003). The sampling process (combined with the second-stage estimates) reveals the average change in expected variable profits from ceasing service in a served market (analogue of "expected potential outcome of treatment") conditional on being a served markets (analogue of "conditional on being treated"), that is $\mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=1\right]$. Likewise, it reveals the average change in profits from serving an unserved market conditional on being an unserved market, that is $\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=0\right]$. The sampling process does *not* provide information on the average change in profits from ceasing service in a served market conditional on being an unserved market, $\mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=0\right]$, denoted $\mathcal{U}$ below (analogue of "conditional on being untreated"), and from serving an unserved market conditional on being a served market, $\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=1\right]$, denoted $\mathcal{L}$ below. Therefore, without additional assumptions, the identified set of γ is simply the set of γ values such that:

$$\begin{aligned}\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=0\right]\Pr(G_{ab,f}=0) + \mathcal{L}\cdot\Pr(G_{ab,f}=1) &\le \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta\text{Q}_{(+ab),f}\right],\\ \mathcal{U}\cdot\Pr(G_{ab,f}=0) + \mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=1\right]\Pr(G_{ab,f}=1) &\ge \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta\text{Q}_{(-ab),f}\right],\end{aligned} \tag{15}$$

for some real numbers $\mathcal{L} \in \mathbb{R}$ and $\mathcal{U} \in \mathbb{R}$.

Obviously, the size of the identified set depends crucially on the assumptions made on $\mathcal{L}$ and $\mathcal{U}$. For example, inspired by Eizenberg (2014), one could adopt a "worst-case" approach and assume that:

$$
\begin{aligned}
\mathcal{L} \geq \underline{\mathcal{L}} &\coloneqq \inf \mathcal{S}_0(\Delta\Pi_{(+ab),f}) > -\infty, \\
\mathcal{U} \leq \overline{\mathcal{U}} &\coloneqq \sup \mathcal{S}_1(\Delta\Pi_{(-ab),f}) < \infty,
\end{aligned} \tag{16}
$$

where $\mathcal{S}_0(\Delta\Pi_{(+ab),f})$ is the support of $\Delta\Pi_{(+ab),f}$ for the unserved markets and $\mathcal{S}_1(\Delta\Pi_{(-ab),f})$ is the support of $\Delta\Pi_{(-ab),f}$ for the served markets, both of which are known, as explained in the above paragraph. (16) imposes that the average change in expected variable profits from serving unserved markets for the served markets must be at least equal to the smallest change in expected variable profits that is registered for the unserved markets. Similarly, the average change in expected variable profits from ceasing service in served markets for the unserved markets must be at most equal to the biggest change in expected variable profits that is registered for the served markets. The more heterogeneous the markets are, the lower $\underline{\mathcal{L}}$ and the higher $\overline{\mathcal{U}}$ tend to be, making the conditions in (16) less binding.

Based on (16), the moment inequalities in (15) imply:

$$
\begin{aligned}
&\mathbb{E}\left[\Delta\Pi_{(+ab),f} | G_{ab,f} = 0\right] \Pr(G_{ab,f} = 0) + \underline{\mathcal{L}} \Pr(G_{ab,f} = 1) \leq \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(+ab),f}\right], \\
&\overline{\mathcal{U}} \Pr(G_{ab,f} = 0) + \mathbb{E}\left[\Delta\Pi_{(-ab),f} | G_{ab,f} = 1\right] \Pr(G_{ab,f} = 1) \geq \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(-ab),f}\right],
\end{aligned} \tag{17}
$$

which can be brought to data.[5]

In our analysis, we extend the solution proposed by Eizenberg (2014) by incorporating the empirical information from exogenous observed covariates. Specifically, we exploit the availability of two (or more) exogenous binary variables, $Z_{(+ab),f}$ (for the unserved markets) and $Z_{(-ab),f}$ (for the served markets), which may vary by both market and firm, and which satisfy a standard exogeneity condition:

$$
\mathbb{E}\big(\eta_{ab,f} \mid Z_{(+ab),f} = 1\big) = 0 \quad \text{and} \quad \mathbb{E}\big(\eta_{ab,f} \mid Z_{(-ab),f} = 1\big) = 0.
$$

These "instruments" can be covariates included in the fixed cost equation, variables constructed from exogenous demand or marginal cost shifters, or external factors that influence product choice without affecting demand, marginal costs, or fixed costs. By conditioning

[5] Although the discussion in Eizenberg (2014) is framed slightly differently, our explanation leads to the same moment inequalities. Specifically, expression (13) of Eizenberg (2014) translates into our notation as follows:

$$
\begin{aligned}
&\mathbb{E}[\Delta\Pi_{(+ab),f}\mathbb{1}\{G_{ab,f} = 0\} + \underline{\mathcal{L}}\mathbb{1}\{G_{ab,f} = 1\}] \leq \gamma_{1,f} + \gamma_{2,f}\mathbb{E}[\Delta Q_{(+ab),f}], \\
&\mathbb{E}[\Delta\Pi_{(-ab),f}\mathbb{1}\{G_{ab,f} = 1\} + \overline{\mathcal{U}}\mathbb{1}\{G_{ab,f} = 0\}] \geq \gamma_{1,f} + \gamma_{2,f}\mathbb{E}[\Delta Q_{(-ab),f}],
\end{aligned}
$$

which is equivalent to our (17).

on $Z_{(+ab),f} = 1$ and $Z_{(-ab),f} = 1$, respectively, the moment inequalities in (13) imply that:

$$\begin{aligned}
&\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=0, Z_{(+ab),f}=1\right]\Pr(G_{ab,f}=0 \mid Z_{(+ab),f}=1)\\
&\qquad\qquad + \mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=1, Z_{(+ab),f}=1\right]\Pr(G_{ab,f}=1 \mid Z_{(+ab),f}=1)\\
&\qquad\qquad\qquad\qquad \le \gamma_{1,f}+\gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(+ab),f}|Z_{(+ab),f}=1\right],\\
&\mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=1, Z_{(-ab),f}=1\right]\Pr(G_{ab,f}=1 \mid Z_{(-ab),f}=1)\\
&\qquad\qquad + \mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=0, Z_{(-ab),f}=1\right]\Pr(G_{ab,f}=0 \mid Z_{(-ab),f}=1)\\
&\qquad\qquad\qquad\qquad \ge \gamma_{1,f}+\gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(-ab),f}|Z_{(-ab),f}=1\right].
\end{aligned}$$

Moreover, like Eizenberg (2014), we impose lower and upper bounds on the unknown counterfactual average changes in profits. Namely, we assume that:

$$\begin{aligned}
\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=1, Z_{(+ab),f}=1\right] &\ge \underline{\mathcal{L}}^{+} := \inf \mathcal{S}_0^{+}(\Delta\Pi_{(+ab),f}) > -\infty,\\
\mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=0, Z_{(-ab),f}=1\right] &\le \overline{\mathcal{U}}^{-} := \sup \mathcal{S}_1^{-}(\Delta\Pi_{(-ab),f}) < \infty,
\end{aligned} \tag{18}$$

where $\mathcal{S}_0^{+}(\Delta\Pi_{(+ab),f})$ is the support of $\Delta\Pi_{(+ab),f}$ for the unserved markets such that $Z_{(+ab),f} = 1$ and $\mathcal{S}_1^{-}(\Delta\Pi_{(-ab),f})$ is the support of $\Delta\Pi_{(-ab),f}$ for the served markets such that $Z_{(-ab),f} = 1$. Therefore, the moment inequalities that we finally use are:

$$\begin{aligned}
&\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f}=0, Z_{(+ab),f}=1\right]\Pr(G_{ab,f}=0|Z_{(+ab),f}=1) + \underline{\mathcal{L}}^{+}\Pr(G_{ab,f}=1|Z_{(+ab),f}=1)\\
&\qquad\qquad\qquad\qquad\qquad\qquad \le \gamma_{1,f}+\gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(+ab),f}|Z_{(+ab),f}=1\right],\\
&\overline{\mathcal{U}}^{-}\Pr(G_{ab,f}=0|Z_{(-ab),f}=1) + \mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f}=1, Z_{(-ab),f}=1\right]\Pr(G_{ab,f}=1|Z_{(-ab),f}=1)\\
&\qquad\qquad\qquad\qquad\qquad\qquad \ge \gamma_{1,f}+\gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(-ab),f}|Z_{(-ab),f}=1\right].
\end{aligned} \tag{19}$$

Note that the approach of Wollmann (2018) is a special case of (19) under the additional assumption

$$\Pr(G_{ab,f}=0 \mid Z_{(+ab),f}=1)=1 \quad \text{and} \quad \Pr(G_{ab,f}=1 \mid Z_{(-ab),f}=1)=1. \tag{20}$$

These conditions require $Z_{(+ab),f}$ and $Z_{(-ab),f}$ to be exogenous variables that select markets where offering (or not offering) direct flights is a *certain* event. In essence, (20) sets the weight assigned to the unknown counterfactual average changes in profits to zero, removing the need to compute those quantities. Moreover, while (20) can be easily verified in the data, if these probabilities are incorrectly assumed to be 1 when they are actually below 1, it leads to a misspecification of the identified set. Our approach avoids this risk. Specifically, the more accurately $Z_{(-ab),f}$ and $Z_{(+ab),f}$ select which markets are served or not—specifically, the closer $\Pr(G_{ab,f} = 1|Z_{(-ab),f} = 1)$ and $\Pr(G_{ab,f} = 0|Z_{(+ab),f} = 1)$ are to one—the more closely the moment inequalities in (19) mirror the moment inequalities in Wollmann (2018), making conditions in (18) less relevant. Conversely, the less accurately $Z_{(-ab),f}$ and $Z_{(+ab),f}$ select which markets are served or not—specifically, the closer

$\Pr(G_{ab,f} = 1|Z_{(-ab),f} = 1)$ and $\Pr(G_{ab,f} = 0|Z_{(+ab),f} = 1)$ are to zero—the more the bounds on γ potentially widen. This widening occurs because $\underline{\mathcal{L}}^+$ (resp. $\overline{\mathcal{U}}^-$) defined in (18) is, by definition, weakly smaller (resp. greater) than $\mathbb{E}\left[\Delta\Pi_{(+ab),f} \mid G_{ab,f} = 0, Z_{(+ab),f} = 1\right]$ (resp. $\mathbb{E}\left[\Delta\Pi_{(-ab),f} \mid G_{ab,f} = 1, Z_{(-ab),f} = 1\right]$, and is assigned an increasingly larger weight.

In practice, we consider a few instruments for both the served and unserved markets and denote them as $Z_{(+ab),f,r}$, $r = 1, \ldots, \mathrm{R}^+$ and $Z_{(-ab),f,r}$ for $r = 1, \ldots, \mathrm{R}^-$. Therefore, the identified set of γ is:

$$
\begin{aligned}
\Gamma_I \coloneqq \Big\{&\gamma \in \Gamma, \forall f \in \mathcal{N}, \\
&\mathbb{E}\left[\Delta\Pi_{(+ab),f}|G_{ab,f} = 0, Z_{(+ab),f,r} = 1\right] \Pr(G_{ab,f} = 0|Z_{(+ab),f,r} = 1) \\
&\qquad + \underline{\mathcal{L}}_r^+ \Pr(G_{ab,f} = 1|Z_{(+ab),f,r} = 1) \leq \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(+ab),f}|Z_{(+ab),f,r} = 1\right], \\
&\text{for } r = 1, \ldots, \mathrm{R}^+, \\
&\mathbb{E}\left[\Delta\Pi_{(-ab),f}|G_{ab,f} = 1, Z_{(-ab),f,r} = 1\right] \Pr(G_{ab,f} = 1|Z_{(-ab),f,r} = 1) \\
&\qquad + \overline{\mathcal{U}}_r^- \Pr(G_{ab,f} = 0|Z_{(-ab),f,r} = 1) \geq \gamma_{1,f} + \gamma_{2,f}\mathbb{E}\left[\Delta \mathrm{Q}_{(-ab),f}|Z_{(-ab),f,r} = 1\right], \\
&\text{for } r = 1, \ldots, \mathrm{R}^- \Big\},
\end{aligned}
\tag{21}
$$

where $\underline{\mathcal{L}}_r^+$ (resp. $\overline{\mathcal{U}}_r^-$) is defined as in (18) with instrument $Z_{(+ab),f,r}$ (resp. $Z_{(-ab),f,r}$).

It is reasonable to wonder whether there are any disadvantages to using as many available exogenous variables or combinations of variables as possible to construct first-stage instruments. The goal would be to generate a large number of instruments and, consequently, more moment inequalities, potentially leading to a tighter identified set. However, using too many instruments risks having in the data only an extremely small number of markets selected for each instrument, which leads to asymptotic results for inference becoming implausible approximations. For this reason, we aim to be as parsimonious and judicious as possible in the number of instruments used while still obtaining informative empirical results. See Table 3 in Section 7.2 for the list of first-stage instruments used, and Table 4 for details on their capacity to predict serving and unserving decisions.

Γ_I is not sharp; that is, it may be a superset of the set of observationally equivalent parameters. This is due to three reasons. First, we derive our moment inequalities from best-response implications, that is, necessary conditions for PSNE. Given the possibility of multiple equilibria, these inequalities do not necessarily exploit all the information embedded in the model assumptions; specifically, they do not capture sufficient conditions for PSNE (Kline et al., 2021). Second, our procedure may neglect other valid first-stage instruments. Third, the moment inequalities in (19) are derived from one-link deviations. However, the airlines may also deviate by deleting/adding more than one link at a time. Although our framework can accommodate such multi-link deviations, doing so is computationally burdensome. Nonetheless, in Appendix E we highlight both theoretically and empirically that several classes of multi-link deviations do not significantly tighten the resulting bounds.

Last, in addition to η, Pakes (2010) and Pakes et al. (2015) suggest including in the revealed-preference inequalities some additive perturbations to account for the fact that the model may be mis-specified. We have not included these terms, but we check for model mis-specification when conducting inference using Stoye (2021).

5 Inference on the Fixed Cost Parameters

Inference on the second-stage parameter, θ, is standard and can be conducted with the theory of the Generalized Method of Moments. More details are provided in Appendix F. In this section, we discuss inference on the first-stage parameters, γ. For simplicity of exposition, in what follows, we assume that θ is known. In Appendix G.3, we show how we incorporate the uncertainty arising from the second-stage estimates.

Observe that the definition of the identified set in (21) can be rewritten equivalently:

$$\begin{aligned}\Gamma_I \coloneqq \Big\{&\gamma \in \Gamma, \forall f \in \mathcal{N},\\ &\mathbb{E}\left[\Delta\Pi_{(+ab),f}(1-G_{ab,f})Z_{(+ab),f,r} + \underline{\mathcal{L}}_r^+ G_{ab,f} Z_{(+ab),f,r}\right]\\ &\qquad\qquad \leq \gamma_{1,f}\mathbb{E}\left[Z_{(+ab),f,r}\right] + \gamma_{2,f}\mathbb{E}\left[\Delta\mathrm{Q}_{(+ab),f} Z_{(+ab),f,r}\right], \text{ for } r = 1, \ldots, \mathrm{R}^+,\\ &\mathbb{E}\left[\Delta\Pi_{(-ab),f} G_{ab,f} Z_{(-ab),f,r} + \overline{\mathcal{U}}_r^-(1-G_{ab,f})Z_{(-ab),f,r}\right]\\ &\qquad\qquad \geq \gamma_{1,f}\mathbb{E}\left[Z_{(-ab),f,r}\right] + \gamma_{2,f}\mathbb{E}\left[\Delta\mathrm{Q}_{(-ab),f} Z_{(-ab),f,r}\right], \text{ for } r = 1, \ldots, \mathrm{R}^-\Big\}.\end{aligned} \tag{22}$$

Further, we streamline the notation of the moment inequalities in (22) as:

$$\mathbb{E}(Z_{m,r}B_m)^\top \gamma - \mathbb{E}(Z_{m,r}A_m) \leq 0, \quad r = 1, \ldots, \mathrm{R}, \tag{23}$$

where r indexes a generic instrument, m is a market $\{a,b\}$, R is the total number of instruments across all firms. Therefore, when r is such that $Z_{m,r} = Z_{(+ab),f,r}$ for some firm f and some market $\{a,b\}$, A_m is equal to $-\left(\Delta\Pi_{(+ab),f}(1-G_{ab,f}) + \underline{\mathcal{L}}_r^+ G_{ab,f}\right)$ and B_m is a vector such that $B_m^\top\gamma = -\gamma_{1,f} - \Delta\mathrm{Q}_{(+ab),f}$—corresponding to the second set of inequalities in (22). Furthermore, when r is such that $Z_{m,r} = Z_{(-ab),f,r}$ for some firm f and some market $\{a,b\}$, A_m is equal to $\Delta\Pi_{(-ab),f}G_{ab,f} + \overline{\mathcal{U}}_r^-(1-G_{ab,f})$ and B_m is a vector composed by 1, $\Delta\mathrm{Q}_{(-ab),f}$ and 0's such that $B_m^\top\gamma = \gamma_{1,f} + \Delta\mathrm{Q}_{(-ab),f}\gamma_{2,f}$—corresponding to the first set of inequalities in (22).

The Support Function

Γ_I is a convex polyhedron because it is defined by linear moment inequalities. It is bounded if the differences in expected variable profits between the realized and counterfactual networks are bounded. It is non-empty if the model is correctly specified. Throughout this section, we assume that Γ_I is bounded and has a non-empty interior.

Convexity has been shown to be a desirable property in the set identification literature (Beresteanu and Molinari, 2008; Bontemps et al., 2012; Kaido and Santos, 2014). It often reduces the computational burden of estimation because, rather than inverting a test to estimate the identified set, one can focus on estimating its support function. In particular, the support function of Γ_I, $\delta(q; \Gamma_I)$, in a given direction q, is equal to the (signed) distance from the origin to the supporting hyperplane of Γ_I with outer normal q.

The support function gathers all the moment inequalities satisfied by the model because γ belongs to Γ_I if and only if $q^\top \gamma \leq \delta(q; \Gamma_I)$ for each $q \in \mathbb{R}^{\mathrm{P}}$. Moreover, inference on a subvector of γ can be easily performed by considering specific directions. For example, if the chosen direction, q, has its p-th component equal to 1 (resp., -1) and the other components equal to 0, then the support function of Γ_I in direction q is equal to the maximum (resp., minus the minimum) value of the p-th component of γ.

Due to the linearity of the moment inequalities in γ, the support function of Γ_I in direction q can be calculated from a linear program:

$$\begin{aligned} \delta(q; \Gamma_I) &\coloneqq \sup_{\gamma \in \Gamma} q^\top \gamma, \\ &\text{s.t. } \mathbb{E}(Z_{m,r} B_m)^\top \gamma - \mathbb{E}(Z_{m,r} A_m) \leq 0, \quad r = 1, \ldots, \mathrm{R}. \end{aligned} \tag{24}$$

$\delta(q; \Gamma_I)$ can be estimated after replacing the expectations in (24) with sample averages. In particular, let the estimated identified set be defined as:

$$\widehat{\Gamma}_I \coloneqq \left\{ \gamma \in \Gamma : \left(\frac{1}{\mathrm{M}} \sum_{m=1}^{\mathrm{M}} Z_{m,r} B_m^\top \right) \gamma \leq \frac{1}{\mathrm{M}} \sum_{m=1}^{\mathrm{M}} Z_{m,r} A_m \text{ for } r = 1, \ldots, \mathrm{R} \right\}. \tag{25}$$

The estimated support function in direction q is the support function of the estimated identified set:

$$\hat{\delta}(q; \Gamma_I) \coloneqq \delta(q; \widehat{\Gamma}_I).$$

Again, $\hat{\delta}(q; \Gamma_I)$ can be calculated from a linear program.

Asymptotic Distribution of the Estimated Support Function

In our setting, the random sample of observations,

$$\{Z_{1,m}, \ldots, Z_{\mathrm{R},m}, A_m, B_m\}_{m=1}^{\mathrm{M}},$$

is identically distributed, M being the number of markets. For each $r = 1, \ldots, \mathrm{R}$, let $W_r(\gamma)$ be the limit in distribution of:

$$\sqrt{\mathrm{M}} \left(\frac{1}{\mathrm{M}} \sum_{m=1}^{\mathrm{M}} (Z_{m,r} B_m^\top \gamma - Z_{m,r} A_m) - (\mathbb{E}(Z_{m,r} B_m)^\top \gamma - \mathbb{E}(Z_{m,r} A_m)) \right).$$

In particular, $W_r(\gamma)$ is a univariate centered normal variable whose variance is the variance of $\frac{1}{\mathrm{M}}\sum_{m=1}^{\mathrm{M}}(Z_{m,r}B_m^\top\gamma - Z_{m,r}A_m)$. Estimating this variance requires accounting for the fact that the data are not mutually independent across markets. Indeed, the observed product characteristics offered by a given airline in different markets are correlated due to spillover variables, which in turn induce correlation among the terms $\{Z_{m,r}B_m^\top\gamma - Z_{m,r}A_m\}_m$. To address this correlation, we use a spatial heteroskedasticity and autocorrelation consistent (SHAC) estimator, as explained in Appendix G.2.

We now introduce some notation that is useful for the next arguments. The Lagrangian of (24) is equal to:

$$L(\gamma, \lambda_1, \dots, \lambda_{\mathrm{R}}) := q^\top\gamma + \sum_{r=1}^{\mathrm{R}} \lambda_r \left(\mathbb{E}(Z_{m,r}B_m)^\top\gamma - \mathbb{E}(Z_{m,r}A_m)\right),$$

where $\lambda = (\lambda_1, \dots, \lambda_{\mathrm{R}})^\top$ is the vector of Lagrange multipliers associated with the R inequality constraints. We denote by $\mathcal{G}_0(q)$ the set of optimal solutions γ, and by $\mathcal{L}_0(q)$ the set (possibly a singleton) of vectors of Lagrange multipliers, $\lambda = (\lambda_1, \dots, \lambda_{\mathrm{R}}^\top)$. Observe that both sets depend on the direction q.

Theorem 1 provides the asymptotic distribution of $\hat{\delta}(q;\Gamma_I)$ in any direction q.

Theorem 1. Assume that the moments of order $2+\tau$ of the random variables exist for some $\tau > 0$. Then:

(i) The estimated support function, $\hat{\delta}(q;\Gamma_I)$, converges in probability to the true support function, $\delta(q;\Gamma_I)$, uniformly in q in the unit ball;

(ii) It holds that, uniformly in q in the unit ball:

$$\sqrt{\mathrm{M}}\left(\hat{\delta}(q;\Gamma_I) - \delta(q;\Gamma_I)\right) \xrightarrow[\mathrm{M}\to\infty]{d} \mathbb{G}(q) := \sup_{\gamma\in\mathcal{G}_0(q)} \sup_{\lambda\in\mathcal{L}_0(q)} \sum_{r=1}^{\mathrm{R}} \lambda_r W_r(\gamma).$$

Moreover,

$$\sup_{q,\ \|q\|\le 1} \sqrt{\mathrm{M}}\left(\hat{\delta}(q;\Gamma_I) - \delta(q;\Gamma_I)\right) - \mathbb{G}(q) \xrightarrow[\mathrm{M}\to\infty]{P} 0.$$

◇

Theorem 1(i) shows that the estimated set converges in probability to the true set, leveraging the isomorphism between the Hausdorff distance between sets and the support function:

$$\sup_{q,\ \|q\|\le 1} |\delta(q,\hat{\Gamma}_I) - \delta(q,\Gamma_I)| = d_H(\hat{\Gamma}_I, \Gamma_I). \tag{26}$$

Theorem 1 also shows that the asymptotic distribution of $\hat{\delta}(q;\Gamma_I)$ depends on the nature of $\mathcal{G}_0(q)$ and $\mathcal{L}_0(q)$. If both sets are singletons, then the estimated support function is asymptotically normal, with a variance that can be estimated from the data, otherwise the asymptotic distribution is generally not a normal distribution.

First, if more than $\mathrm{P}+1$ linear inequalities intersect at a vertex of the identified set, then $\mathcal{L}_0(q)$ is not a singleton when q corresponds to one of the multiple outer normal vectors of the supporting hyperplanes at that vertex. In turn, the limit of the estimated support function in this direction is not a normal distribution. Note that this case constitutes a violation of the classical LICQ (Linear Independence Constraint Qualification) assumption usually made in the optimization literature.

Second, if the direction q corresponds to the outer normal of an exposed face, then $\mathcal{G}_0(q)$ is the set of points on that exposed face (Bontemps et al., 2012). In turn, the estimated support function in those directions is no longer asymptotically normal. It is generally impossible to learn from the data whether q is the outer normal of an exposed face and, therefore, further developments are needed to build a valid confidence region.

For example, Fang and Santos (2019) show that the presence of an exposed face causes non-differentiability of the support function and propose a modified bootstrap procedure to address this issue. Cho and Russell (2024) and Gafarov (2025) perturb the linear program, Chandrasekhar et al. (2019) replace discrete explanatory variables with continuous ones, and Chernozhukov et al. (2015) as well as Bontemps et al. (2025) smooth the identified set. Here, we adopt an approach following the latter strategy, which allows us to construct confidence regions from standard tools after appropriately smoothing the identified set.

Smoothing the Identified Set

If the identified set is strictly convex and no moment inequality is redundant, then the support function is differentiable everywhere and $\mathcal{G}_0(q)$ and $\Lambda_0(q)$ are singletons. To obtain strict convexity, we replace our collection of linear inequality constraints with a smooth convex constraint. To do so, we consider a convex approximation of the function $x_+ = \max(x, 0)$.

Following Chen and Mangasarian (1995), for any $\kappa > 0$, the function:

$$f_\kappa(x) = x + \frac{1}{\kappa}\log(1+\exp(-\kappa x)) = \frac{1}{\kappa}\log(1+\exp(\kappa x)),$$

is strictly convex and lies above x_+. The maximum distance between the two functions is equal to $\log(2)/\kappa$ (at $x = 0$) and it holds that:

$$f_\kappa(x) - \log(2)/\kappa \leq x_+ \leq f_\kappa(x),$$

for each x in $\mathbb{R}$. As a result, f_κ converges uniformly to x_+ as κ goes to infinity. By using this insight, we replace the original R inequality constraints with:

$$g_\kappa(\gamma) \coloneqq \sum_{r=1}^{\mathrm{R}} f_\kappa(\mathbb{E}(Z_{m,r}B_m)^\top\gamma - \mathbb{E}(Z_{m,r}A_m)) - \mathrm{R}\log(2)/\kappa \leq 0. \tag{27}$$

Let Γ_I^κ be the collection of γ values that satisfy (27). Observe that Γ_I^κ is strictly convex because g_κ is strictly convex. Moreover, the Hausdorff distance between Γ_I^κ and Γ_I is bounded above by V/κ, where V depends on Γ_I (Chen and Mangasarian, 1995).

By using (27), we rewrite (24) as:

$$\begin{aligned} \delta(q;\Gamma_I^\kappa) \coloneqq \sup_{\gamma\in\Gamma} & \, q^\top\gamma, \\ \text{s.t. } & \, g_\kappa(\gamma) \leq 0. \end{aligned} \tag{28}$$

In Appendix H.1, we show that (28) is a linear optimization problem with exponential cone constraints. This problem can be efficiently solved with any solver used in convex optimization (RMOSEK in our case).

Let the estimate of Γ_I^κ be defined as:

$$\widehat{\Gamma}_I^\kappa \coloneqq \Big\{\gamma\in\Gamma : \sum_{r=1}^{\mathrm{R}} f_\kappa\Big(\frac{1}{\mathrm{M}}\sum_{m=1}^{\mathrm{M}} Z_{m,r}B_m^\top\gamma - \frac{1}{\mathrm{M}}\sum_{m=1}^{\mathrm{M}} Z_{m,r}A_m\Big) - \mathrm{R}\log(2)/\kappa \leq 0\Big\}, \tag{29}$$

and its support function be:

$$\hat{\delta}(q;\Gamma_I^\kappa) \coloneqq \delta(q;\widehat{\Gamma}_I^\kappa).$$

Theorem 2 provides the asymptotic distribution of $\hat{\delta}(q;\Gamma_I^\kappa)$ in any direction q for a fixed κ.

Theorem 2. Assume that the moments of order $2+\tau$ of the random variables exist for some $\tau > 0$. Then,

(i) The estimated support function, $\hat{\delta}(q;\Gamma_I^\kappa)$, converges in probability to the true support function, $\delta(q;\Gamma_I)$, uniformly in q in the unit ball, when κ tends to infinity;

(ii) It holds that:

$$\sqrt{\mathrm{M}}\left(\hat{\delta}(q;\Gamma_I^\kappa) - \delta(q;\Gamma_I^\kappa)\right) = \mathbb{G}_\kappa(q) + o_P(1),$$

where

$$\mathbb{G}_\kappa(q) \coloneqq \lambda_q\left(\sum_{r=1}^{\mathrm{R}} \frac{W_r(\gamma_q)}{1+\exp\left(-\kappa\left(\mathbb{E}(Z_{m,r}B_m)^\top\gamma_q - \mathbb{E}(Z_{m,r}A_m)\right)\right)}\right),$$

γ_q is the unique point of the frontier of Γ_I^κ achieving the supremum of $q^\top\gamma$ and λ_q, the (unique) Lagrange multiplier related to the constraint $g_\kappa(\gamma) \leq 0$.

$\diamond$

As $\delta(q, \Gamma_I^\kappa)$ is differentiable everywhere in q—thus solving the issue of exposed faces and restoring LICQ—Theorem 2 shows that the asymptotic distribution of $\hat{\delta}(q; \Gamma_I^\kappa)$ is normally distributed with a variance that can be estimated from the data, as detailed in Appendix G.2 using a SHAC estimator.

Theorem 2 allows us to construct a confidence region for the outer set Γ_I^κ that is valid for the identified set. In particular, as in Bontemps et al. (2012), the following test statistic can be used to test the null hypothesis that a given parameter γ belongs to the identified set:

$$T_{\mathrm{M}}(\gamma) := \sqrt{\mathrm{M}} \frac{\hat{\delta}(q_{\mathrm{M}}; \Gamma_I^\kappa) - q_{\mathrm{M}}^\top \gamma}{\sqrt{v_\kappa(q_{\mathrm{M}})}}, \tag{30}$$

where $q_{\mathrm{M}} = \arg\inf_{q, \|q\|=1} (\hat{\delta}(q; \Gamma_I^\kappa) - q^\top \gamma)$ and $v_\kappa(q)$ is the variance of $\mathbb{G}_\kappa(q)$. Let n_α be the α-th quantile of the standard normal distribution. We define our confidence region with significance level α as the set of points that are not rejected by our test:

$$CR_{1-\alpha}^{\mathrm{M}} := \{\gamma \in \Gamma,\ T_{\mathrm{M}}(\gamma) \geq n_\alpha\}.$$

Recall that if the support function of the identified set were known, then the numerator in (30) would be expected to be positive. Consequently, a candidate γ is included in the confidence region if the test statistic is not excessively negative.

The following corollary shows the asymptotic validity of our procedure.

Corollary 1. Under the assumptions of Theorem 2, we have:

$$\lim_{\mathrm{M} \to +\infty} \inf_{\gamma \in \Gamma_I} \Pr\left(\gamma \in CR_{1-\alpha}^{\mathrm{M}}\right) \geq 1 - \alpha.$$

Equipped with the asymptotic distribution derived in Theorem 2, we can also construct a confidence interval for each component or linear combination of components of γ, using an appropriate choice of q. See Appendix H.3 for more details.

An important aspect of constructing our confidence region is the selection of the tuning parameter κ. When κ is small, the outer set Γ_I^κ is larger than the true identified set, leading to a conservative confidence region. In contrast, when κ is large, the gap between Γ_I^κ and Γ_I diminishes, but the accuracy of the asymptotic approximation worsens.

Deriving the optimal value of κ is beyond the scope of this paper. However, an intuition for the role of smoothing is as follows. Recall that the estimated support function of the identified set is not asymptotically normal when the direction considered, q, is the outer normal of an exposed face. The problem arises from the fact that all points on this exposed face reach the supremum, leading to a non-standard asymptotic distribution. The smoothing approach implicitly "chooses" one point from this exposed surface and restores an asymptotically normal distribution. Indeed, there is only one constraint, r_0, which is binding and the variance of $\hat{\delta}(q;\Gamma_I^{\kappa})$ is approximately equal to the variance of $W_{r_0}(\gamma_q)$, a point on the exposed face, when κ is large. In turn, the optimal choice of κ depends on the magnitude and variability of the moments $\frac{1}{\mathrm{M}}\sum_{m=1}^{\mathrm{M}}(Z_{m,r}B_m^{\top}\gamma - Z_{m,r}A_m)$, $r = 1,\ldots,\mathrm{R}$ considered. We expand on this discussion in Appendix H.2. See also (Bontemps et al., 2025) for a complete technical discussion in a different setting. In our empirical application, we implement a heuristic to select κ, expanding the estimated identified set by a given 3% to ensure a valid yet not overly conservative confidence region. Moreover, we show in Appendix H.2 that such 3% leads to a value of κ which guarantees the accuracy of the approximation we use for our inference procedure.

6 Data

Our data are from the Airline Origin and Destination Survey (DB1B) and consist of a 10% random sample of all tickets issued in the United States during the third quarter of 2011. By then, the merger between United Airlines and Continental Airlines had been completed and American Airlines and US Airways had not yet announced their intention to merge. We restrict the sample to the flights operated between the 85 largest metropolitan statistical areas (MSAs) in the United States.[6] Hereafter, we refer to MSAs as cities.

A market is defined as a non-directional pair of two cities. If a city has multiple airports, we combine them into a single entity. We delete tickets with multiple operating carriers or multiple ticketing carriers; tickets with different inbound and outbound itineraries; tickets that are not round-trip; and connecting tickets with more than one stop (only 1.47% of all passengers). A product is a combination between an itinerary between two cities with, at most, one stop and an airline ticketing this trip. We consider tickets featuring the same airline-itinerary combination but with different fares as the same product. We compute the corresponding price as the trimmed average price, weighted by the number of passengers.[7][8]

We compute the market sizes using data from the US Census Bureau on MSA population. In particular, we calculate the size of a market as the geometric mean of the populations at the endpoints. We compute the share of a product as the total number of passengers buying that product divided by the market size.

The major carriers in the sample are United Airlines (UA), Delta Airlines (DL), Amer-

[6] We focus on domestic operations because the effects of the merger between American Airlines and US Airways were mainly felt domestically.

[7] We delete tickets with fares in the highest and lowest percentiles and tickets with fares below $25.

[8] We do not delete flights connecting via non-hub airports and use them to estimate the second stage. However, they only account for around 19% of all connecting passengers and 17% of all connecting flights. Almost all of these flights are offered by Southwest Airlines and by airlines classified as LCC or Others. For the four legacy carriers (American Airlines, Delta Airlines, United Airlines, and US Airways), connecting flights through non-hub airports account for 1.5% of all connecting passengers and 3.2% of all connecting flights. In practice, we do not delete data on flights connecting via non-hub airports; we use them to estimate the second stage of the model. However, we do not use these flights to construct our moment inequalities for the first stage, where we focus exclusively on connecting flights via hubs.

ican Airlines (AA), US Airways (US), and Southwest Airlines (WN). The four legacy carriers rely on hub-and-spoke operations. Southwest Airlines does not exploit a pure hub-and-spoke business model, but a hybrid system in which a small number of airports are focus cities offering some of the services generally found at hubs. For our model estimation, we treat these focus cities as hubs. Nevertheless, as discussed in Table 3 below, we account for Southwest's different business model when defining the first-stage instruments, which differ from those used for the legacy carriers. All the other carriers in the sample are put either in a group called Low-Cost Carriers (LCC), or in a group called Other.[9] These carriers are considered to be fringe competitors, differing only in whether or not they can be classified as low cost. Further, to enhance computational tractability, we do not consider the fixed costs of LCC and Other when estimating the first-stage parameters, and we assume that their networks are exogenously determined before the game begins. See Table 1 for the list of hubs and focus cities in 2011.

Figures 3 to 6 illustrate the evolution of hubs from 1978 to 2011. As indicated by these figures, the 1980s saw a rapid expansion of hubs driven by consolidation, during which airlines absorbed competitors' hubs and established new ones to support network expansion. However, few hubs changed between 2000 and 2011. Only a small number of hubs disappeared between 1990 and 2011—mainly due to earlier merger activity—and no new hubs emerged. Overall, the patterns shown in Figures 3 to 6 support our decision to treat hub locations as exogenously predetermined in the model.[10]

Table 1: Hubs of the legacy carriers and focus cities of Southwest Airlines in 2011.

AA	**DL**	**UA**	**US**	**WN**
Dallas	Atlanta	Washington DC	Charlotte	Washington DC
New York	Cincinnati	Denver	Washington DC	Denver
Los Angeles	Detroit	Houston	Philadelphia	Houston
Miami	New York	New York	Phoenix	Las Vegas
Chicago	Memphis	Los Angeles		Chicago
	Minneapolis-Saint Paul	Chicago		Phoenix
	Salt Lake City	San Francisco		

[9]The group LCC contains Frontier Airlines, Alaska Airlines, Spirit Airlines, JetBlue Airlines, Virgin Airlines, Sun County Airlines, and Allegiant Air. The group Other contains AirTran Airways and USA3000 Airlines.

[10]Note that hub-and-spoke systems primarily developed after the deregulation of the airline industry in 1978. For 1978, our maps highlight cities where airports already had a major presence and, in some cases—such as Delta Airlines in Atlanta—were operating full-fledged hub-and-spoke systems. Moreover, the maps display the hubs of all airlines that eventually merged with the airline shown in the figure. For example, the maps for Delta Airlines include former Northwest Airlines hubs, as Delta and Northwest merged in 2008, with Delta retaining some of Northwest's hubs.

Figure 3: Evolution of American Airlines hubs

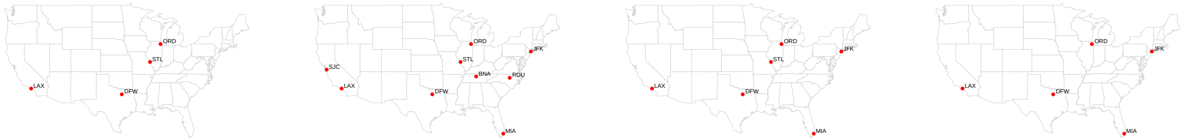

1978 1990 2000 2011

Figure 4: Evolution of Delta Airlines hubs

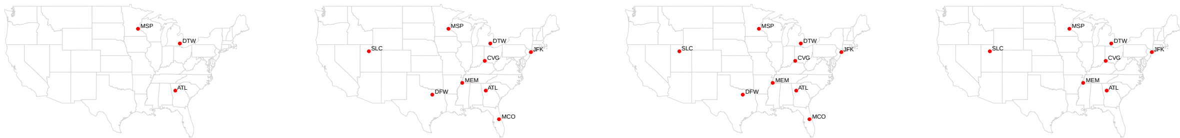

1978 1990 2000 2011

Figure 5: Evolution of United Airlines hubs

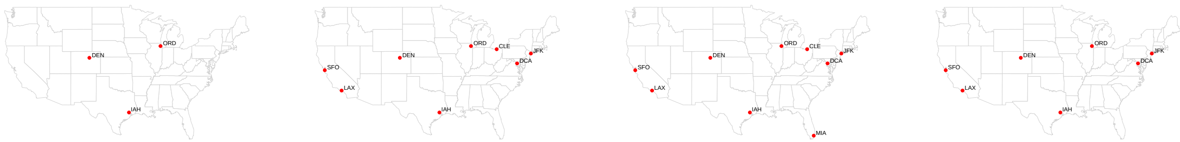

1978 1990 2000 2011

Figure 6: Evolution of US Airways hubs

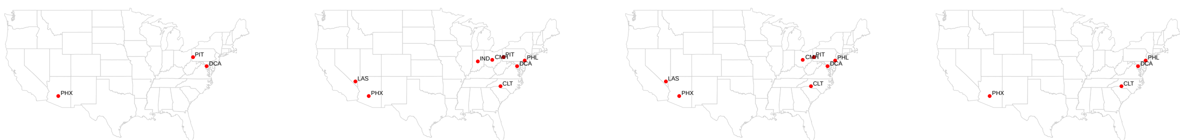

1978 1990 2000 2011

The observed demand shifters, $X_{j,m}$, include the number of stops (which is 1 or 0, "*Indirect*"), the maximum number of direct flights offered at the itinerary's endpoints by the same carrier offering itinerary j ("*Nonstop Origin*"), the distance flown in thousands of miles, ("*Distance*"), and its squared value ("*Distance2*"). We also add to $X_{j,m}$ carrier and city fixed effects, in order to capture unobserved brand- and city-specific features. We allow the marginal cost parameters to differ between short-haul and long-haul flights, which are defined as flights covering up to 1,500 miles and flights covering more than 1,500 miles, respectively. The observed marginal cost shifters, $W_{j,m}$, include the number of stops ("*Indirect*"), the average number of cities that are reachable from the endpoints of itinerary j with the same carrier offering itinerary j either with direct flights or one-stop flights ("*Connections*").[11] We also add to $W_{j,m}$ carrier fixed effects.

Table 2 provides some average statistics. Panel (b) reveals that American Airlines and US Airways are the two smallest carriers among the major airlines; however, when combined, they represent the largest airline in terms of pre-merger market share. Panel (c) shows that hub cities are much more connected than non-hub cities, as evidenced by the average degree and density.

Table 3 outlines the instruments used in the first stage, constructed from factors such as hub status (included in both the fixed-cost equation and the demand and marginal-cost equations), population size and distance (acting as demand and marginal-cost shifters), and the market's historical continuity of service over many years (an external element that can influence serving decisions without directly affecting fixed costs). Table 4 demonstrates the effectiveness of these instruments in predicting serving and non-serving decisions. Table 3 shows that the first-stage instruments for Southwest Airlines differ from those used for the legacy carriers. This distinction arises because Southwest employs a hybrid system with focus cities rather than a pure hub-and-spoke model (see Section 6). Consequently, we slightly modified the first-stage instruments for Southwest Airlines to ensure better predictive power, as indicated in Table 4. Importantly, this modification does not result in misspecification since we fully account for the instruments' predictive power in our moment inequalities.

Lastly, the instruments used in the second stage are functions of the characteristics of the competing products which is the standard in the literature.[12]

[11] We define "*Nonstop Origin*" for a direct or one-stop flight offered by airline f in market $\{a, b\}$ as

$$\max\left\{\sum_{c\in\mathcal{C}} G_{ac,f},\ \sum_{c\in\mathcal{C}} G_{bc,f}\right\},$$

where $\mathcal{C}$ is the set of cities. We define "*Connections*" as

$$\frac{1}{2}\left[\sum_{c\in\mathcal{C}} G_{ac,f} + \sum_{h\in\mathcal{H}}\sum_{c\in\mathcal{C}} G_{ah}G_{hc} + \sum_{c\in\mathcal{C}} G_{bc,f} + \sum_{h\in\mathcal{H}}\sum_{c\in\mathcal{C}} G_{bh}G_{hc}\right],$$

where $\mathcal{H}$ is the set of hubs. In the computation, we avoid double counting of itineraries.

[12] See Berry and Jia (2010), section IV.B, for a complementary discussion. In particular, our second-stage instruments include: the number of firms present in the market, the number of itineraries offered in the market, the number of products offered in the market, a dummy variable for destination being a hub, a dummy variable for the market being a monopoly, the number of rival firms offering direct flights in the market, the square of the number of rival firms offering direct flights in the market, a dummy variable for route to/from Florida/Las Vegas and some interactions between these instruments.

Table 2: Summary statistics.

(a) Sizes		
Number of firms	7	
Number of products	17,481	
Number of markets	3,146	
Fraction of direct flights	0.14	
Fraction of hub itineraries	0.83	
Number of passengers	25.33	
Fraction of direct passengers	0.85	
Fraction of passengers in hub markets	0.57	
Fraction of markets served	0.93	
(b) Market shares by airline		
AA	0.12	
DL	0.19	
UA	0.15	
US	0.09	
WN	0.24	
LCC	0.16	
Other	0.05	
(c) Network statistics	**Mean**	**St.Dev**
Degree (Hub)	49.86	13.03
Density (Hub)	0.61	0.16
Degree (No hub)	7.21	7.72
Density (No hub)	0.09	0.09
Clustering (Global)	0.17	
(d) Demand and marginal cost variables	**Mean**	**St.Dev**
Price (100 USD)	4.32	1.20
Indirect	0.86	0.34
Nonstop Origin (100)	0.20	0.19
Connections (100)	0.56	0.15
Distance (1,000 miles)	1.44	0.68
Product share	4.61e-04	1.48e-03
Market Size (1 million)	2.55	1.85
(e) Market-level statistics	**Mean**	**St.Dev**
Number of firms	3.59	1.81
Number of products	5.56	4.43
Number of direct flights	0.75	1.20
Number of hub itineraries	4.62	3.43
Number of passengers (1,000)	8.05	24.43
Number of direct passengers (1,000)	6.82	23.98
Number of passengers in hub markets (1,000)	4.60	15.39

Note: The degree is the number of links out of a node. The density is the ratio between the actual number of links and the total number of potential links. The clustering coefficient is the ratio between the number of closed triplets and the total number of triplets. A hub market for airline f is defined as a market in which one or both endpoints are hubs for airline f.

Table 3: First-stage instruments.

	$Z_{(-ab),f,r} = 1$ if
$Z_{(-ab),f,1}$	a has a population $> 10^6$; b is the closest hub to a; distance between a and b is > 150 miles
$Z_{(-ab),f,2}$	a has a population $> 10^6$; b is the second-closest hub to a; distance between a and b is > 150 miles
$Z_{(-ab),f,3}$	a has a population $> 10^6$; b is the third-closest hub to a[i]
$Z_{(-ab),f,4}$	$\{a, b\}$ is a non-hub market; it has been continuously served for the last 60 quarters[ii]
	$Z_{(+ab),f,r} = 1$ if
$Z_{(+ab),f,1}$	a has a population $< 10^6$; b is the furthest hub from a[iii]
$Z_{(+ab),f,2}$	a has a population $< 10^6$; b is the second-furthest hub from a[iv]
$Z_{(+ab),f,3}$	$\{a, b\}$ is a non-hub market; a has a population $< 10^6$; b is a competitor's hub[v]

[i] Not used for Southwest Airlines.
[ii] 45 out of the last 60 quarters for Southwest Airlines.
[iii] b is the second-furthest hub from a for Delta Airlines; b is the third-furthest hub from a for United Airlines.
[iv] b is the third-furthest hub from a for Delta Airlines; b is the fourth-furthest hub from a for United Airlines.
[v] a has a population $> 2 \times 10^6$ for Southwest Airlines.

Table 4: Predictive capacity of first-stage instruments.

	AA	**DL**	**UA**	**US**	**WN**
$\widehat{\Pr}(G_{ab,f} = 1 \mid Z_{(-ab),f,1} = 1)$	89.5%	100.0%	100.0%	88.1%	87.5%
$\widehat{\Pr}(G_{ab,f} = 1 \mid Z_{(-ab),f,2} = 1)$	86.7%	97.7%	95.2%	80.4%	84.4%
$\widehat{\Pr}(G_{ab,f} = 1 \mid Z_{(-ab),f,3} = 1)$	82.2%	72.1%	95.3%	67.4%	—
$\widehat{\Pr}(G_{ab,f} = 1 \mid Z_{(-ab),f,4} = 1)$	100.0%	80.0%	100.0%	100.0%	98.4%
$\widehat{\Pr}(G_{ab,f} = 0 \mid Z_{(+ab),f,1} = 1)$	100.0%	97.0%	72.7%	97.0%	65.4%
$\widehat{\Pr}(G_{ab,f} = 0 \mid Z_{(+ab),f,2} = 1)$	75.8 %	78.8%	54.5%	60.6%	55.8%
$\widehat{\Pr}(G_{ab,f} = 0 \mid Z_{(+ab),f,3} = 1)$	98.5%	96.3%	95.7%	92.1%	83.7%

7 Results

7.1 Results from the Demand and Supply

The second-stage results are presented in Table 5. The standard errors are computed to account for the fact that the data are not mutually independent across markets due to the spillover variables. In particular, we use a SHAC estimator of the variance as explained in Appendix G.1.

We find significant spillovers in entry on the demand side. Specifically, passengers benefit from having a large number of direct flights offered by an airline at the itinerary's endpoints ("*Nonstop Origin*"). Hence, denser networks increase consumers' willingness to pay for an airline's flights, all the rest being constant. To give an idea of the magnitude of the spillovers, note that, on average, an origin city allows passengers to reach 20 destinations with a given airline. Doubling this number generates an increase in utility equivalent to a decrease in price of around \$30.[13]

We estimate that the nesting parameter, λ, is around 0.6. Therefore, we can conclude that there is substitution between the inside goods and the outside option.

The price coefficient is negative, lying between the coefficients for the two consumer types considered by Berry and Jia (2010) and within the range reported in other studies.

[13] Remember that both "*Nonstop Origin*" and "*Price*" are measured in hundreds.

The implied elasticities further corroborate these findings: the mean own-price elasticity across products is estimated at -3.782, with a standard deviation of 1.092, and the aggregate elasticity—which measures the percentage change in the inside product share when all products' prices rise by 1%—is estimated at -2.104. These figures are consistent with previous findings in the literature. See, for example, Ciliberto and Williams (2014) and Li et al. (2022).

Passenger utility is an inverted U-shaped function of the distance flown. This means that, as distance increases, air travel becomes more pleasant relative to the outside option. However, as distance increases further, travel becomes less enjoyable and demand starts to decrease.

Passengers exhibit a strong disutility for connecting flights. In particular, passengers would be willing to pay $306 to turn their flight from a connecting flight into a direct one. This figure aligns with previous studies. Both Berry and Jia (2010) and Ciliberto and Williams (2014) consider two passenger types—“business” and “leisure”. According to Berry and Jia (2010), the estimated willingness to pay to switch from a connecting to a direct flight is $68 for leisure passengers and $443 for business passengers in 1999 (with a weighted average of $184), and $56 and $510, respectively, in 2006 (with a weighted average of $224). In Ciliberto and Williams (2014), the estimates are $98 for leisure travelers and $741 for business travelers in 2006-2008 (with a weighted average of $330).[14]

Table 5: Second-stage estimates.

Utility			**Marginal Cost**		
	Coefficient	SE		Coefficient	SE
Mean utility			**Short-haul flights**		
Intercept	-5.598	(0.309)	Intercept	3.118	(0.291)
Price	-0.587	(0.085)	Indirect	0.031	(0.041)
Indirect	-1.794	(0.088)	Distance	0.474	(0.086)
Nonstop Origin	0.868	(0.088)	Connections	-1.245	(0.174)
Distance	0.289	(0.124)	**Long-haul flights**		
Distance2	-0.093	(0.029)	Intercept	3.703	(0.295)
Nesting Parameter (λ)	0.623	(0.035)	Indirect	-0.189	(0.055)
			Distance	0.667	(0.053)
			Connections	-2.016	(0.192)
Carrier FEs			**Carrier FEs**		
DL	-0.168	(0.029)	DL	0.082	(0.043)
UA	-0.387	(0.026)	UA	0.050	(0.031)
US	0.142	(0.039)	US	0.079	(0.043)
WN	-0.519	(0.035)	WN	-0.363	(0.035)
LCC	-0.348	(0.063)	LCC	-1.509	(0.077)
Other	-0.074	(0.067)	Other	-1.398	(0.062)
Statistics					
J-statistic (p-value)	15.627	(0.156)			

Note: “*Price*” and marginal costs are in hundreds of USD. “*Connections*” and “*Nonstop Origin*” are in hundreds. “*Distance*” is in thousands of miles. City fixed effects are included in the demand. The number of over-identifying restrictions is 11.

We also find significant spillovers in entry on the marginal-cost side. Specifically, the marginal cost of an itinerary decreases with the average number of cities that an airline allows to reach from the itinerary's endpoints and intermediate stops (“*Connections*”). This

[14] These figures are based on our own calculations from the results in the tables in the two cited papers.

is due to economies of density: the larger the number of final destinations consumers can reach, the more the opportunities for an airline to pool passengers from several itineraries into the same planes, and so the more an airline can efficiently use larger aircraft that typically have lower unit costs. Hence, denser hub-and-spoke structures lead to marginal cost savings, all the rest being constant. To give an idea of the magnitude of the spillovers, observe that an additional connection reduces marginal costs by \$1.2 for short-haul and \$2 for long-haul flights. These magnitudes are comparable to increasing the distance flown by around 30 miles. The impact of the variable "*Connections*" is more pronounced for long-haul flights, as the efficiency of large planes is especially evident in long routes. Further, long-haul one-stop flights have lower marginal costs than long-haul direct flights, again by virtue of economies of density. Short-haul one-stop flights are not significantly cheaper or more expensive than short-haul direct flights. This is because economies of density may be offset by the extra take-off and landing, which uses a large volume of fuel.[15] The marginal costs of both long-haul and short-haul flights increase with the distance flown. This is because, as distance increases, more fuel is needed to cover the itinerary. Lastly, as expected, Southwest Airlines, LCC, and Other have lower marginal costs than the legacy carriers.

A layover increases the marginal cost of short-haul flights by \$3.1 and decreases the marginal cost of long-haul flights by \$18.9 (see the variable "*Indirect*" in the marginal cost). Relative to the Intercept, these numbers represent a marginal cost increase of 1% for short-haul flights and a marginal cost decrease of 5.11% for long-haul flights. Berry and Jia (2010) find cost decreases of 3.85% (short-haul) and 5.77% (long-haul) for 1999 and increases of 4.52% (short-haul) and 3.87% (long-haul) for 2006, which are comparable in magnitude to our findings.

Table 6 shows the estimated mean variable profits, mean prices, mean marginal costs, and mean markups for each airline.[16] Our results are in line with previous findings in the literature studying competition in airline markets.

Table 6: Profits by firms.

Firm	Profits (100k)	Price	Marginal cost	Markup	Lerner Index
AA	1.78	453.36	335.20	118.16	0.28
DL	1.41	436.45	310.40	126.05	0.31
UA	1.25	445.56	328.43	117.13	0.28
US	1.30	453.43	336.77	116.67	0.27
WN	2.79	419.43	299.51	119.92	0.31

Note: Quantities are in USD.

[15] See Berry and Jia (2010), page 25.

[16] See Appendix I.1 for a presentation of the same figures at different levels of aggregation.

7.2 Results from Entry

The first-stage results are presented in Table 7. The second and third columns of Table 7 report the projection of the estimated identified set of the first-stage parameters, $\widehat{\Gamma}_I$, defined in (25).[17] While we allow for a firm-specific $\gamma_{2,f}$ (i.e., heterogeneous congestion effects on the fixed costs), we distinguish the intercept $\gamma_{1,f}$, representing the baseline fixed costs, only between the four legacy carriers (United Airlines, Delta Air Lines, American Airlines, and US Airways) and Southwest Airlines. This decision is motivated by finite sample considerations, namely, the fact that the number of markets selected by each instrument for each airline would be very small if considered separately across airlines, rendering the asymptotic approximations used for inference implausible.

We observe that, in the absence of congestion costs at hubs, the baseline fixed costs, $\gamma_{1,f}$, for offering a direct service between two endpoints range from \$570,640 to \$1,010,410 for the four legacy carriers. In contrast, Southwest Airlines exhibits higher baseline fixed costs, suggesting that it requires a larger passenger volume to operate a direct service between two cities. Specifically, for a mean price of \$430 and a mean Lerner index of 30%, a legacy carrier will offer a direct flight between two non-hub cities if it can transport between 4,423 and 7,832 passengers in a quarter, while Southwest Airlines requires between 6,600 and 15,448 passengers to do so.

To interpret the congestion cost parameters, $\gamma_{2,f}$, consider the second and third columns of Table 8, which report the estimated increase in fixed costs when adding a spoke to a hub with 49 spokes (the mean degree of a hub). The results vary across firms, although we cannot reject the hypothesis that the increases in fixed costs are equal. For instance, US Airways and Southwest Airlines record higher minimum fixed costs than the other airlines, indicating that they potentially face the highest total fixed costs for a hub with 50 spokes. This may be due to the fact that Southwest Airlines employs a business model that combines hub-and-spoke and point-to-point operations, while US Airways operates a relatively small network concentrated mainly in the Eastern US. These higher fixed costs are counterbalanced by lower marginal costs, as highlighted in Table I.1 of Appendix I.1.

We assess our fixed cost estimates by computing each airline's variable cost share of operating costs and comparing these with the FAA's 2018 estimates (see Appendix I.2), which we find to be closely aligned with our results.

The other columns of Tables 7 and 8 report the quantities discussed above, but now correspond to the estimated superset of the identified set, $\widehat{\Gamma}_I^\kappa$, defined in (29), for $\kappa = 100$ and $\kappa = 500$.[18] As expected, the projections of $\widehat{\Gamma}_I^\kappa$ widen as κ decreases. From Table 8 onward, we report our results based on $\kappa = 100$. This value of κ expands the estimated identified set by 3% and guarantees the accuracy of the approximation we use for our

[17] As discussed in Section 5, we compute the projections of $\widehat{\Gamma}_I$ by solving the linear program (24).

[18] As discussed in Section 5, we compute the projections of $\widehat{\Gamma}_I^\kappa$ by solving the linear program with exponential cone constraints (28).

inference procedure, as discussed in Section 5 and in Appendix H.2.

The fourth and fifth columns of Table 9 show the 95% confidence intervals for each component of γ, obtained from the application of Theorem 2.[19] Implementing the misspecified adaptive confidence intervals of Stoye (2021) does not alter these intervals because the estimated difference between the upper and lower bounds of each component is large. Therefore, we do not report it.

To further validate our smoothing approach, instead of $T_{\mathrm{M}}(\gamma)$ defined in (30), we consider the following test statistic for the null hypothesis that a given parameter belongs to the identified set:

$$\xi_{\mathrm{M}}(\gamma) = \max_{r=1,\ldots,R} \frac{\sqrt{\mathrm{M}}\,\frac{1}{\mathrm{M}}\sum_{m=1}^{\mathrm{M}}\left(Z_{m,r}B_m^{\top}\gamma - Z_{m,r}A_m\right)}{\sqrt{W_r(\gamma)}}.$$

The test statistic $\xi_{\mathrm{M}}(\gamma)$ is frequently used in empirical work in conjunction with the conservative yet competitive critical value proposed by Chernozhukov et al. (2018). We use this procedure to construct the confidence region for $(\gamma_{1,f}, \gamma_{2,f})$ for all legacy carriers as well as for Southwest Airlines.[20] Specifically, we collect the values of the first-stage parameters that are not rejected by the test using the one-step Self-Normalized critical value at a 5% significance level (see Chernozhukov et al., 2018). The last two columns of Table 9 present the projection of these confidence regions on each axis. Although projecting the confidence region is a conservative approach, our confidence intervals are of the correct order of magnitude.

We provide further evidence of the performance of our inference method in Appendix H.2.

Table 7: First-stage estimates.

	$\widehat{\Gamma}_{I}$		$\widehat{\Gamma}_{I}^{\kappa}$, $\kappa = 100$		$\widehat{\Gamma}_{I}^{\kappa}$, $\kappa = 500$	
	LB	UB	LB	UB	LB	UB
Baseline fixed costs ($\gamma_{1,f}$)						
Legacy carriers (AA, US, DL, UA)	570,941	1,010,411	568,393	1,042,329	570,432	1,016,798
WN	851,306	1,992,769	844,624	2,013,831	849,970	1,996,981
Congestion costs ($\gamma_{2,f}$)						
AA	11,297	35,529	10,651	36,180	11,168	35,659
DL	7,953	27,991	7,560	28,486	7,874	28,090
UA	8,173	17,941	7,753	18,152	8,089	17,983
US	16,382	34,450	15,877	34,973	16,281	34,555
WN	18,889	36,348	18,450	36,720	18,801	36,423

Note: Quantities are in USD.

[19]As discussed in Section 5, we construct a confidence interval for each component of the vector γ by solving the linear program with exponential cone constraints (28), using critical values from the Normal distribution as stated in Theorem 2, and standard errors based on a SHAC estimator as described in Appendix G.2.

[20]Recall that our first-stage instruments are airline-specific and that $\gamma_{1,f}$ is only distinguished between the legacy carriers and Southwest Airlines.

Table 8: Estimated fixed costs for adding a 50^{th} spoke to a hub.

	$\widehat{\Gamma}_I$		$\widehat{\Gamma}_I^{\kappa},\ \kappa = 100$	
	LB	UB	LB	UB
AA	2,129	4,088	2,063	4,153
DL	1,714	3,410	1,673	3,459
UA	1,732	2,425	1,689	2,446
US	2,632	3,982	2,581	4,034
WN	3,857	4,450	3,812	4,487

Note: Quantities are in USD 1 million.

Table 9: First-stage inference.

	$\widehat{\Gamma}_I^{\kappa},\ \kappa = 100$		**95% CI**		**95% CR - CCK**	
	LB	UB	LB	UB	LB	UB
Baseline fixed costs ($\gamma_1.f$)						
Legacy carriers (AA, US, DL, UA)	568,393	1,042,329	554,483	1,299,833	555,000	1,375,000
WN	844,624	2,013,831	808,732	2,173,262	810,000	2,185,000
Congestion costs ($\gamma_{2,f}$)						
AA	10,651	36,180	5,150	41,203	1,150	42,050
DL	7,560	28,486	4,158	31,740	1,350	32,050
UA	7,753	18,152	3,542	19,780	3,700	20,100
US	15,877	34,973	11,199	39,159	7,950	39,950
WN	18,450	36,720	14,339	39,804	12,300	40,500

Note: Quantities are in USD.

8 Counterfactuals

This section studies the impact on firm and market outcomes of a merger between two of the four legacy carriers in our sample, American Airlines and US Airways. These two firms did in fact merge in 2013. They first expressed an interest in merging in January 2012 and officially announced their plans to merge in February 2013. At the time they expressed their interest in merging, American Airlines' holding company (AMR Corporation) was in Chapter 11 bankruptcy.[21] Further, the DoJ, along with several state attorney generals, sought to block the merger, as they were concerned that it would substantially lessen competition and hurt consumers. In 2013, a settlement was reached, whereby the merging parties pledged to give up landing slots or gates at seven major airports and to maintain the same level of operations in the hub markets out of Charlotte, New York (Kennedy), Los Angeles, Miami, Chicago (O'Hare), Philadelphia, and Phoenix for a period of three years.[22] Below, we refer to this settlement as the 2013 settlement. According to articles from the time the merger was announced, the parties expected the merger to make the new entity the largest airline in the world in terms of passenger numbers and to generate

[21]Recall that we use data from the third quarter of 2011, which is before the two parties expressed an interest to merge and corresponds to the last quarter before AMR Corporation filed for Chapter 11 bankruptcy.

[22]`https://www.justice.gov/opa/pr/justice-department-requires-us-airways-and-american-airlines-divest-facilities-seven-key`, `https://americanairlines.gcs-web.com/news-releases/news-release-details/amr-corporation-and-us-airways-announce-settlement-us-department`.

annual cost savings of around $1 billion per year.[23] In addition, the merger was seen by analysts as an opportunity for American Airlines to expand its footprint in markets along the East Coast, where US Airways had a strong presence.[24] The merger was the last in a series of mergers between large airlines and reduced the number of legacy carriers to four (Delta Airlines, United Airlines, Southwest Airlines, and the new American Airlines).

8.1 Set-up

When simulating a merger between airlines, the canonical approach consists of relying on demand and supply models, where the firms best respond to competitors by adjusting their prices, while holding the networks fixed. Given that the networks are held fixed, before running the simulation, the researcher needs to take a stand on which network the merged entity will inherit. In turn, this choice determines the list of products offered by the merged entity and their observed characteristics.

In particular, the previous literature has considered different scenarios, for example, a "base-case scenario," where the merging firms maintain their pre-merger products and behave as if they have colluded, and a "merged scenario," where the new merged entity inherits the network resulting from combining the pre-merger networks of the merging firms. These are just two possible scenarios, and nothing suggests that they should be taken as extreme reference cases, as there are many ways in which the merged entity may revisit its entry decisions. Our methodology eliminates such ambiguity because it allows the merged entity to best respond by adjusting both its network and prices. Depending on the dominating forces, the merged entity may find it convenient to exit some markets in order to downsize the higher total congestion costs from managing a denser network. Alternatively, it may enter new markets so as to exploit the marginal cost savings from denser hub-and-spoke structures and consumers' willingness-to-pay for flying from dense nodes. Further, our methodology allows the competitors to re-optimize both their networks and prices. For example, the competitors may find it opportune to exit markets where the merged firm has acquired excessive market power, or to enter markets where the merger has created space for other companies.

To highlight the advantages of our framework, we compare the counterfactual predictions arising from our model with those obtained using ad-hoc assumptions on the post-merger network. In particular, we consider two ad-hoc scenarios:

1. Networks fixed - base. After the merger, the merging airlines remain separate entities. All firms maintain the pre-merger networks and products. They play the simultaneous pricing game described in Section 3.1 and new equilibrium prices arise. The merging firms choose the prices that maximize their joint profits; that is, they behave as if they have

[23] `https://www.reuters.com/article/uk-americanairlines-merger-idUSLNE91D02020130214`.

[24] `https://money.cnn.com/2013/02/14/news/companies/us-airways-american-airlines-merger/index.html`.

colluded.

2. Networks fixed - updated. After the merger, all firms, except American Airlines and US Airways, maintain the pre-merger networks and products. We treat the merged entity as a new firm that gets the more favorable firm fixed effect of the merging airlines—specifically, that of US Airways on the demand side and that of American Airlines on the marginal cost side. This is akin to assuming a "best-case" scenario in which the merged entity is more efficient—producing at lower marginal cost and perceived as "better" by consumers—than the individual merging parties. The rationale behind this assumption is that if the merger does not benefit consumers even in this most favorable setting, it is even less likely to do so in a neutral or "worst-case" scenario where the merged entity has higher marginal costs and is perceived as a worse version of the individual merging parties. We assign the merged entity the network resulting from merging the pre-merger networks of American Airlines and US Airways. The products of the merged entity and their covariates are constructed from the merged network. The demand and supply shocks of the itineraries that were offered both by American Airlines and US Airways before the merger are replaced by their mean values. After such rearrangements, the firms play the simultaneous pricing game described in Section 3.1 and new equilibrium prices arise.

These two ad-hoc scenarios are compared with the simulations from our full model. When simulating our full model, we consider three possible scenarios:

1. Networks vary - without remedies. After the merger, we treat the merged entity as a new firm that gets the more favorable firm fixed effect of the merging airlines. We let the firms play the two-stage game described in Section 3. New equilibrium networks and prices arise;

2. Networks vary - with remedies. An important advantage of our approach is that it allows us to evaluate the impact on firm and market outcomes of the 2013 settlement. To do so, after the merger, we treat the merged entity as a new firm that gets the more favorable firm fixed effect of the merging airlines. We let the firms play the two-stage game described in Section 3. New equilibrium networks and prices arise. However, in contrast to the *Networks vary - without remedies* scenario, now we incorporate as binding constraints some of the DoJ's remedies contained in the 2013 settlement. In particular, we force the merged entity to continue serving all the markets that were served by the merging firms before the merger out of Charlotte, New York, Los Angeles, Miami, Chicago, Philadelphia, and Phoenix. Recall that these were the cities signalled by the DoJ, as discussed at the beginning of Section 8.[25,26]

[25]Note that this scenario differs from the *Networks fixed - Updated case* scenario because, first, the competitors of the merged entity are allowed to re-optimize their networks; second, the merged entity is allowed to increase its operations in the markets out of the hubs signalled by the DoJ and to increase/decrease its operations in the markets out of the hubs not signalled by the DoJ (Washington DC and Dallas).

[26]As highlighted at the beginning of Section 8, the 2013 settlement also required the merged airline to give up landing slots or gates at seven major slot constrained airports, in order to facilitate the expansion of low-cost carriers, such as Southwest Airlines. This part of the settlement is not incorporated in the

3. Networks vary - PHX dehubbed. In contrast to the *Networks vary - w/o remedies* scenario, we assume that the merged entity removes the hub status of Phoenix and can only offer a direct service from Phoenix to the remaining hubs and no longer to non-hub cities. We focus on Phoenix because it presents an interesting case for analysis, as it was one of the smallest hubs and was located in a relatively small city, somewhat comparable to Memphis, Cincinnati, and Cleveland—cities that lost their hub status in earlier mergers. Our framework allows us to explicitly quantify the impact of such hub closures on consumer surplus.

8.2 Results

Table 10 shows the impact of the merger on consumer surplus. Before commenting on the results, we clarify that, under the *Networks fixed* columns, we report each quantity's percentage changes for the *Networks fixed - base* and *updated* scenarios, respectively. Under the *Networks vary* columns, each item is reported in three rows. The first row displays the median percentage change across different parameter values in the estimated identified set of the fixed cost parameters and across different first-stage equilibria constructed by the counterfactual algorithms. The second row shows (within brackets) the minimum and maximum percentage changes across different parameter values in the estimated identified set of the fixed cost parameters and across different first-stage equilibria constructed by the counterfactual algorithms. The third row reports (within brackets) the minimum and maximum percentage changes across different parameter values in the confidence region of the fixed cost parameters and across different first-stage equilibria constructed by the counterfactual algorithms.[27]

The other tables in this section have a similar structure and sometimes replace percentage changes with actual values. For simplicity of exposition, our discussion will be often based on the median values or median percentage changes. Hereafter, the merged entity is also referred to as American Airlines.

The first item of Table 10 reports the impact of the merger on total consumer surplus. Absent any remedies, the merger leads to a median decrease in consumer surplus of approximately 0.95%—a modest decline, yet more pronounced than the 0.58% decrease

simulations, because our framework does not distinguish between airports in the same cities and, hence, does not explicitly model the process of slot assignment. See also Section 3.2 for a discussion on this issue.

[27] As the *Networks vary* columns account for the first stage of our model, we have an *interval* of counterfactual results for two reasons. First, a counterfactual outcome value must be computed for each parameter value within the confidence region (or estimated identified set). Second, for a given value of the parameters, multiple first-stage equilibria can lead to many counterfactual networks for each parameter value and, in turn, many counterfactual outcome values. Addressing the first issue, we forego the standard gridding approach, which is computationally burdensome, and instead discuss in Appendix H.4 a method to draw points uniformly from the confidence region (or estimated identified set). Regarding the second issue, for a given value of the parameters, we use an algorithm that generates a probability distribution of possible equilibrium networks, by iterating best-response dynamics among firms until convergence for many sequences of firms and markets, as in Lee and Pakes (2009) and Wollmann (2018) and detailed in Appendices J.1 and J.2.

observed when the airlines' networks are held fixed. Moreover, when considering the minimum percentage changes, our analysis does not rule out a decrease in consumer surplus of up to 5.31%. With the remedies discussed above, the median change in consumer surplus turns positive to 3.90%. Importantly, when examining the minimum percentage changes, the remedies rule out any consumer surplus losses. Finally, had Phoenix been dehubbed, the merger would have led to a decrease in consumer surplus. Overall, these results suggest that the remedies helped prevent consumer surplus losses and that hub closures can cause significant consumer surplus losses.

The second and third items of Table 10 show the impact of the merger on consumer surplus when we distinguish between two groups of markets. We call "new markets" the markets where the merging parties do not offer direct flights pre-merger. and where the merged entity offers direct flights post-merger. We call "old markets" the markets where the merging parties offer direct flights pre-merger. Old markets are those on which antitrust authorities typically focus their merger analysis. We can see that while the overall impact of the merger on consumer surplus is small, there is an important tension between old and new markets. On the one hand, due to the entry-exit patterns explained below, old markets undergo consumer surplus losses. If the merger's effect on consumer surplus in old markets was the relevant criterion, then the merger should have been blocked, even with the remedies we are considering. This is in line with the DoJ's initial attempt to stop the merger. On the other hand, new markets experience a considerable increase in consumer surplus, which reveals substantial positive effects of the merger and can be used to legitimize its implementation. Further, the remedies, which were tailored to prevent the exit of American Airlines from the old markets, reduce consumer surplus losses in the old markets but, at the same time, weaken consumer surplus gains in new markets. This highlights the need for antitrust authorities to carefully balance these two effects when designing network interventions. To the best of our knowledge, this tension between consumer surplus losses in old markets and consumer surplus gains in new markets is a novel empirical finding that has major implications for policymakers. Contrary to the fixed network approach, our model is best suited to study such a trade-off.

Table 12 sheds some light on what drives these effects. The last two items report the number of markets served with direct flights by American Airlines and its competitors before and after the merger out of American Airlines' and US Airways' hubs.[28] Absent any remedies, the merger would have led American Airlines and the other major airlines to reduce the size of their networks, resulting in an overall decrease in competition, fewer flight offerings, and significant consumer surplus losses. The remedies prevent American Airlines from shrinking its network and even encourage an expansion.

In particular, Table 11 sheds light on the hub-specific entry and exit patterns. Airlines

[28]Before the merger, we take the sum of the number of markets served with direct flights by US Airways and American Airlines.

Table 10: Percentage change in consumer surplus across different scenarios.

	Merger				
	Networks fixed		Networks vary		
	base	updated	w/o remedies	w/ remedies	PHX dehubbed
Total consumer surplus	-0.56	-0.58	-0.95	+3.09	-2.27
			[-2.22, +0.48]	[+2.12, +4.02]	[-3.53, -0.40]
			[-5.31, +1.40]	[+1.15, +5.31]	[-6.06, +0.39]
New markets			+68.14	+55.17	+62.38
			[+46.18, +98.51]	[+45.57, +66.24]	[+46.38, +75.26]
			[+43.75, +98.63]	[+34.11, +80.00]	[+36.97, +85.17]
Old markets	-0.56	-0.58	-7.52	-2.74	-8.68
			[-8.70, -6.08]	[-3.20, -2.10]	[-10.03, -7.40]
			[-11.05, -6.07]	[-4.16, -1.11]	[-11.67, -7.22]

Note: Consumer surplus is calculated using the log-sum formula and is expressed in USD 1 million up to the integration constant. Percentage changes with respect to the pre-merger scenario are reported.

face a trade-off when making entry and exit decisions. On the one hand, entering additional markets increases variable profits by boosting consumers' willingness to pay for flights (recall the variable "*Nonstop Origin*" in the utility function) and reducing marginal costs (recall the variable "*Connections*" in the marginal cost equation). On the other hand, expanding into additional markets raises fixed costs by increasing congestion, as reflected in the quadratic term of the fixed cost equation. Absent any remedies, American Airlines expands its presence at some hubs to lower marginal costs through economies of density and capitalize on consumers' willingness to pay for flights departing from densely connected nodes. To mitigate rising congestion costs, American Airlines reduces its flight offerings at other hubs—most notably at Los Angeles, Miami, Charlotte, Phoenix, and Philadelphia—which results in an overall smaller network. Under the *Networks vary - with remedies* scenario, however, these offsetting reductions are no longer allowed, leading to a larger overall network.

By contrast, post-merger entry by rivals does not happen even under the remedies. This may seem counter-intuitive, as one would expect reduced competition due to the elimination of one firm to free up room for competitors. We do not observe this mechanism for two reasons. First, there was relatively little overlap between American Airlines and US Airways, hence there was almost no space created for post-merger entry by the other airlines. Second, as mentioned above, by expanding its network, American Airlines increases consumers' willingness to pay for its flights and decreases its marginal costs. This makes it more challenging for the other carriers to compete with such a powerful player and may cause them to exit.

As the decline in consumer surplus in the old market observed in Table 10 suggests, the increased consumer utility from American Airlines' larger network under the remedies is offset by the contraction of competitors' operations. In Section 8.3, we show that the expansion of American Airlines' network and the reduction in competitors' networks align with the actual entry-exit dynamics observed after 2013.

Table 11: Changes in direct flights offered in the hub markets of AA and US.

	Before			Merger								
				w/o remedies			w/ remedies			PHX dehubbed		
	AA/US	Others	Avg. presence	AA/US	Others	Avg. presence	AA/US	Others	Avg. presence	AA/US	Others	Avg. presence
AA hubs												
DFW	68	55	1.6	69	57	1.52	68	57	1.52	69	57	1.54
				[68, 70]	[56, 57]	[1.51, 1.55]	[68, 70]	[57, 57]	[1.52, 1.55]	[68, 72]	[56, 57]	[1.51, 1.57]
				[67, 71]	[56, 57]	[1.50, 1.56]	[67, 70]	[56, 57]	[1.50, 1.55]	[67, 72]	[56, 57]	[1.50, 1.57]
LAX	28	90	1.51	24	90	1.41	29	90	1.47	26	90	1.44
				[22, 26]	[89, 91]	[1.37, 1.43]	[29, 30]	[89, 91]	[1.46, 1.49]	[23, 32]	[89, 91]	[1.40, 1.51]
				[20, 27]	[89, 91]	[1.35, 1.44]	[29, 30]	[89, 91]	[1.46, 1.49]	[21, 33]	[89, 91]	[1.36, 1.52]
ORD	59	129	2.35	63	107	2.06	61	112	2.12	61	110	2.09
				[61, 66]	[101, 113]	[1.99, 2.17]	[60, 64]	[105, 117]	[2.01, 2.17]	[60, 65]	[99, 115]	[1.95, 2.18]
				[60, 68]	[88, 114]	[1.84, 2.17]	[60, 65]	[101, 118]	[1.99, 2.18]	[57, 65]	[92, 118]	[1.89, 2.21]
MIA	40	51	1.17	19	52	0.87	40	52	1.12	20	52	0.88
				[18, 24]	[51, 52]	[0.85, 0.93]	[40, 44]	[51, 52]	[1.11, 1.17]	[19, 24]	[52, 52]	[0.87, 0.93]
				[18, 25]	[51, 52]	[0.85, 0.94]	[40, 44]	[51, 52]	[1.11, 1.17]	[15, 24]	[51, 52]	[0.80, 0.93]
JFK	41	113	2	48	104	1.87	51	102	1.87	47	105	1.84
				[39, 56]	[101, 111]	[1.77, 1.96]	[45, 53]	[99, 107]	[1.76, 1.91]	[33, 54]	[101, 112]	[1.66, 1.95]
				[34, 57]	[101, 112]	[1.67, 1.98]	[42, 54]	[99, 107]	[1.73, 1.93]	[33, 58]	[101, 112]	[1.66, 1.99]
US hubs												
CLT	61	41	1.29	45	42	1.06	62	41.5	1.26	44	42	1.05
				[41, 48]	[42, 42]	[1.01, 1.10]	[61, 65]	[41, 42]	[1.24, 1.29]	[41, 49]	[42, 42]	[1.01, 1.11]
				[35, 50]	[42, 42]	[0.94, 1.12]	[61, 65]	[41, 42]	[1.24, 1.29]	[35, 51]	[42, 42]	[0.94, 1.13]
PHX	41	74	1.49	23	68	1.11	41	65.5	1.3	8	68	0.93
				[21, 26]	[66, 69]	[1.07, 1.15]	[41, 41]	[62, 67]	[1.26, 1.32]	[8, 8]	[68, 70]	[0.93, 0.95]
				[13, 27]	[66, 69]	[0.99, 1.16]	[41, 41]	[62, 67]	[1.26, 1.32]	[8, 8]	[67, 70]	[0.91, 0.95]
DCA	40	130	2.16	45	130	2.13	35	129	2	40.5	130	2.09
				[35, 60]	[128, 131]	[2.01, 2.30]	[28, 45]	[127, 130]	[1.91, 2.10]	[32, 53]	[129, 131]	[1.98, 2.23]
				[29, 61]	[128, 133]	[1.95, 2.32]	[23, 49]	[127, 131]	[1.87, 2.17]	[29, 65]	[129, 133]	[1.95, 2.37]
PHL	52	53	1.33	27	59	1.05	53.5	55	1.34	28	59	1.07
				[21, 36]	[58, 60]	[0.99, 1.17]	[52, 55]	[54, 56]	[1.32, 1.37]	[20, 36]	[58, 60]	[0.98, 1.19]
				[15, 37]	[58, 63]	[0.90, 1.19]	[52, 55]	[54, 56]	[1.32, 1.37]	[15, 37]	[56, 64]	[0.91, 1.20]
Total												
Total	394	670	1.66	332	644	1.46	406	639.5	1.56	312	649	1.44
				[320, 345]	[635, 651]	[1.44, 1.48]	[401, 413]	[625, 647]	[1.54, 1.57]	[304, 325]	[636, 655]	[1.42, 1.46]
				[292, 353]	[628, 654]	[1.38, 1.49]	[387, 420]	[625, 648]	[1.53, 1.58]	[276, 340]	[633, 657]	[1.38, 1.47]

Table 12 also shows that American Airlines' markups increase substantially. In the scenario without remedies, this can be attributed to greater market power. Once the remedies are imposed and American expands its network, markups rise further both to cover the higher fixed costs of a larger post-merger network and because American only partially passes through the marginal cost savings generated by the merger.[29]

The entry-exit patterns generated by the merger lead to noticeable heterogeneity across American Airlines' hub markets. Table 13 highlights that while some hub markets experience large consumer surplus gains, others suffer from the merger. The negative effect of the merger is especially pronounced in the markets out of New York, Chicago, Miami, and Phoenix. Table 11 in Appendix J.3 also reports the hub-level changes in the number of direct flights offered by American Airlines and the other major airlines. It reveals that the consumer surplus decrease in the markets out of New York and Chicago was driven by the exit of competitors, whereas the consumer surplus decrease in the markets out of Miami and Phoenix was driven by the exit of both American Airlines and competitors. Remember that the 2013 settlement required the merged airline to divest slots at Chicago O'Hare and gates in Miami and at La Guardia Airport in New York. Such remedies—which our framework does not allow to properly model as explained in Footnote 26—likely prevented the significant rival exit that we predicted in these markets.

[29] Table J.1 in Appendix J.3 examines the changes in price and marginal cost in more detail.

Table 12: Outcomes across different scenarios.

	Before	**Merger**				
		Networks fixed		Networks vary		
		base	updated	w/o remedies	w/ remedies	PHX dehubbed
Total consumer surplus	2,684.45	-0.56	-0.58	-0.95 [-2.22, +0.48] [-5.31, +1.40]	+3.09 [+2.12, +4.02] [+1.15, +5.31]	-2.27 [-3.53, -0.40] [-6.06, +0.39]
Mean consumer surplus	3.91	-0.56	-0.58	-1.94 [-3.06, -0.24] [-6.13, +0.52]	+1.86 [+1.24, +3.02] [-0.06, +3.65]	-3.11 [-4.23, -0.97] [-6.74, -0.48]
Markups: American	119.2	+5.98	+7.34	+7.91 [+7.19, +8.49] [+6.35, +9.02]	+10.26 [+10.12, +10.52] [+9.59, +10.90]	+7.42 [+6.44, +7.88] [+5.56, +8.49]
Markups: Others	116.22	+0.07	-0.45	+0.76 [+0.52, +0.92] [+0.41, +1.54]	-0.5 [-0.64, -0.46] [-0.78, -0.31]	+0.98 [+0.78, +1.11] [+0.43, +1.70]
Segments: American	394	394	394	332 [320, 345] [292, 353]	406 [401, 413] [387, 420]	312 [304, 325] [276, 340]
Segments: Others	670	670	670	644 [635, 651] [628, 654]	639 [625, 647] [625, 648]	649 [636, 655] [633, 657]

Note: Consumer surplus is calculated using the log-sum formula and is expressed in USD 1 million up to the integration constant. Percentage changes with respect to the pre-merger scenario are reported for total consumer surplus, mean consumer surplus, and markups.

8.3 Comparison with Post-merger Data

Table 14 shows a comparison of our *Networks vary - with remedies* scenario with post-merger real data on the markets served with direct flights by American Airlines and its competitors before and after the merger out of American Airlines' and US Airways' hubs. Note that such a comparison is always difficult, as other changes—such as shifts in preferences, costs (e.g., the significant drop in kerosene prices in the 2010s), and market structure—occurred around the same time the merger was consummated.[30] Nevertheless, our model does a good job of capturing the actual entry-exit dynamics. In particular, we correctly predict the post-merger expansion of American Airlines' network and the reduction of competitors' networks. Moreover, the observed number of markets served by American Airlines with direct flights closely matches our median prediction in 2015 and remains near our upper bound until 2017. Meanwhile, the observed number of markets served by competitors with direct flights remains close to our lower bound for all the years reported.

8.4 The Importance of Spillovers

Table 15 shows how spillovers in the demand, marginal cost, and fixed cost equations affect the counterfactual outcomes. The column *Network vary - with spillovers* reports changes in various outcomes under the *Network vary - without remedies* scenario, as shown in Table 12.

[30] For more details, see, for instance, Bontemps et al. (2022).

Table 13: Percentage change in consumer surplus in the hub markets of AA and US.

	Before	**Merger**				
		Networks fixed		Networks vary		
		base	updated	w/o remedies	w/ remedies	PHX dehubbed
AA hubs						
DFW	327.86	-1.98	-1.98	+18.67	+18.59	+18.62
				[+17.38, +20.90]	[+18.38, +21.13]	[+16.53, +22.02]
				[+15.98, +21.12]	[+16.54, +21.17]	[+15.16, +22.17]
LAX	530.73	-0.29	-0.29	-1.98	+0.52	-1.34
				[-3.13, -0.95]	[+0.29, +1.80]	[-3.46, +1.08]
				[-4.74, -0.25]	[+0.02, +1.80]	[-4.90, +2.07]
ORD	502.42	-0.36	-0.36	-9.92	-8.4	-9.8
				[-12.42, -6.40]	[-12.19, -6.89]	[-14.34, -6.61]
				[-16.56, -6.35]	[-12.92, -6.51]	[-15.87, -5.90]
MIA	309.62	-0.65	-0.65	-35.47	-26.26	-35.3
				[-36.13, -33.41]	[-26.49, -24.40]	[-35.84, -33.78]
				[-36.52, -32.83]	[-26.49, -24.22]	[-37.9, -33.69]
JFK	619.04	-0.33	-0.33	-22.88	-22.58	-23.62
				[-26.49, -19.15]	[-26.44, -20.95]	[-30.3, -20.31]
				[-30.36, -18.81]	[-27.56, -20.56]	[-30.3, -18.67]
US hubs						
CLT	103.55	-2.5	-2.5	+9.27	+32.4	+8.36
				[+4.23, +12.71]	[+30.01, +37.55]	[+3.76, +13.73]
				[-2.59, +14.97]	[+29.97, +37.71]	[-2.73, +16.17]
PHX	214.94	-0.83	-0.83	-31.15	-22.58	-40.5
				[-33.58, -29.36]	[-24.02, -21.48]	[-41.01, -39.66]
				[-38.29, -28.99]	[-24.19, -21.46]	[-41.32, -39.58]
DCA	445.04	-0.42	-0.42	+29.9	+23.52	+26.78
				[+22.97, +39.78]	[+20.91, +28.78]	[+21.07, +34.23]
				[+19.37, +40.69]	[+18.28, +32.80]	[+19.61, +42.99]
PHL	173.36	-1.36	-1.36	-2.21	+24.71	-0.57
				[-6.50, +7.05]	[+22.43, +26.34]	[-7.43, +6.73]
				[-12.81, +7.96]	[+22.3, +26.39]	[-11.43, +7.90]

Note: Consumer surplus is computed using the log-sum formula and it is in USD 1 million up to constant of integration. Percentage changes with respect to the pre-merger scenario are reported.

The three columns labeled *Network vary - without spillovers* then show how these outcomes change when each type of spillover is individually shut down. Specifically, we construct the column *No demand spillovers* by preventing the merged entity from updating the demand variable "*Nonstop Origin*" (i.e., preventing it from capitalizing on consumers' willingness to pay for flights departing from densely connected nodes), keeping this variable at its pre-merger level during the post-merger network re-optimization.[31] Likewise, we construct the column *No MC spillovers* by preventing updates to the marginal cost variable "*Connection*"

[31] The pre-merger level of the demand variable "*Nonstop Origin*" is computed after combining the pre-merger networks of American Airlines and US Airways.

Table 14: Comparison of merger prediction with data from 2015-2019.

	2011	**Prediction**	**2015**	**2016**	**2017**	**2018**	**2019**	**Mean 15-19**
Segments: AA/US	394	406	409	429	427	442	452	431.8
		[401, 413]						
		[387, 420]						
Segments: Others	670	639	607	614	629	617	618	617
		[625, 647]						
		[625, 648]						

(i.e., preventing the merged entity from exploiting economies of density). Finally, we construct the column *No FC spillovers* by preventing updates to the quadratic expression in the fixed cost equation (i.e., preventing the merged entity from adjusting its congestion costs).

The demand spillover has the largest impact on the merger's outcome. Eliminating this spillover—that is, preventing the merged entity from leveraging the higher willingness to pay associated with a larger network—leads to a drastic reduction in network size, with the post-merger network being, on median, only one-fourth the size of the pre-merger network. This reduction is far more pronounced than when all spillover variables are accounted for, and it leaves consumers substantially worse off (with a median decrease of 17.75% in consumer surplus relative to the pre-merger network, compared to 0.95% when all spillovers are considered).

Similarly, the marginal cost spillover exerts a significant, though somewhat smaller, effect. When this spillover is removed—that is, when the merged entity cannot benefit from increased economies of density—the network shrinks to one-half the size of the pre-merger network on median, and consumer surplus falls markedly, with a median reduction of 11.86% relative to the pre-merger network.

The impact of shutting down the fixed cost spillover is less conclusive. In some counterfactual runs, consumer surplus and the merged entity's network size decrease substantially. In those runs, pre-merger fixed costs are sufficiently low to encourage network expansion and allow the merged entity to benefit from higher willingness to pay and economies of density without incurring additional fixed costs, leading to a steep increase in consumer surplus. In other runs, however, pre-merger fixed costs are high enough that the gains from higher willingness to pay and economies of density when expanding the network do not fully offset these costs, prompting the merged entity to reduce its network and resulting in lower consumer surplus.

Overall, Table 15 underscores the crucial role these different spillovers play in shaping the merger's outcome, as well as the importance of accounting for post-merger re-optimization of these spillovers. All three significantly affect the merged entity's post-merger entry and exit decisions, resulting in substantial differences in the merger's impact on consumer surplus.

9 Conclusions

In this paper, we build and estimate a two-stage game model of airline competition in which the airlines choose the networks of markets to serve in the first stage and compete in prices in the second stage. Our model allows for spillovers in entry decisions across markets on the demand, marginal cost, and fixed cost sides. We estimate the model using data from US domestic fares from the third quarter of 2011 and find significant spillovers in entry. We use

Table 15: Outcomes across different scenarios when shutting down the spillovers.

	Before	**Merger**			
		Networks vary - w/ spillovers	Networks vary - w/o spillovers		
			No demand spillovers	No MC spillovers	No FC spillovers
Total Consumer Surplus	2,684.45	-0.95 [-2.22, +0.48] [-5.31, +1.40]	-17.75 [-22.17, -10.28] [-24.49, -10.03]	-11.86 [-13.80, -9.77] [-15.90, -9.22]	-3.24 [-6.24, +9.77] [-6.24, +9.77]
Mean Consumer Surplus	3.91	-1.94 [-3.06, -0.24] [-6.13, +0.52]	-17.25 [-21.48, -10.15] [-23.71, -10.04]	-11.78 [-13.60, -9.64] [-15.53, -9.09]	-3.80 [-6.79, +8.98] [-6.79, +8.98]
Markups: American	119.2	+7.91 [+7.19, +8.49] [+6.35, +9.02]	+0.35 [-1.35, +2.41] [-5.55, +3.39]	+1.26 [+0.29, +2.22] [-1.00, +2.52]	+6.39 [+5.60, +12.62] [+5.07, +12.65]
Markups: Others	116.22	+0.76 [+0.52, +0.92] [+0.41, +1.54]	+5.21 [+3.18, +6.28] [+2.95, +6.75]	+3.87 [+3.47, +4.52] [+3.33, +4.93]	+1.55 [-0.52, +1.84] [-0.54, +2.04]
Segments: American	394	332 [320, 345] [292, 353]	106 [56, 185] [19, 191]	182 [162, 203] [138, 210]	288 [268, 446] [268, 446]
Segments: Others	670	644 [635, 651] [628, 654]	670 [656, 675] [649, 680]	659 [651, 674] [642, 678]	645 [614, 680] [614, 683]

Note: Consumer surplus is calculated using the log-sum formula and is expressed in USD 1 million up to the integration constant. Percentage changes with respect to the pre-merger scenario are reported for total consumer surplus, mean consumer surplus, and markups.

the estimates to counterfactually evaluate the 2013 merger between American Airlines and US Airways. Our counterfactual analysis reveals that, absent any remedies, the merger leads to a decrease in consumer surplus, whereas with the remedies in place, consumer surplus increases markedly. Further, we uncover two important sources of heterogeneity in the merger's impact: first, some hubs enjoy large gains in consumer surplus, while other hubs suffers substantial losses; and second, consumer surplus decreases in markets in which the merging parties served pre-merger and increases substantially in markets where the merged entity enters post-merger.

Hence, the decision of whether or not to allow for the merger depends significantly on which markets the antitrust authority focuses on. Last, we show that such differences are driven by the expansion of American Airlines' network in an attempt to leverage spillovers in entry and the exit of rivals. Overall, these findings have important implications for antitrust authorities because they underlie the relevance of endogenizing post-merger network readjustments and accounting for spillovers when evaluating mergers.

There are several directions for further research. For instance, it would be interesting to consider whether capacity constraints and intertemporal price discrimination may generate dynamics in the pricing strategies of the airlines. Notably, we also abstract away from frequency choices by airlines. Flight frequency is another margin by which the airlines can respond to a merger and that may have a direct impact on marginal costs and consumer utility. We leave these extensions to future work.

References

Aguirregabiria, V. and Ho, C.-Y. (2012). A dynamic oligopoly game of the US airline industry: Estimation and policy experiments. *Journal of Econometrics*, 168(1):156–173.

Aguirregabiria, V., Iaria, A., and Sokullu, S. (2020). Demand models with endogenous product availability and multiple equilibria. Working Paper.

Benkard, C. L., Bodoh-Creed, A., and Lazarev, J. (2020). Simulating the dynamic effects of horizontal mergers: U.S. airlines. Working Paper.

Beresteanu, A. and Molinari, F. (2008). Asymptotic properties for a class of partially identified models. *Econometrica*, 76(4):763–814.

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

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262.

Berry, S. T., Carnall, M., and Spiller, P. T. (1996). Airline hubs: costs, markups and the implications of customer heterogeneity. NBER Working Papers 5561.

Berry, S. T. and Haile, P. A. (2014). Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797.

Berry, S. T. and Jia, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, 2(3):1–43.

Berry, S. T., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Bjornerstedt, J. and Verboven, F. (2016). Does merger simulation work? Evidence from the swedish analgesics market. *American Economic Journal: Applied Economics*, 8(3):125–64.

Bontemps, C., Kumar, R., and Lesellier, M. (2025). Smoothed inference for moment inequality models. Working Paper Toulouse School of Economics.

Bontemps, C., Magnac, T., and Maurin, E. (2012). Set identified linear models. *Econometrica*, 80(3):1129–1155.

Bontemps, C., Remmy, K., and Wei, J. (2022). Ex-post evaluation of the american airlines–us airways merger a structural approach. *Journal of Transport Economics and Policy*, 56(2):129–155.

Chandrasekhar, A. G., Chernozhukov, V., Molinari, F., and Schrimpf, P. (2019). Best linear approximations to set identified functions: With an application to the gender wage gap. NBER Working Papers 25593.

Chen, C. and Mangasarian, O. L. (1995). Smoothing methods for convex inequalities and linear complementarity problems. *Mathematical Programming*, 71(1):51–69.

Chernozhukov, V., Chetverikov, D., and Kato, K. (2018). Inference on causal and structural parameters using many moment inequalities. *The Review of Economic Studies*, 86(5):1867–1900.

Chernozhukov, V., Kocatulum, E., and Menzel, K. (2015). Inference on sets in finance. *Quantitative Economics*, 6(2):309–358.

Cho, J. and Russell, T. M. (2024). Simple inference on functionals of set-identified parameters defined by linear moments. *Journal of Business & Economic Statistics*, 42(2):563–578.

Ciliberto, F., Murry, C., and Tamer, E. (2021). Market structure and competition in airline markets. *Journal of Political Economy*, 129(11):2995–3038.

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

Ciliberto, F. and Williams, J. W. (2014). Does multimarket contact facilitate tacit collusion? Inference on conduct parameters in the airline industry. *The RAND Journal of Economics*, 45(4):764–791.

Crawford, G. S., Shcherbakov, O., and Shum, M. (2019). The welfare effects of endogenous quality choice in cable television markets. *American Economic Review*, 109(3):956–995.

de Paula, A. (2013). Econometric analysis of games with multiple equilibria. *Annual Review of Economics*, 5:107–131.

Eizenberg, A. (2014). Upstream innovation and product variety in the U.S. home pc market. *The Review of Economic Studies*, 81(3):1003–1045.

Fan, Y. and Yang, C. (2024). Estimating discrete games with many firms and many decisions: An application to merger and product variety. *Journal of Political Economy (forthcoming)*.

Fang, Z. and Santos, A. (2019). Inference on directionally differentiable functions. *The Review of Economic Studies*, 86(1):377–412.

Gafarov, B. (2025). Simple subvector inference on sharp identified set in affine models. *Journal of Econometrics*, 249:105952.

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

Graham, B. and de Paula, A. (2020). *The Econometric Analysis of Network Data*. Academic Press.

Ho, K. (2009). Insurer-provider networks in the medical care market. *American Economic Review*, 99(1):393–430.

Holmes, T. J. (2011). The diffusion of wal-mart and economies of density. *Econometrica*, 79(1):253–302.

Houde, J.-F., Newberry, P., and Seim, K. (2023). Economies of density in e-commerce: A study of Amazon's fulfillment center network. *Econometrica*, 91(1):147–190.

Jia, P. and Yuan, Z. (2024). Network competition in the airline industry: An empirical framework. *Rand Journal of Economics (forthcoming)*.

Kaido, H. and Santos, A. (2014). Asymptotically efficient estimation of models defined by convex moment inequalities. *Econometrica*, 82(1):387–413.

Kline, B., Pakes, A., and Tamer, E. (2021). Moment inequalities and partial identification in industrial organization. In *The Handbook of Industrial Organization*, volume 4, pages 345–431. Elsevier.

Kuehn, J. (2018). Spillovers from entry: the impact of bank branch network expansion. *The RAND Journal of Economics*, 49(4):964–994.

Lee, R. S. and Pakes, A. (2009). Multiple equilibria and selection by learning in an applied setting. *Economics Letters*, 104(1):13–16.

Li, S. Y., Mazur, J., Park, Y., Roberts, J. W., Sweeting, A., and Zhang, J. (2022). Repositioning and Market Power After Airline Mergers. *The RAND Journal of Economics*, 53(1):166–199.

Manski, C. P. (2003). *Partial Identification of Probability Distributions*. Springer.

Mazzeo, M. J. (2002). Product choice and oligopoly market structure. *The RAND Journal of Economics*, 33(2):221–242.

Miller, N. H. and Weinberg, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791.

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

Nocke, V. and Schutz, N. (2018). Multiproduct-firm oligopoly: An aggregative games approach. *Econometrica*, 86(2):523–557.

Pakes, A. (2010). Alternative models for moment inequalities. *Econometrica*, 78(6):1783–1822.

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

Park, Y. (2020). Structural remedies in network industries: An assessment of slot divestitures in the american airlines/us airways merger.

Reiss, P. C. and Spiller, P. T. (1989). Competition and entry in small airline markets. *The Journal of Law & Economics*, 32(2):S179–S202.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics*, 37(3):619–640.

Stoye, J. (2021). A simple, short, but never-empty confidence interval for partially identified parameters. CeMMAP Working Papers CWP23/21.

Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *American Economic Review*, 58(1):18–36.

Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406.