

Peters, Craig

Working Paper

Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry

CSIO Working Paper, No. 0033

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Peters, Craig (2003) : Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry, CSIO Working Paper, No. 0033, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/23462>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE CENTER FOR THE STUDY
OF INDUSTRIAL ORGANIZATION
AT NORTHWESTERN UNIVERSITY

Working Paper #0032

**Evaluating the Performance of Merger
Simulation:
Evidence from the U.S. Airline Industry**

By
Craig Peters*
Antitrust Division, US Department of Justice

January 2003

* Antitrust Division, US Department of Justice, Economic Analysis Group; 600 E Street, NW; Suite 10000; Washington DC 20530. E-mail: Craig.Peters@usdoj.gov. Thanks to Robert Porter, Michael Whinston, Shane Greenstein, Aviv Nevo, David Barth, and seminar participants at Northwestern, Duke, Rochester, and the Department of Justice for helpful comments and suggestions; and to John Panzar and Oliver Richard for their assistance in obtaining the data. Financial support from the Center for the Study of Industrial Organization and the Transportation Center is gratefully acknowledged. The views expressed in this paper do not necessarily reflect those of the US Department of Justice.

Abstract

This paper uses merger simulations to predict post-merger prices for six major airline mergers from the 1980's, and compares these predictions with actual post-merger prices. Simulations which incorporate varying degrees of post-merger information allow us to decompose the actual price changes into its component effects. The results suggest that standard simulation methods, which measure the effect of the ownership transfer on pricing incentives, can account for a large component of the post-merger price change, but should not be expected to account for all of it. Changes in marginal cost or firm conduct generally account for most of the remainder, while post-merger entry and changes in observed and unobserved demand-side variables typically have a relatively small effect. A comparison of two alternative demand models indicates that the cross-price elasticities, and hence simulated prices, are quite sensitive to the demand specification.

1 Introduction

Over the past ten years or so, economists and antitrust practitioners have devoted considerable attention to a new methodology for predicting the price effects of horizontal mergers in industries with differentiated products. In its basic form, this methodology combines an estimated structural model of consumer demand with an assumed model of pricing behavior, which allows post-merger equilibrium prices to be “simulated.” By now, there is a large body of literature applying these techniques to hypothetical or actual mergers in a variety of industries. Examples include beer (Baker and Bresnahan (1985), Hausman, Leonard and Zona (1994)); long distance telecommunications (Werden and Froeb (1994)); HMOs (Feldman (1994)); ready-to-eat cereals (Nevo (2000)); and carbonated soft drinks (Dube (2000)). Moreover, merger review by antitrust regulators has increasingly relied on simulations to identify potentially anticompetitive mergers, with the technique featuring prominently in merger cases such as L’Oreal-Maybelline and General Mills-Chex.

Despite this widespread interest in the methodology, however, there has been very little research comparing the predicted price changes from merger simulation with the observed price changes following an actual merger. One motivation for making such a comparison is that, to the extent that simulated prices are offered as forecasts of post-merger outcomes, we would like to know how accurate these forecasts are. Because simulation inevitably relies on strong assumptions about the post-merger behavior of demand, costs, and conduct, a formal comparison of the simulated and actual prices can also be viewed as a test of these assumptions. If these assumptions are violated, it is more appropriate to interpret the simulated price changes as a measure of one component of the merger’s effect on prices, rather than as a forecast of the

overall effect.¹ The availability of post-merger data then allows us to determine the importance of this component relative to the total price change, and also to identify other relevant changes that had been erroneously assumed away. Identifying the variables that most significantly influence post-merger prices is of interest because it may offer guidance on where future efforts to improve predictions are most likely to be fruitful.

The U.S. airline industry of the late 1980's offers an ideal environment to evaluate the performance of merger simulation, for several reasons. First, there was a wave of consolidation during a short period of time, with at least fourteen mergers occurring over a three-year span. Secondly, antitrust enforcement was unusually lax, with every proposed airline merger receiving regulatory approval. As a result, we are able to observe the effects of mergers that might otherwise have been prevented. Third, a rich dataset of prices and quantities is publicly available, with considerable variation across a large cross-section of markets.

Using this data, I aim to answer four broad questions for each of six mergers: (i) *What effect did the merger have on prices?* This issue has been researched by other authors, notably Borenstein (1990) and Kim and Singal (1993). The results presented here are somewhat novel in that I focus on individual cases rather than industry averages (as in Kim and Singal (1993)), and that I do not restrict attention to routes to and from the merging carriers' primary hubs (as in Borenstein (1990)). (ii) *What price changes are predicted by merger simulations?* This can be interpreted as measuring the loss of competition due to merger. (iii) *How sensitive are the predictions to alternative modeling assumptions?* In particular, I will compare simulation with two alternative demand models, as well as prediction with two simple linear models. (iv) *What*

¹ A similar point is made by Baker (1997).

accounts for the difference between the actual and predicted price changes? To answer this, I will combine merger simulation methodology with post-merger data to decompose the observed price changes into individual components. This will allow us to infer the relative contributions of various post-merger changes to the overall price effect.

The remainder of the paper is organized as follows: Section 2 discusses the background on the airline mergers of the 1980's. Section 3 provides details on the simulation methodology, including demand estimation, recovering the implied marginal costs, and post-merger prediction, as well as the linear models that will be considered for comparison. Section 4 describes the dataset and the post-merger changes in price and other observed variables. Section 5 reports the estimated demand parameters, the results of the simulations, and the decomposition of the price changes into their component effects. Section 6 concludes with a brief summary and suggestions for further research.

2 Background on the Airline Mergers of the 1980's

For each of six major airline mergers from the late 1980's, the empirical analysis below will examine price changes in city-pair markets where the two merging carriers competed with each other prior to the merger. This focus on markets with pre-merger overlap is not based on a presumption that the mergers' effects in other markets were negligible: Merger-induced changes in cost, quality, or conduct might very well spill over across markets. For example, for the Northwest-Republic and TWA-Ozark mergers, Brueckner, Dyer, and Spiller (1992) use regressions of price on network variables to predict significant price decreases in markets where only one of the merging firms offered service prior to the merger, due to post-merger economies

of traffic density. By contrast, standard methods of simulating horizontal mergers will generally capture only those price changes which stem from a loss of within-market competition. Because this research is motivated by an interest in the performance of these methods, I restrict attention to markets where both firms had a significant pre-merger presence.

For all of the cases, I will define the “pre-merger” period as the four quarters of 1985, pre-dating the January 1986 announcement of Northwest’s bid to acquire Republic. Using a single time period for all of the mergers simplifies the analysis, by making it unnecessary to estimate the pre-merger demand model more than once. The “post-merger” period is then defined as the first four quarters after the operations of the merged firms were fully integrated.

The first two mergers examined here, Northwest-Republic (NW-RC) and TWA-Ozark (TW-OZ), have been the subject of considerable attention in the economics literature (see Borenstein (1990), Werden, Joskow, and Johnson (1991), Beutel and McBride (1992)). This is largely because in both cases, the merging airlines shared at least one major hub airport (Minneapolis and Detroit in the former case, St. Louis in the latter), and because both mergers had been approved by the Department of Transportation over the objections of the Department of Justice.

The third case, Continental-People Express (CO-PE), implicitly includes two other mergers: In late 1985, People Express acquired Frontier Airlines, but their operations were not integrated until they both merged with Continental. At the same time, the operations of New York Air were integrated with Continental’s, although this merger did not involve a transfer of

ownership, since both airlines were subsidiaries of Texas Air². Thus, when I compare pre-merger and post-merger prices for the Continental-People Express merger, the set of pre-merger prices includes prices charged by Frontier and New York Air. Similarly, when I define the set of markets where Continental and People Express competed against one another, I actually identify markets where Continental *or* New York Air competed with People Express *or* Frontier. This is relevant because People Express's Newark hub allowed it to compete with New York Air in a number of markets, while Frontier and Continental both had major hubs at Denver.

It should also be noted that this merger is unique among the six considered here in that the acquired airlines were both suffering serious financial difficulties during the period prior to the merger. As discussed by Busse (2002), there is a strong correlation between financial distress and the likelihood of initiating a fare war in this industry, suggesting that firm conduct may play an important role in determining the effects of mergers involving failing airlines. As noted in the next section, I will formally restrict attention to studying the implications of static models with non-cooperative price-setting conduct. In interpreting the results, however, I will highlight areas where this assumption seems likely to be violated, including the case of highly aggressive pricing by financially distressed firms.

In each of the next two cases, a major carrier expanded its network into the western U.S.: Delta-Western (DL-WA) and American-Air Cal (AA-OC).³ Because both acquiring firms were predominantly based in the East and Midwest prior to the mergers, there were relatively few city-

² Texas Air Corp. was not a separate airline during this time period, but was essentially a holding company for Continental, New York Air, and eventually Eastern.

pair markets with significant pre-merger competition between the merging airlines. A more extreme case is the merger between USAir and PSA, in which the pre-merger route structures had no significant overlap. As a result, this case is excluded from the analysis here.

The sixth and final case, USAir-Piedmont (US-PI), valued at \$1.6 billion, remains the largest airline merger in U.S. history. While the two airlines did not share any hubs, and their networks were largely complementary (USAir's was based in the Northeast, Piedmont's in the Southeast), the extent of premerger overlap was comparable to that of Northwest-Republic and TWA-Ozark. In part, this overlap reflects Piedmont's February 1986 acquisition of Syracuse, N.Y.,-based Empire Air. As with the Continental-People Express merger, I define Piedmont's pre-merger (1985) prices and markets to include Empire's. The USAir-Piedmont merger also stands out for the unusually long period of time between the completion of the acquisition (November 1987) and the full integration of operations (August 1989). (See Kole and Lehn (2000) for a detailed discussion of the difficulties that contributed to this delay).

One major merger which I do not examine here is Texas Air's acquisition of Eastern. This case is anomalous in that the operations of the acquired airline were not integrated into those of the acquirer (or its principal subsidiary, Continental). Rather, Eastern remained operationally independent until its bankruptcy and liquidation in the early 1990's. The sheer scale of the acquired firm's network (Eastern was one of the four largest airlines in the U.S. in the early 1980's) also makes this something of a special case, so for simplicity I will restrict attention to the six cases described above.

³ I use the FAA two-letter airline codes to abbreviate each merger. Usually, these are quite predictable, but occasional anomalies arise, such as "OC" for Air Cal.

Two smaller airline mergers during this period are also worth mentioning: In March, 1985, Southwest announced its bid to acquire Muse Air, and completed the acquisition three months later. There was considerable pre-merger overlap between the two airlines, so that ideally this case would be included here. However, my current dataset only extends as far back as the first quarter of 1985, and hence does not include an adequate pre-merger period for this merger. Finally, Braniff's 1988 acquisition of Florida Express is excluded for the same reason as USAir-PSA, the absence of any pre-merger overlap.

For each of the six mergers of interest, I identify the set of markets where both carriers provided service during all four quarters of 1985. Also, I exclude from this set any markets where, in any quarter of 1985, the merger would be considered “unlikely to have adverse competitive consequences” under the 1992 DOJ/FTC Horizontal Merger Guidelines.⁴ The markets excluded under this criterion are generally those where at least one of the merging airlines has less than 5% market share. Table I indicates the number of overlap markets identified in this way, along with the dates that the acquisition and operational integration were completed, for each merger. An obvious point to emerge from this table is that, with only two overlap markets, the American – Air Cal merger by itself should not be expected to yield much in the way of statistically significant results. I will include it in the analysis as an additional case study, essentially for completeness.

⁴ A market satisfies this condition if the post-merger Herfindahl-Hirschman Index (HHI) is less than 1000; if the post-merger HHI is less than 1800 and the change in the HHI is less than 100; or if the change in the HHI is less than 50. For the purposes of computing the HHI, I use market shares based on revenues, although very similar results are obtained with passenger shares.

Table I**Mergers with Significant Pre-merger Overlap**

Merger:	Acquisition Completed:	Operations Integrated:	# overlap markets:
NW-RC	08/86	09/86	78
TW-OZ	09/86	10/86	50
CO-PE	12/86	02/87	67
DL-WA	12/86	03/87	11
AA-OC	04/87	07/87	2
US-PI	11/87	08/89	60

3 Simulation Methodology

In general, merger simulation consists of three components: estimation of consumer demand using pre-merger data; recovery of pre-merger marginal costs under a maintained assumption about firm conduct; and prediction of post-merger prices under specific assumptions about what does and does not change as a result of the merger. This section describes each of these components in turn.

3.1 Demand Estimation

To estimate the demand relationship, I use a familiar discrete choice framework, based on an underlying random utility model. Defining a market as an ordered origin-destination city-pair (where a city is a Metropolitan Statistical area, possibly including more than one airport), and a product as an airline-itinerary combination⁵, I specify consumer utility as a simple function of five observable product characteristics: price, flight frequency, airport presence, whether the

itinerary is nonstop, and distance. Price is computed as a passenger-weighted average of all fares reported for a given product. Flight frequency on a nonstop route is the number of quarterly flights offered by the airline on the outbound portion of the itinerary; on a connecting route, it is the geometric mean of the frequencies on each of the outbound segments. Airport presence is the airline's share of all traffic at the origin airport, and is included in the utility function to reflect the observation that "the dominant airline at an airport attracts a disproportionate share of the traffic that originates at the airport." (Borenstein (1991))⁶ Consumers will generally prefer nonstop itineraries over connecting itineraries, other things being equal, and within any market, we would also expect travel over shorter distances to be preferred.

Adopting a log-linear functional form yields the following expression for the utility of consumer i from product j at time t :

$$u_{ijt} = \beta_p \ln p_{jt} + \beta_f \ln f_{jt} + \beta_{ap} \ln Apt\ Pres_{jt} + \beta_n Nonstop_j + \beta_{mi} \ln Miles_j + \mu_{jt} + \varepsilon_{ijt} \quad (1)$$

where μ_{jt} = unobserved product quality

ε_{ijt} = the consumer-specific deviation from mean utility⁷

Also, it will generally be desirable to control for components of unobserved product quality which are constant within significant subsets of the data. Accordingly, I include fixed effects for airlines, markets, time periods, and airports (in multi-airport cities). Formally:

⁵ An itinerary is defined as an ordered sequence of airports through which a passenger travels.

⁶ Explanations for this finding include consumer preferences for familiarity, economies of scale in local advertising, and the use of marketing devices such as frequent flyer programs and travel agent commission overrides.

⁷ The use of the term "mean utility" to denote the component of utility that does not vary across consumers, due to Berry (1994), is actually a bit of a misnomer, since the extreme value distributions that are typically assumed for ε

$$\mu_{jt} = \text{Airline}_j \alpha_{al} + \text{Time}_t \alpha_T + \text{Mkt}_j \alpha_m + \text{OrigApt}_j \alpha_o + \text{DestApt}_j \alpha_d + \xi_{jt} \quad (2)$$

where the variables in this expression are vectors of indicator variables.

For the purposes of estimating the demand parameters, there is no reason to distinguish between the coefficients on the observed characteristics (β) and the fixed effects (α). However, when we begin using these estimates to simulate post-merger outcomes, an important distinction arises. Because β measures consumers' preferences for specific attributes, it seems reasonable to assume that these parameters do not vary over time in general, and are not affected by mergers in particular. In contrast, α measures the utility contribution of unobserved characteristics which are fixed during the pre-merger period. Precisely because these characteristics are unobserved, it seems less appropriate to make strong assumptions about their post-merger behavior. This is especially true of α_{al} , the vector of airline-specific fixed effects: the unobserved component of product quality that is specific to a particular airline may very well change when that airline merges with another. Accordingly, throughout this paper, I will maintain the assumption that β is fixed at the pre-merger estimates, but will explicitly consider possible changes in μ .

Different assumptions about the distribution of ε correspond to different ways of modeling heterogeneity in consumer preferences. A simple form frequently used in antitrust analysis is the logit model, in which ε is assumed to be i.i.d. extreme-value. More general models include the nested logit; Bresnahan, Stern, and Trajtenberg's (1997) Principles of Differentiation (PD) GEV model; and Berry, Levinsohn, and Pakes' (1995) random coefficients

do not have mean zero. "Mode utility" would be more precise, but I will stick to convention and use the standard term.

model. For the simulations described below, I compare the results from two models: a simple nested logit, with a nest for the “outside good”; and a “nested PD-GEV” model – a generalization of Bresnahan, Stern, and Trajtenberg’s model. This model includes an outside good nest and two PD-type market segments: one which groups itineraries according to whether they are nonstop or connecting, and another which groups them based on distinct airport-pairs within a market (if either endpoint city has more than one airport)⁸. These groupings reflect the intuitive notion that products within the same group are closer substitutes for each other than products in different groups.

Given an assumed distribution of ε (defined by a vector of estimable parameters, ρ), and the assumption that each consumer purchases the product which maximizes her utility, product j ’s share of potential market sales can be expressed as a function of the characteristics for all products in the market. For the nested PD-GEV model estimated here, this function is:

$$s_j = \left\{ \frac{a \cdot \exp(\delta_j / \rho_{AP})}{\left[\sum_{j' \in AP_j} \exp(\delta_{j'} / \rho_{AP}) \right]^{1-\rho_{AP}/\rho_0}} + \frac{(1-a) \cdot \exp(\delta_j / \rho_{NS})}{\left[\sum_{j' \in NS_j} \exp(\delta_{j'} / \rho_{NS}) \right]^{1-\rho_{NS}/\rho_0}} \right\} \cdot$$

$$\frac{\left\{ a \sum_{AP} \left[\sum_{j' \in AP} \exp(\delta_{j'} / \rho_{AP}) \right]^{\rho_{AP}/\rho_0} + (1-a) \sum_{NS} \left[\sum_{j' \in NS} \exp(\delta_{j'} / \rho_{NS}) \right]^{\rho_{NS}/\rho_0} \right\}^{-(1-\rho_0)}}{1 + \left\{ a \sum_{AP} \left[\sum_{j' \in AP} \exp(\delta_{j'} / \rho_{AP}) \right]^{\rho_{AP}/\rho_0} + (1-a) \sum_{NS} \left[\sum_{j' \in NS} \exp(\delta_{j'} / \rho_{NS}) \right]^{\rho_{NS}/\rho_0} \right\}^{\rho_0}}$$

⁸ Although cities with more than one airport are relatively rare (9 out of 114 cities in my data), these tend to be large cities which appear disproportionately in the data. In the full 1985-1990 sample, 66% of the observations (and 70% of the passengers) are in markets where travel is available to or from more than one airport in at least one of the endpoint cities.

where $0 < \rho_{AP} \leq \rho_0 \leq 1$, $0 < \rho_{NS} \leq \rho_0 \leq 1$, $a = \frac{\rho_0 - \rho_{AP}}{2\rho_0 - \rho_{AP} - \rho_{NS}}$,

$\delta_j = u_{ij} - \varepsilon_{ij}$ for any i (i.e., the mean utility from product j)

AP_j = the set of products serving the same airport-pair as product j

NS_j = the set of products in the same nonstop/connecting group as product j

The ρ parameters are inversely related to the strength of the correlation between the consumer-specific preferences across products within the same cluster. When $\rho_0 = 1$, this expression reduces to the PD-GEV model of Bresnahan, Stern, and Trajtenberg (1997), which can be interpreted as a weighted average of two nested logits, with one nest for airport-pairs, the other for nonstop/connecting. Alternatively, the simple nested logit with a nest for the outside good is obtained when $\rho_{AP} = \rho_{NS} = \rho_0$.

Equating the market shares defined by this function with observed quantities requires an estimate of the total market potential, or equivalently, of the share of the outside good (i.e. the probability that a potential consumer chooses not to purchase any of the available products in the market). Two distinct approaches to obtaining such an estimate are available. First, Werden and Froeb (1994) suggest estimating (or assuming) a value for the aggregate elasticity, defined as the percentage change in total sales implied by a marginal percentage change in the prices of all products in the market. For any set of parameter values, then, this aggregate elasticity implies a particular share for the outside good, which turns out to be simple to compute, even for complicated GEV models.⁹ The second approach, suggested by Berry (1994) and widely used in

⁹ For both models estimated in this paper, the outside good's share equals the ratio of the aggregate elasticity to β_p .

the subsequent literature, is to assume some relationship (which may or may not include parameters to be estimated) between the total market potential and observed data on market population. The Werden-Froeb method may be more appropriate in applications where we are more confident in our ability to estimate aggregate elasticities than in the accuracy of the assumed relationship between market potential and population, while Berry's method would be preferable if the reverse is true. A detailed comparison of the two methods has not, to my knowledge, been attempted, and is beyond the scope of the present paper. In any case, I use Berry's approach, assuming that the market potential for any city-pair is equal to the geometric mean of the population of the endpoint cities.

In order to estimate the parameters of the utility function specified above, we need to account for potential correlation between the observed and unobserved characteristics. Price, flight frequency, and airport presence are all endogenous, and we will therefore need to find instruments for them. The basic identifying assumption that allows us to construct a set of instruments is that each airline's overall network structure is exogenous. That is, the set of routes served by each airline is assumed to be independent of the unobserved component of utility (ξ). This is analogous to the standard assumption in this literature, that variation in the set of available products across markets and over time is exogenous.

To construct instruments for flight frequency and airport presence, I exploit the overlap between itineraries in different markets. In particular, I expect the number of flights on a segment to increase with the number of distinct itineraries which include that segment, and with the average population of the endpoint cities across all overlapping itineraries on a segment.

Similarly, an airline's airport presence will increase with the total number of itineraries offered by the airline out of the airport, and with the fraction of those itineraries which are nonstop. Moreover, it will decrease with the total number of itineraries offered by rival airlines out of the airport, and with the fraction of those which are nonstop. Finally, to construct instruments for price, I use numbers of products and means of product characteristics within markets and within PD-type groupings, loosely following Bresnahan, Stern, and Trajtenberg (1997).¹⁰ The intuition behind the expected correlation between these instruments and price is that, in the absence of perfectly collusive conduct, price will tend to fall with an increase in the number of alternative products, the number of relatively close substitutes, or the attractiveness of those alternatives' characteristics. Given this set of instruments, estimation follows Berry's (1994) procedure.

3.2 Firm Conduct and Marginal Cost

Once we have estimated the parameters of the demand function, an assumption about firm conduct is sufficient to allow us to recover marginal costs. The standard approach in the literature on merger simulation with differentiated products is to assume static Nash Bertrand conduct: firms choose prices non-cooperatively to maximize a short-run profit function. The widespread use of this assumption is partly due to its mathematical tractability, and partly due to the belief that collusion is more difficult to sustain when products are differentiated (see, for instance, Werden and Froeb (1994)). Formally, then, I will maintain this assumption. As will be noted in the discussion of the results below, if the Bertrand assumption is violated, the inferred changes in marginal costs may in fact reflect changes in conduct, or errors in the estimated markup due to misspecified conduct. Explicitly considering the implications of alternative

¹⁰ Since some of the available product characteristics are themselves endogenous (flight frequency and airport presence), I instead use means of their corresponding instruments.

models of firm behavior, including quantity-setting (Cournot) and collusive conduct, as well as attempting to identify mergers' effect on conduct, would be interesting areas for future research.

Given the Bertrand assumption, the first-order conditions for firms' profit maximization decision allow us to solve for marginal cost. For product j :

$$q_j + \sum_{k \in S_j} (p_k - c_k) \frac{\partial q_k}{\partial p_j} = 0 \quad (3)$$

where p , q , and c are prices, quantities, and marginal costs, respectively; and S_j is the set of all products produced by the same firm as product j . Price and quantity are observed in the data, and the derivatives of the demand function can be computed from the estimated demand parameters and product characteristics (both observed and unobserved)¹¹. For a market with N products, this yields a system of N linear equations, which can easily be solved for c :

$$c = p + \Delta^{-1} q \quad (4)$$

where Δ is an $N \times N$ matrix such that $\Delta_{jk} = \frac{\partial q_k}{\partial p_j}$ if $k \in S_j$, 0 otherwise.

Note that marginal costs are essentially computed as a residual: the difference between observed prices (p) and estimated markups ($-\Delta^{-1}q$). Just as the unobserved product quality, μ , measures the component of consumer utility that is unaccounted for by observed characteristics, the marginal cost, c , represents the component of price that is unaccounted for by the estimated markup. As a result, it seems reasonable to treat the post-merger behavior of marginal costs in a way similar to that of unobserved product quality. While I maintain the Bertrand conduct assumption throughout, I will explicitly consider possible post-merger changes in c .

3.3 Post-Merger Prediction

Once we have estimated the demand parameters, unobserved product quality, and marginal cost, we can use these estimates to predict the post-merger equilibrium prices. The approach developed in the merger simulation literature (Berry and Pakes (1993), Werden and Froeb (1994)) involves changing the ownership structure in each market, while holding the set of products and all of their characteristics (including μ and c) constant. Formally, the changed ownership structure affects the post-merger pricing equations by redefining the set S_j for all products j that were initially produced by either of the merging firms. For each market, the first-order conditions given in equation (3), together with the estimated demand function, define a system of non-linear equations in price. Standard numerical methods allow this system to be solved for the post-merger equilibrium prices.

As pointed out by Baker and Bresnahan (1985), this approach does not distinguish a merger between two firms from a bilateral collusive arrangement between them. Hence, the price changes generated by this type of simulation can be viewed as measuring the loss of competition associated with the changed incentives of the combined firm. Competitive incentives, however, may not be the only change prompted by the merger. If other changes also influence the equilibrium outcome, then the simulated prices will generally not be equal to the actual prices, even after correcting for changes in the overall price level. This raises the question of what these other merger-induced changes are. I will explicitly consider two broad classes of post-merger change: changes that can be directly observed *ex post*, and those that cannot. The former category includes demand-side observables, particularly flight frequency and airport

¹¹ The analytical derivatives for the market share function are reported in Appendix I.

presence; and the set of available products, which is influenced by the integration of the two merging airlines into a single firm, by entry or exit, and by any changes in the available itineraries within a particular market. Potential changes which are not directly available in the post-merger data include the unobserved component of product quality, μ ; marginal cost, c ; and, in principle, firm conduct. As noted above, for the purposes of the formal simulations, I maintain the assumption of non-cooperative price-setting conduct.

In order to measure the extent to which the observed changes contribute to the differences between simulated and actual post-merger prices, I perform the following “updated simulation”: For each quarter in the post-merger period, I take the *actual* data on the set of available products, flight frequency, and airport presence. Given this post-merger data, along with the estimated demand parameters and the conduct assumption, the only requirements for simulating equilibrium prices and market shares are values of μ and c for each product. Since the aim here is to measure the effect of the observed changes alone, it is desirable to hold the unobserved components constant at pre-merger levels.¹² However, changes in the set of products require that some of these values be estimated. For any itinerary offered post-merger by the merged airline, and also offered pre-merger by either or both of the merging airlines, the post-merger values of μ and c are computed as the passenger-weighted means of the pre-merger values. For itineraries offered by other airlines, both before and after the merger, post-merger μ and c are simply the average over the four pre-merger quarters. Finally, for products which are new to the post-merger sample (due to entry, for instance), I estimate post-merger μ and c as the out-of-sample

¹² At this point, I am allowing for the *demand-side* effects of changes in flight frequency and airport presence. Cost effects of these changes, especially the former, are likely to occur as well. These are implicitly accounted for when I compute updated marginal costs.

predictions from linear regressions of pre-merger values on route distance and an indicator for nonstop itineraries, with airport, market, and airline fixed effects¹³. Given these values for the unobservables, I use the standard numerical procedure to simulate post-merger prices. A comparison of these prices with those obtained from the previous simulation (where only ownership was changed) will indicate the effect of the observed changes on prices, while a comparison with the actual post-merger prices will indicate the effect that is still unaccounted for – i.e., that is due to changes in the unobserved quantities.

Finally, we may wish to determine the relative importance of changes in costs and product quality. Given the maintained conduct assumption, and holding the demand parameters (β and ρ) fixed at the pre-merger estimates, it is straightforward to recover the changes in μ and c from post-merger data on prices, characteristics, and market shares. In particular, we first use the Berry inversion to transform market shares into mean utilities (for any value of ρ , there is a unique vector of mean utilities δ which equates the predicted and observed market shares). Next, using β , we recover μ by subtracting the observed component of mean utility (i.e., the linear combination of price and product characteristics defined by equation (1)) from δ . Once we know μ , marginal cost can be recovered as before, as the difference between price and the estimated markup defined by equation (4), where the markup is computed using the recovered post-merger mean utilities. Finally, to evaluate the relative effect of the unobserved changes on post-merger

¹³ For any pair of merged airlines, the pre-merger values used in these regressions are averages across both firms, so that a single fixed effect is used for each merged airline.

prices, I perform another simulation, with mean utility updated to the post-merger values, while marginal cost is held constant.¹⁴ The results of these comparisons are discussed in Section 4.

3.4 Linear Prediction

An alternative to the formal simulation methods described above is to estimate a simple linear regression of price on market structure variables. The estimated coefficients on the relevant variables can then be used to predict post-merger prices under an assumed merger-induced change in market structure. This approach is often viewed as problematic for two reasons. First, extrapolating from an empirical relationship between price and concentration implicitly relies on strong assumptions about the underlying supply equation. Secondly, the assumed change in market structure typically involves combining the pre-merger market shares of the merging firms, without accounting for the effect of increased prices on post-merger sales. (See Farrell and Shapiro (1990) for a discussion of these concerns). Nonetheless, the approach remains popular among antitrust practitioners due to its simplicity. In order to evaluate the relative performance of this methodology in the present context, I use the 1985 data to estimate the following linear regression for the price of product j in market k at time t :

$$\ln p_{jt} = \beta_H HHI_{kt} + \beta_m \ln Miles_j + \beta_n Nonstop_j + \alpha_t + \epsilon_{jt}$$

where α_t is a time fixed effect.

¹⁴ I have also computed the reverse, simulating with c updated and μ held constant. The results are essentially the same. By construction, if *both* μ and c are updated, then the simulated prices will be identical to the actual prices.

Under this model, the predicted percentage price change for a merger between two firms with pre-merger market shares s_1 and s_2 is:

$$\% \Delta p = \exp(2 \beta_H s_1 s_2) - 1$$

Because the predictions depend critically on the estimated value of β_H , it is important to consider possible sources of bias in this coefficient. In particular, HHI may be correlated with the error term, depending as it does on market shares. Accordingly, I estimate the model two ways: first, by ordinary least squares; second, by two-stage least squares, with the log number of airlines serving the market as an instrument for HHI . As shown below, the results are quite sensitive to whether we instrument for HHI .¹⁵

4 Data and Observed Post-Merger Changes

In this section, I provide details on the data used in estimation, and discuss the observed changes in price and other observed variables.

4.1 Data Description

The data used here is drawn primarily from the DOT's Origin and Destination (O&D) survey, a 10% sample of all tickets sold to domestic U.S. airlines each quarter. Information on flight frequency is from the DOT's Service Segment Databank. The dataset is quarterly, covering the “pre-merger” period (1985), and the first year post-merger for each of the six mergers. Each observation corresponds to an airline-itinerary-date combination, with total passengers and a passenger-weighted average price for each product. The sample is restricted to round-trip coach-class travel between a set of 131 major airports, with at most one connection in each direction. A

¹⁵ By contrast, the inclusion of other controls, such as the population of the endpoint cities, flight frequency, and airport presence, had no significant effect on β_H (and hence the predicted price changes).

market is defined as a directional city-pair, where airports are grouped into cities according to Census Bureau definitions of Metropolitan Statistical Areas (MSA). City-pair markets with fewer than 50 passengers in any quarter are removed from the data, as well as observations with fewer than five passengers or with extremely low or high prices. Also, I remove itineraries with a connection that cannot plausibly be a simple layover,¹⁶ and routes where a reporting airline does not report at least five flights during the quarter.¹⁷ Table II provides some summary information about the samples for the pre-merger and post-merger periods.

Table II
Summary Statistics

	Pre-		Post-merger				
	merger	NW-RC	TW-OZ	CO-PE	DL-WA	AA-OC	US-PI
Time Period	1985 Q1 – Q4	1986 Q4 1987 Q3	1987 Q1 – Q4	1987 Q1 – Q4	1987 Q2 1988 Q1	1987 Q3 1988 Q2	1989 Q4 1990 Q3
Airlines	37	34	28	28	26	24	17
Cities	114	114	114	114	114	114	114
Airports	130	130	130	130	130	130	131
Markets	1,758	1,758	1,758	1,758	1,758	1,758	1,758
Airportpairs	3,418	3,607	3,615	3,615	3,632	3,659	3,666
Itineraries	10,172	12,542	12,759	12,759	12,861	13,012	13,606
Observations	42,613	49,046	49,783	49,783	50,132	50,253	50,433
Passengers	4,780,204	5,418,447	5,436,686	5,436,686	5,471,247	5,454,820	5,503,304

¹⁶ For example, JFK-LAX-BOS is not a connecting flight in the New York - Boston market.

¹⁷ The Service Segment Data from this time period does not include flight frequency for the smallest or newest airlines. In these cases, I use an imputed estimate of flight frequency based on other observables.

4.2 Actual Post-merger Price Changes

As described in Section 2 above, I focus on the set of markets where each of the two merging airlines faced significant competition from the other (“overlap markets”). The reasons for this focus are, again, that these are the markets where we would expect a merger’s effects to be most pronounced, and hence somewhat easier to disentangle from other changes influencing prices; and also that, in practice, these markets would generally be the predominant focus of a legal or regulatory challenge to this sort of merger. The merger may have effects on prices in other markets if the cost structure or overall service quality of the combined firm is altered; if multimarket contact facilitates collusion; or if firms’ pricing decisions involve dynamic considerations, such as entry deterrence, that may be influenced by merger. The existing models used for merger simulation, however, are not well-suited to predicting these other changes, so the focus on overlap markets seems appropriate.

To measure the apparent effect of the mergers on prices in these markets, I define the “relative price change” for each market as the difference between the actual percentage price change and an average percentage price change, conditional on route distance, across all markets within each quarter. The motivation here is to remove the portion of the observed price change which is properly attributed to other sources of variation, such as overall inflation, fluctuations in fuel prices, wages, interest rates, and other factor prices, demand seasonality, and growth over time. To the extent that the price effects of these other changes may depend on route distance, it seems appropriate to compute a conditional mean as the basis for comparison. The approach used by Kim and Singal (1993) and Borenstein (1990) is to compute the mean across the subset of routes which are similar in distance, or equivalently, to estimate the conditional mean non-

parametrically, using running-mean smoothing. Alternatively, we can estimate a parametric conditional mean, specifying a log-linear relationship between price and distance. Figure I (in Appendix II) compares the predictions from this simple parametric estimator with those from a locally-weighted running-mean with bandwidth 0.2, using the data from the first quarter of 1985. The reasonably close fit between these two predictions ($R^2 = 0.97$) suggests that relatively little is lost by using the less computationally intensive regression approach.

A concern with this definition of the relative price change is that the overall industry average may itself be influenced by the merger. Ideally, we might define a “control group” of markets unaffected by merger, and compute the average across this subsample. The problem with this, however, is that there are only five markets in the dataset which are not served by any of the airlines involved in the six mergers here, all of which are monopoly markets served by United. Alternatively, we might allow the control group for a given merger to include other firms which also merged during this time period (as in Kim and Singal (1993)), or to include markets where there was no pre-merger overlap between the merging airlines. While these might offer a slightly more accurate measure of the isolated effect of the merger, I do not expect that using the industry average would produce significantly different results, at least qualitatively. In any case, I will measure price changes relative to the industry as a whole, rather than relative to a control group.

Formally, the relative price change in market m for a given merger is:

$$\% \Delta rel p_m = \frac{\bar{p}_m^{post} - \bar{p}_m^{pre}}{\bar{p}_m^{pre}} - \frac{\hat{p}_m^{post} - \hat{p}_m^{pre}}{\hat{p}_m^{pre}}$$

where \bar{p}_m^{pre} and \bar{p}_m^{post} are the passenger-weighted averages across all pre- and post-merger prices, respectively, for itineraries offered by the merging airlines in market m ; and where \hat{p}_m^{pre} and \hat{p}_m^{post} are the passenger-weighted averages of the predicted prices (from the log-linear regression of price on distance) for market m across each of the four quarters of the pre- and post-merger periods, respectively. Once I have computed $\% \Delta rel p_m$ for each market, the “average relative price change” for each merger is simply the mean of $\% \Delta rel p_m$ across a specific set of markets. Table III reports the average relative price changes for each of the six mergers, across all overlap markets, and also across the set of “duopoly markets,” defined here as markets where the two merging firms had a combined pre-merger share of at least 90%, each with at least 10%.

Table III
Average % Relative Price Change in Overlap Markets
1985 vs. 1st Year Post-Merger

	Overlap		Duopoly	
	#	% $\Delta relp$	#	% $\Delta relp$
NW-RC	78	7.2 (0.81)	16	11.0 (1.69)
TW-OZ	50	16.0 (2.26)	17	19.5 (7.06)
CO-PE	67	29.4 (3.59)	0	- -
DL-WA	11	11.8 (3.56)	1	26.6 -
AA-OC	2	6.5 (4.62)	0	- -
US-PI	60	20.3 (2.63)	14	21.8 (5.42)

Standard errors for the mean, under normality, are in parentheses.

The key points to note from these results are: (i) As we would expect, the average price change across all overlap markets is positive for all cases, significantly so for all except American – Air Cal. (ii) The magnitude of the changes are economically significant, but vary considerably across cases, ranging from less than 8% to almost 30%. (iii) Again as we would expect, for the three cases with significant numbers of duopoly markets, the average price changes are larger within this subsample, although the difference is only statistically significant for Northwest – Republic. (iv) Relative to the other major cases, Northwest-Republic stands out as having the lowest average price change. This is somewhat surprising, considering that two hubs were shared by the merging airlines. (v) Also notable is the unusually high average price

change following the Continental – People Express merger, despite the absence of any duopoly markets. This finding may be partially accounted for by the financial difficulties experienced by the acquired airlines, as discussed previously.

4.3 Other Observed Post-Merger Changes

In addition to price, there are several other relevant changes which may be observed in the post-merger data. In particular, changes in flight frequency and airport presence are likely to have demand-side effects. Also, the integration of the merging carriers' operations, as well as any entry or exit precipitated by the merger, implies that the set of products available post-merger will not be identical. As described above, I will explicitly consider how these observed changes affect prices by performing simulations with updated post-merger data. First, though, I examine what these other changes were.

Table IV reports the percentage change in flight frequency for segments which fall within the overlap markets for each merger. In addition to reporting the overall average for each merger, I also divide the set of segments into two groups: those where both carriers operated flights pre-merger, and those where only one did. The change in flight frequency for each

segment s is computed as: $\Delta f_s = \frac{\bar{f}_s^{post} - \bar{f}_s^{pre}}{\bar{f}_s^{pre}}$, where \bar{f}_s^{post} = the average quarterly number of

flights offered by the merged carrier on segment s during the post-merger period, and \bar{f}_s^{pre} = the average quarterly number of flights offered by both carriers combined on segment s during the pre-merger period. This measure does not correct for overall industry changes in flight frequency, so I also report the average expected change in frequency for each group of segments.

As with the computation for the average relative price, the expectation is based on an industry-wide conditional mean, where in this case, the conditional mean is specified as a log-linear regression of frequency on the total segment flights in 1985 and the total number of airlines serving the segment in 1985. In other words, the “industry average” ($E(\% \Delta f_s)$) for each group of segments reported in Table IV represents the average change in flight frequency over the same time period for all segments with a similar number of pre-merger flights and carriers. The net effect of a merger on flight frequency ($\text{Net} \% \Delta f_s$) is then simply the difference between the absolute percentage change and the corresponding industry average.

Table IV
Average % Change in Flight Frequency in Overlap Markets
1985 vs. 1st Year Post-Merger

	All segments in overlap mkts				Overlap segs		Non-Overlap	
	#	$\% \Delta f_s$	$E(\% \Delta f_s)$	$\text{Net} \% \Delta f_s$	#	$\% \Delta f_s$	#	$\% \Delta f_s$
NW-RC	103	7.0 (5.0)	17.0 (2.4)	-10.0 (4.9)	60	-17.2 (3.2)	43	40.8 (8.9)
TW-OZ	49	-4.4 (4.0)	14.8 (2.1)	-19.2 (3.4)	39	-14.7 (3.0)	10	36.1 (6.8)
CO-PE	78	0.9 (6.5)	8.2 (2.0)	-7.3 (5.3)	39	-32.1 (2.1)	39	33.8 (10.5)
DL-WA	30	23.0 (8.6)	15.7 (3.1)	7.3 (8.4)	4	29.1 (10.6)	26	22.0 (9.9)
AA-OC	6	40.2 (22.6)	18.2 (2.5)	22.0 (24.4)	2	-23.1 (6.8)	4	72.0 (15.8)
US-PI	128	34.7 (7.0)	32.5 (3.0)	2.3 (6.2)	29	-5.6 (7.5)	99	46.6 (8.5)

Standard errors for the mean, under normality, are in parentheses.

Based on the overall changes, the effect of the mergers on flight frequency seems somewhat ambiguous. While flight frequency increased significantly on routes served by Delta-Western and USAir-Piedmont, these changes can be fully accounted for by industry-wide trends. Relative to the industry as a whole, the only significant changes were relative decreases for Northwest-Republic and TWA-Ozark. When we consider the two subsamples separately, however, a clearer pattern is evident. On segments where both carriers operated pre-merger, there is a strong tendency to reduce frequency post-merger, as we would expect. When the direct competition on a segment is reduced, the incentive to maintain high levels of capacity declines as well.

One notable exception is the significant increase on the small number of overlap segments for Delta-Western. This case illustrates how a merger between two networks with relatively little overlap can have pro-competitive effects even in markets where the two firms competed directly with each other. Consider one of the overlap segments for Delta and Western: Salt Lake City-to-Dallas. With Salt Lake City as its major hub, Western used this segment for Dallas-bound passengers originating on the West Coast. Similarly, Delta offered connections in Dallas for passengers traveling from Salt Lake City to destinations in the Northeast. After the merger, this segment served to link the two networks, so that passengers on transcontinental itineraries might pass through both these hubs. As a result, Delta introduced large increases in the number of flights on this segment, in both directions, accounting for the observed frequency increases in Table IV.

By contrast, on segments where only one carrier operated, the effect on flight frequency is always positive (and greater than the underlying industry trend). This suggests that the effect of the merger-induced network expansion dominates the loss of within-market competition for different segments which serve the same markets. Also, to the extent that airlines' frequency decisions are driven by competitive pressures, the relevant competition seems to be from rivals which operate flights on the same segments. Overall, then, the differences across mergers can be explained, at least in part, by the extent to which competition between the merging airlines occurred on shared segments. The two cases where the net change in flight frequency is significantly negative (NW-RC and TW-OZ) are also the only two where more than half of the relevant segments were served by both airlines. *A priori*, we would expect to see a significant overall effect of flight frequency changes on post-merger prices only in these two cases.

Post-merger changes in airport presence and the number of firms serving the market are reported in Table V. For the first, the figures reported ($\% \Delta Sh$) are the average percentage changes in the merging firms' combined share of traffic at each origin airport within an overlap market, net of the change in the industry-wide average over the same time period. Similarly, the second column (ΔNF) reports the average change in the number of rival firms per market for all overlap markets, net of the industry-average.¹⁸ This change does not include the immediate effect of the merger – a decrease by one in the total number of firms in the market – but captures any additional entry or exit.¹⁹

¹⁸ The industry average used for these changes is an unconditional mean across all airports or markets.

¹⁹ This effect seems more convenient to display as a unit change, rather than a percentage change. The qualitative interpretation is not, of course, affected.

Table V
Average Change in Combined Airport Shares
and Number of Firms in Overlap Markets
1985 vs. 1st Year Post-Merger

	Airports		Markets	
	#	%ΔSh	#	ΔNF
NW-RC	31	-4.4 (7.1)	78	0.146 (0.094)
TW-OZ	24	10.3 (13.3)	50	0.517 (0.115)
CO-PE	33	-19.8 (7.6)	67	0.057 (0.088)
DL-WA	10	-21.9 (6.6)	11	0.152 (0.161)
AA-OC	4	30.5 (46.2)	2	1.872 (0.125)
US-PI	26	-1.3 (16.7)	60	0.499 (0.191)

Standard errors for the mean, under normality, are in parentheses.

As with the overall changes in flight frequency, the averages reported here tend to be small relative to their standard errors, suggesting that the systematic effects of the merger are small relative to other sources of variation in these variables. In part, this may reflect ambiguity in the direction of the merger's effects. For example, post-merger price increases and capacity reductions are likely to reduce market shares, and hence airport presence; but at the same time, the marketing advantages of combining the two firms may increase the overall attractiveness of the merged airline, so that increased airport shares will tend to be self-reinforcing. In fact, the only significant changes in airport presence are decreases of about 20% for Continental-People Express and Delta-Western. For the first case, with its high price increases (documented in the previous section), it is not surprising that the first effect should predominate. Also, the Delta-Western merger is unique among these six cases in one important way: there was not a single

airport where *both* of the firms had at least a 5% share during 1985. By contrast, TWA and Ozark shared their primary hub, while each of the other cases involved at least four shared airports. The self-reinforcing nature of combining airport shares is likely to be effective only when both airlines have significant pre-merger share, so that again, it does not seem surprising that the average net effect of the Delta-Western merger on airport presence is negative.

The evidence on post-merger entry is similarly mixed, with insignificant results for three of the six cases. USAir-Piedmont and TWA-Ozark each faced, on average, entry by one rival into half of their overlap markets, while there were two post-merger entrants into both American-Air Cal overlap markets. Without placing too much significance on the outcomes in such a small subsample, the relatively substantial entry into these two markets offers at least an anecdotal explanation for the low post-merger price increases after this merger.

Taken together, it is interesting to note that each of the six cases shows a significant post-merger change for at least one of the three variables considered here (flight frequency, airport presence, and the number of rivals). Only TWA-Ozark has a significant change for two. As a result, we would expect to see some significant improvement in the accuracy of our simulated prices if we correct for these changes, particularly for the latter case.

5 Empirical Results

5.1 Demand Estimates

Turning now to the demand estimates, Table VI reports the estimated parameters from the two models: the simple nested logit (NL) with an outside good nest, and the nested PD-GEV with groupings for airport-pairs and nonstop/connecting.

Table VI
Estimated Demand Parameters

	NL	GEV
<i>lnp</i>	-2.54 (0.34)	-1.66 (0.20)
<i>lnf</i>	0.968 (0.057)	0.685 (0.044)
<i>lnAptPres</i>	0.0958 (0.0088)	0.0663 (0.0061)
<i>Nonstop</i>	2.22 (0.12)	1.60 (0.10)
<i>lnMiles</i>	-0.748 (0.083)	-0.575 (0.060)
r_0	0.595 (0.037)	0.557 (0.035)
r_{AP}	ρ_0 -	0.380 (0.027)
r_{NS}	ρ_0 -	0.478 (0.032)

Standard errors in parentheses,
robust to heteroskedasticity and
within-market autocorrelation.

Several points are worth noting. First, all parameters in both models are highly significant with the expected sign, and the elements of ρ satisfy the necessary restrictions ($0 < \rho_h \leq \rho_0 \leq 1$, $h=AP, NS$). Second, the restrictions of the nested logit ($\rho_{AP} = \rho_{NS} = \rho_0$) are easily rejected, indicating the significance of heterogeneity in consumer preferences for flying

direct and, especially, for traveling through particular airports in multi-airport cities. As a result of this rejection, the comparison between out-of-sample predictions of the two models will not be motivated by a desire to choose which is a more accurate representation of consumer demand: this is already possible using only pre-merger data. Rather, the comparison will be intended to demonstrate the effect on prediction of using a less realistic demand model.

For the purposes of merger simulation (without updating for changes in flight frequency and airport presence), the only parameters in these models that influence the price elasticities, and hence the post-merger predictions, are ρ and β_p . One convenient feature of this class of models is that the aggregate elasticity in any market is especially simple to compute:

$$\left. \frac{\partial \left(\sum_{j \neq 0} q_j(\lambda \mathbf{p}) \right)}{\partial \lambda} \frac{1}{\sum_{j \neq 0} q_j(\lambda \mathbf{p})} \right|_{\lambda=1} = \beta_p s_0$$

where s_0 is the outside good's share of the potential market.

This relationship has been documented for the simple logit model by Werden and Froeb (1994),²⁰ and it is straightforward to extend it to the nested logit. Comparable analytical results for the GEV model are considerably more complicated, but the relationship can be verified numerically by computing the aggregate elasticity directly for each market. Again, the importance of the price coefficient is evident. In fact, because the outside good's share is

²⁰ Werden and Froeb's model has price enter the utility function linearly, rather than log-linearly, so their formula differs slightly from the one given here.

generally close to one²¹, β_p approximately equals the aggregate elasticity for most markets. This highlights one significant difference between the NL and GEV results: The NL estimates ($\beta_p = -2.54$) imply a considerably more elastic aggregate demand than the GEV ($\beta_p = -1.66$). A comparison with other research on demand elasticity in the airline industry suggests that the GEV estimate is more accurate. For example, using estimates from Oum, Gillen, and Noble (1986), Brander and Zhang (1990) calibrate a model of airline competition using a “base case” elasticity of -1.6 , and argue that values of -1.2 and -2.0 “probably cover the reasonable range.” Hence, the formal rejection of the nested logit model against the alternative of the GEV is consistent with prior beliefs about the aggregate demand elasticity.²²

The more relevant quantities for merger simulation, however, are the individual product elasticities. The own-price elasticity of the demand curve facing an individual product indicates the extent to which consumers are willing to switch to alternatives after a marginal price increase, and hence measures the overall competition faced by that product. By contrast, the cross-price elasticity between two products in the same market measures how well the products substitute for each other. Intuitively, we would expect the anticompetitive effect of a merger to increase with the degree of substitution between the products of the merging firms. Table VII summarizes the own- and cross- price elasticities implied by the demand estimates, for all products in the pre-merger data, and for products offered by the merging carriers in overlap

²¹ The outside good’s share will usually be close to one when we define the market size as proportional to market population, with our choice of the constant of proportionality constrained by the requirement that total market sales must not exceed total market potential (i.e. the outside good’s share is never negative). This is a result of the highly skewed distribution of the ratio of market sales to market population.

²² Morrison and Winston (1996) assume a baseline elasticity of -0.7 , with reasonable range of -0.5 to -1.0 . Thus, the GEV results reported here can be interpreted as supporting Brander and Zhang’s (1990) assumption that demand for air travel tends to be more elastic.

markets. For the cross-price elasticities, pairs of products offered by the same carrier are excluded from the merger-specific statistics.

Table VII
Price Elasticity Summary Statistics (1985)
(Means and Standard Deviations)

	Own-price			Cross-price		
	#	NL	GEV	#	NL	GEV
Full Sample	42,613	-3.99 (0.376)	-3.45 (0.502)	391,046	0.158 (0.246)	0.196 (0.296)
NW-RC	818	-3.79 (0.386)	-3.24 (0.489)	1,148	0.443 (0.384)	0.571 (0.512)
TW-OZ	479	-3.73 (0.359)	-3.18 (0.470)	588	0.493 (0.359)	0.675 (0.509)
CO-PE	797	-3.99 (0.252)	-3.43 (0.372)	1348	0.248 (0.237)	0.293 (0.311)
DL-WA	116	-3.99 (0.326)	-3.55 (0.442)	178	0.234 (0.299)	0.302 (0.400)
AA-OC	32	-4.10 (0.081)	-3.45 (0.335)	48	0.152 (0.079)	0.288 (0.309)
US-PI	843	-3.94 (0.398)	-3.30 (0.546)	1,624	0.296 (0.363)	0.360 (0.427)

Standard deviations of the estimated elasticities are in parentheses.

Merger-specific results are for the merging firms' products in overlap markets only.

The unit of observation for the own-price elasticities is a product in a quarter.

For the cross-price elasticities, it is a product-pair in a quarter.

These statistics allow us to make two comparisons: between the NL and GEV models; and between the elasticities facing the merging carriers and industry-wide averages. The first comparison is clear: overall and across all subsamples, the GEV model implies smaller own-price elasticities (in absolute value) than the NL, although the difference is not as large as the difference between the aggregate elasticities might suggest; and larger cross-price elasticities. In other words, under the GEV model, when a product's price increases, fewer of its consumers are driven away, but a larger share of those that are choose other products, rather than leaving the

market entirely. The difference in the cross-price elasticities suggests that the post-merger price changes will tend to be under-predicted by the NL model.

A comparison between the own-price elasticities facing merging firms and industry averages indicates that the demand for products affected by merger is generally no more or less elastic than a typical product in the industry. When we look at the cross-price elasticities, however, we obtain a very different result: in the markets identified as overlap markets, the cross-price elasticity between the merging firms' products is significantly greater than the average. To some extent, this simply reflects the concentration criterion used to select the markets: two products which both have significant market share will generally have higher cross-price elasticities, other things being equal. Indeed, in the full sample of all product pairs offered by two different firms, the raw correlation between the estimated cross-price elasticity and the hypothetical post-merger change in the HHI is 0.38 (for both models). However, this does not account for the full effect. To see this, I construct the subsample of all firm-pairs and markets that would satisfy the concentration criterion, identifying 614 hypothetical mergers, with an average of 43 overlap markets for each. In this subsample, with 208,726 observations²³, the average cross-price elasticity is 0.225 (0.278) for the NL (GEV) model, considerably higher than the industry-wide average, but still lower than the average for each of the six actual cases, (with the exception of the anomalous American-Air Cal merger). In short, the products of the airlines involved in merger tended to be relatively close substitutes for one another.

While directly examining the estimated elasticities can be informative about the extent of competition between the merging firms, the averages reported here remain an incomplete picture.

In particular, we have not considered the cross-price elasticities between the merging firms and other rivals, and, when firms offer more than one product in a market, the combined effect of the merger may not be accurately represented by single-product elasticities. This raises the question of how to summarize the elasticities in a way that captures the full expected effect of the merger on pricing incentives. In some sense, this is precisely what is accomplished through formal merger simulation. As noted by Baker and Rubinfeld (1999), when merger simulation does not involve predicting changes in cost or conduct, it primarily consists of “transforming the demand elasticities into a more informative metric.”

5.2 Predicted Post-Merger Price Changes

Table VIII reports average predicted relative price changes for overlap markets for each of the mergers, computed in the same way as the actual changes discussed in Section 4.2. Predictions are reported for simulations with a change in ownership only, under the NL and GEV demand models, as well as for the linear model, with and without instrumenting for concentration. For comparison, actual price changes (from Table III) are reported again here.

²³ As in the cross-price column of Table VII, each observation corresponds to a product-pair in a quarter of 1985.

Table VIII

**Average Predicted % Relative Price Change in Overlap Markets
1985 vs. 1st Year Post-Merger**

	#	Simulation		Linear Prediction		Actual
		NL	GEV	OLS	IV	
NW-RC	78	7.0 (0.5)	19.8 (1.9)	8.8 (0.6)	21.6 (1.5)	7.2 (0.8)
TW-OZ	50	7.2 (0.6)	20.8 (2.5)	8.7 (0.7)	21.5 (1.8)	16.0 (2.3)
CO-PE	67	3.4 (0.3)	6.4 (0.6)	4.1 (0.4)	9.8 (1.1)	29.4 (3.6)
DL-WA	11	3.3 (1.0)	7.6 (3.1)	3.0 (1.1)	7.2 (2.7)	11.8 (3.6)
AA-OC	2	1.3 (0.1)	4.7 (0.0)	1.1 (0.2)	2.5 (0.5)	6.5 (4.6)
US-PI	60	4.5 (0.6)	12.7 (2.2)	6.4 (0.8)	15.7 (1.9)	20.3 (2.6)

Standard errors for the mean, under normality, are in parentheses.
Simulations assume no change in flight frequency, airport presence,
the unobserved component of mean utility, or marginal costs.

It should be noted that the standard errors reported in this table do not take into account the variance of the estimated demand parameters, but are simply computed from the empirical variation across markets in the predicted prices. Because the simulated prices are computed numerically, it is not possible to analytically compute the true standard errors. For this reason, it is typical in the merger simulation literature (e.g. Nevo (2001)) to report average predicted price changes without standard errors. One possible approach to correcting the standard errors would be to use a Monte Carlo procedure, simulating post-merger prices for a large number of draws from the estimated distribution of demand parameters. Unfortunately, the computational

demands of this approach are prohibitive.²⁴ Hence, the reported standard errors are to some extent underestimated.

These results allow us to draw five conclusions. (i) *Simulation with the GEV demand model predicts significantly greater price increases than with the nested logit.* This is exactly as we would expect, given the larger estimated cross-price elasticities obtained from the former model. It is important to note that the different elasticity estimates appear to be the driving influence behind the difference in the predictions. This contrasts with the focus of Crooke, *et al* (1999), where the predictions of four alternative demand models are compared. That paper used Monte Carlo methods to demonstrate that differences in the curvature of the demand systems can yield significantly different predictions, even when the pre-merger elasticities are calibrated to be equal across models. By contrast, the result here suggests that, in practice, the critical distinction between alternative demand models is more likely to be the different elasticity estimates which they provide.

(ii) *The linear model estimated with instrumental variables predicts significantly greater price increases than the linear model estimated by OLS.* This, of course, is a statement about the estimated coefficient on *HHI* in these two regressions. For the IV, $\beta_H = 0.766$, (s.e.= 0.015); for OLS, $\beta_H = 0.335$, (s.e.= 0.008). The large difference between the coefficients indicates the presence of unobserved variables, uncorrelated with the number of firms in the market (the instrument), but related to both price and concentration. Hence, we might expect the IV estimate

²⁴ A possible alternative is to estimate numerical derivatives of the simulated prices with respect to the relevant parameters, by running simulations with small perturbations in the parameter estimates, and then computing standard errors with the delta-method. Because the simulated prices only depend on the nonlinear parameters and the price coefficient, this would require only four additional simulations for the GEV model, two for the nested logit.

to yield a more accurate measure of the marginal effect of an exogenous change in concentration. Note that this does not necessarily imply more accurate predictions of post-merger price changes.

(iii) *The linear model with instrumental variables predicts price changes that are similar to those predicted by the GEV simulation.* It is also true that predictions from the linear model estimated by OLS are similar to those from the nested logit. Given the rejection of the nested logit model against the alternative of the GEV, however, this comparison is somewhat less interesting. The similarity between the linear IV model and the GEV is not perfect: the linear predictions for the Continental-People Express merger are, on average, more than 50% higher than the GEV's simulated prices, while the predictions for the American-Air Cal merger are almost 50% lower. For the other four cases, though, the predictions seem surprisingly close. This suggests that the empirical relationship between price and concentration may be closely related to the unilateral market power identified through formal simulation.

(iv) *The ordering of the simulated price changes across cases is perfectly matched by the estimated cross-price elasticities.* (See Table VII) This is true for both the NL and GEV models. Moreover, the relative magnitudes are very similar as well. This underscores Baker and Rubinfeld's (1999) point that merger simulation with ownership changes alone is essentially a means of summarizing the information contained in the elasticities.

(v) *In all cases except Continental-People Express, the price changes predicted by the GEV simulation or IV linear model account for a large share of the actual price change (or exceed it), but significant differences remain.* The first part of this finding supports the belief

that the unilateral incentive to raise price after a merger is an important component of the overall effect of the merger on prices. The second part indicates that it is not the only important component. The evidence about the importance of unilateral pricing incentives offers empirical justification for the use of these methods in merger analysis, while the significance of the prediction errors suggests the need for further refinements. The remainder of this section combines the simulation methodology with evidence from post-merger data in order to account for these differences between the predicted and actual prices, and hence to offer guidance in future efforts to improve prediction.

5.3 Accounting for the Observed Price Changes

Table IX reports the incremental effect on prices of four types of post-merger change, using the various updated simulations described in section 3, with the GEV demand model. The first column (“Ownership change”) is identical to the GEV column from Table VIII, and indicates the effect of the changed pricing incentives due to the ownership transfer. In the second column (“Observed Changes”), I take the changed ownership as given, and compute the combined effect of other observed post-merger changes: entry, exit, and other changes in the set of available products (including the occurrence of other mergers), as well as the demand-side effect of changes in flight frequency and airport presence. Taking these changes as given, the third column (“Change in μ ”) reports the price effect of changes in unobserved product quality. The fourth column (“Change in c ”) reports the price effect of changes in the estimated marginal costs, given all other changes. Each of these effects is reported as a percentage change relative to the prices that would otherwise hold, under the other assumed changes. As a result, the percentages do not add up exactly to the total actual price change (reported in the final column). Computing

them relative to the original prices, while it would allow the individual changes to sum to the total, has the undesirable property that the changes would be sensitive to the order in which they are computed (i.e., simulation with updated marginal costs and pre-merger product quality would imply different price effects for each component than simulation with updated quality and pre-merger costs).

Table IX
Component Effects of
Average % Relative Price Change in Overlap Markets

	#	Ownership Change	Observed Changes	Change in <i>m</i>	Change in <i>c</i>	Actual % <i>Dp</i>
NW-RC	78	19.8	-1.4	0.9	-10.1	7.2
TW-OZ	50	20.8	-2.2	-0.8	-1.0	16.0
CO-PE	67	6.4	0.7	0.2	20.5	29.4
DL-WA	11	7.6	-1.5	-0.5	6.0	11.8
AA-OC	2	4.7	-3.6	-1.8	7.6	6.5
US-PI	60	12.7	2.0	-1.9	6.7	20.3

This table provides the key finding of this paper: *With the notable exception of TWA-Ozark, the combined effect of all observed changes and the change in unobserved product quality is small relative to the effect of the change in marginal cost.* The exception for TWA-Ozark is wholly consistent with the earlier finding that this merger was followed by significant entry as well as significant frequency reductions. For the other cases, the implications for merger analysis – under the assumptions of the model – are clear: If we wish to improve significantly on the predictions implied by standard merger simulation, we need to focus our attention on anticipating changes in marginal costs.

The caveat about the modeling assumptions is critical, however, in view of how marginal costs are recovered. In particular, recall that the estimated marginal cost is equal to the difference between price and the estimated markup. This will be a consistent estimate of the true marginal cost only if the estimated markup is itself a consistent estimate of the true markup, a requirement that leans heavily on the assumption of Bertrand conduct with static profit-maximization.²⁵ In other words, changes in firm conduct may also account for the estimated changes in marginal cost. For example, consider the large changes in marginal costs following the Northwest-Republic and Continental-People Express mergers. In the former case, the marginal cost interpretation would suggest large efficiencies resulting from the merger, while pre-merger collusion between the two carriers might also explain the change. In the latter case, the results suggest that marginal cost rose dramatically after the merger, while a richer dynamic interpretation might allow for financially distressed airlines (People Express and Frontier) to set their prices below the short-run profit-maximizing level in order to meet near-term financial obligations. This sort of strategy is made possible by the fact that marginal revenue is received at the time a ticket is purchased, while marginal cost is incurred at the time of the flight. For firms with a sufficiently high discount factor (such as those facing imminent bankruptcy), this lag may drive a significant wedge between marginal revenue and cost.²⁶

While the large changes in the estimated marginal cost for these two cases raise concerns about the validity of the conduct assumption, it should also be noted that the marginal cost increases of 6 – 8% for the last three mergers are not unreasonable. While prospective merger analysis often focuses on the likelihood of cost efficiencies being realized, the actual experience

²⁵ Of course, it also requires that our demand model be well-specified and consistently estimated.

with the airline mergers of the 1980's suggests that inefficiencies may have been more common. As discussed by Kole and Lehn (2000), the integration of two distinct workforces in a heavily unionized industry can lead to significant increases in post-merger labor costs.

Nonetheless, the results presented here do not distinguish between changes in marginal costs and changes in conduct. Hence, a more accurate statement of the implications for merger analysis would be: to improve simulation methods, we need to focus on changes in marginal cost *or* allow for more flexible conduct assumptions. While this leaves a great deal of uncertainty about precisely what accounts for the post-merger price changes, the results are informative about what does *not* account for them, at least not to a significant degree. In short, efforts to predict post-merger entry or the demand-side effects of changes in product characteristics should not be expected to yield large gains in the accuracy of price predictions, at least in situations that resemble the airline industry of the 1980's.

²⁶ See Busse (2002) for other explanations for the relationship between financial distress and pricing conduct.

4 Conclusion

To summarize, the evidence presented in this paper supports the following conclusions: (i) Post-merger price increases in overlap markets were significant, with considerable variation across mergers. (ii) Standard merger simulation methods, which measure the effect of the ownership transfer on unilateral pricing incentives, can account for a large component of the post-merger price change, but should not be expected to account for all of it. (iii) The sensitivity of the simulated price changes to the estimated elasticities suggests that unduly restrictive demand models, such as the logit or nested logit, should be used with caution. (iv) Linear prediction based on the empirical relationship between price and market structure can yield results which are reasonably close to the predictions from formal simulation. (v) Changes in marginal costs or firm conduct play a significantly larger role in explaining post-merger price increases than do changes in the set of products or changes in observed and unobserved demand variables.

This final conclusion suggests two important directions for future research aimed at improving the predictive performance of merger simulation. First, it is desirable to predict changes in marginal cost, perhaps by estimating the relationship between marginal costs and other variables influenced by the merger, such as capacity or segment-level traffic. In practice, this would require predictions for post-merger changes in these variables, which might be accomplished by either reduced form or structural methods. Secondly, efforts to accommodate alternative conduct assumptions, and to distinguish between them, are much needed in merger analysis. Based on the evidence from the airline mergers of the 1980's, these lines of research have significant potential to improve the accuracy of predicted post-merger prices.

References

- Baker, J. (1997) "Contemporary Empirical Merger Analysis," *George Mason Law Review*, 5(3), 347-361.
- Baker, J., and T. Bresnahan (1985) "The Gains from Merger or Collusion in Product-Differentiated Industries," *Journal of Industrial Economics*, 33(4), 427-444.
- Baker, J., and D. Rubinfeld (1999) "Empirical Methods in Antitrust Litigation: Review and Critique," *American Law and Economics Review*, 1(1-2), 386-435.
- Berry, S. (1994) "Estimating Discrete-Choice Models of Product Differentiation," *Rand Journal of Economics*, 25(2), 242-262.
- Berry, S., and A. Pakes (1993) "Some Applications and Limitations of Recent Advances in Empirical Industrial Organization: Merger Analysis," *American Economic Review*, 83(2), 247-252.
- Berry, S., J. Levinsohn, and A. Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841-890.
- Beutel, P. and M. McBride (1992) "Market Power and the Northwest-Republic Airline Merger: A Residual Demand Approach," *Southern Economic Journal*, 58(3), 709-720.

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

Borenstein, S. (1991) "The Dominant-Firm Advantage in Multiproduct Industries: Evidence from the U.S. Airlines," *Quarterly Journal of Economics*, 106, 1237-1266.

Brander, J., and A. Zhang (1990) "Market Conduct in the Airline Industry: An Empirical Investigation," *Rand Journal of Economics*, 21(4), 567-583.

Bresnahan, T., S. Stern, and M. Trajtenberg (1997) "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s," *Rand Journal of Economics*, 28(0), S17-S44.

Brueckner, J., N. Dyer, and P. Spiller (1992) "Fare Determination in Airline Hub-and-Spoke Networks," *Rand Journal of Economics*, 23(3), 309-333.

Busse, M., (2002) "Firm Financial Condition and Airline Price Wars," *Rand Journal of Economics*, 33(2), 298-318.

Crooke, P., L. Froeb, S. Tschantz, and G. Werden (1999) "Effects of Assumed Demand Form on Simulated Postmerger Equilibria," *Review of Industrial Organization*, 15, 205-217.

Dube, J.-P. (2000) "Product Differentiation and Mergers in the Carbonated Soft Drink Industry," working paper, University of Chicago.

Feldman, R. (1994) "The Welfare Economics of a Health Plan Merger," *Journal of Regulatory Economics*, 6(1), 67-86.

Hausman, J., G. Leonard, and D. Zona (1994) "Competitive Analysis with Differentiated Products," *Annales d'Economie et de Statistique*, 0(34), 159-180.

Kim, E.H., and V. Singal (1993) "Mergers and Market Power: Evidence from the U.S. Airline Industry," *American Economic Review*, 83, 549-569.

Kole, S., and K. Lehn (2000) "Workforce Integration and the Dissipation of Value in Mergers: The Case of USAir's Acquisition of Piedmont Aviation," in *Mergers and Productivity*, ed. by S. Kaplan, Chicago: University of Chicago Press.

Morrison, S., and C. Winston (1996) "Causes and Consequences of Airline Fare Wars," *Brookings Papers on Economic Activity, Microeconomics*, Vol. 1996, 85-123.

Nevo, A. (2000) "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *Rand Journal of Economics*, 31(3), 395-421.

Oum, T., D. Gillen, and D. Nobel (1986) "Demand for Fareclass and Pricing in Airline Markets," *Logistics and Transportation Review*, Vol. 23, 195-222.

Werden, G. and L. Froeb (1994) "The Effects of Mergers in Differentiated Products Industries: Logit Demand and Merger Policy," *Journal of Law, Economics, and Organization*, 10(2), 407-426.

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

Appendix I

For the nested PD-GEV model, the derivative of the market share of product j with respect to the price of product k is given by:

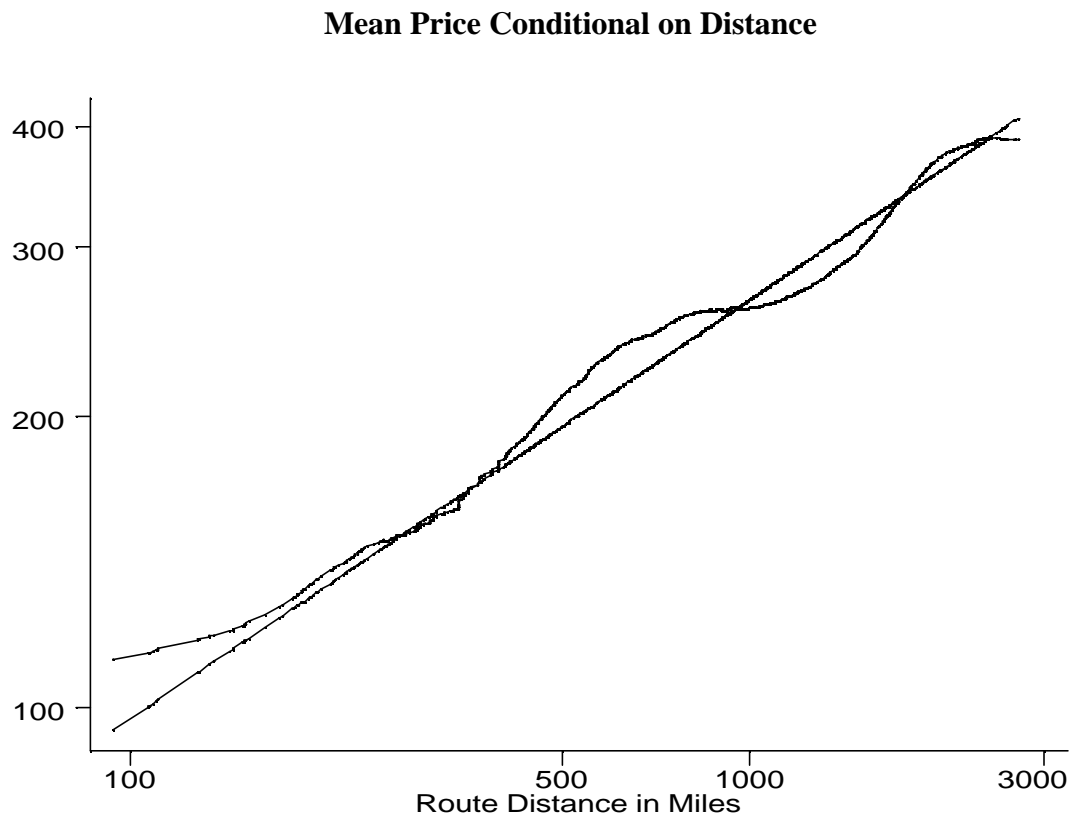
$$\frac{\partial s_j}{\partial p_k} = \frac{\beta_p}{p_k} \cdot \left\{ \begin{aligned} & \left[\frac{(a/\rho_{AP}) \cdot \exp(\delta_j/\rho_{AP})}{\left[\sum_{j' \in AP_j} \exp(\delta_{j'}/\rho_{AP}) \right]^{1-\rho_{AP}/\rho_0}} + \frac{((1-a)/\rho_{NS}) \cdot \exp(\delta_j/\rho_{NS})}{\left[\sum_{j' \in NS_j} \exp(\delta_{j'}/\rho_{NS}) \right]^{1-\rho_{NS}/\rho_0}} \right] \cdot \frac{\mathbf{H}^{\rho_0-1}}{1+\mathbf{H}^{\rho_0}} \cdot \mathbf{1}(j=k) \\ & + \frac{a(\rho_0^{-1} - \rho_{AP}^{-1}) \cdot \exp(\delta_j/\rho_{AP}) \cdot \exp(\delta_k/\rho_{AP})}{\left[\sum_{j' \in AP_j} \exp(\delta_{j'}/\rho_{AP}) \right]^{2-\rho_{AP}/\rho_0}} \cdot \frac{\mathbf{H}^{\rho_0-1}}{1+\mathbf{H}^{\rho_0}} \cdot \mathbf{1}(k \in AP_j) \\ & + \frac{(1-a)(\rho_0^{-1} - \rho_{NS}^{-1}) \cdot \exp(\delta_j/\rho_{NS}) \cdot \exp(\delta_k/\rho_{NS})}{\left[\sum_{j' \in NS_j} \exp(\delta_{j'}/\rho_{NS}) \right]^{2-\rho_{NS}/\rho_0}} \cdot \frac{\mathbf{H}^{\rho_0-1}}{1+\mathbf{H}^{\rho_0}} \cdot \mathbf{1}(k \in NS_j) \\ & - s_j s_k \left[\frac{1-\rho_0}{\rho_0} (\mathbf{H}^{-\rho_0} + 1) + 1 \right] \end{aligned} \right\}$$

where $\mathbf{H} = a \sum_{AP} \left[\sum_{j' \in AP} \exp(\delta_{j'}/\rho_{AP}) \right]^{\rho_{AP}/\rho_0} + (1-a) \sum_{NS} \left[\sum_{j' \in NS} \exp(\delta_{j'}/\rho_{NS}) \right]^{\rho_{NS}/\rho_0}$

and where the other terms are as defined in the text.

Appendix II

Figure I



Straight line = predicted values from log-linear regression

Curved line = locally-weighted running mean, with bandwidth = 0.2

Data from first quarter, 1985