The Effect of the Internet on Pricing in the Airline Industry

Jeremy A Verlinda[1]
University of California, Irvine
Department of Economics
3151 Social Science Plaza A
Irvine, CA 92697
verlinda@uci.edu

Leonard Lane
University of California, Irvine
Graduate School of Management
GSM 415
Irvine, CA 92697
llane@gsm.uci.edu

November 2004

Abstract

We examine the relationship between nonlinear pricing and consumer search costs in the airline industry during a period of increasing Internet use. We find that higher Internet usage increases the spread between unrestricted and restricted fares. We also find that the ratio of unrestricted to restricted ticket fares increases with more competition, but the effect is smaller as a result of a larger population searching for airline travel online, implying that the Internet is inducing more competitive firm behavior. The data support models of price discrimination through brand differentiation, but do not support standard consumer search cost models well.

Keywords: Airlines, Internet, price discrimination, consumer search.

JEL Codes: L11, L15, L93

[1]Corresponding author. We would like to thank Jun Ishii, Linda Cohen, and Bo Carlsson for their help in developing this paper. All mistakes are, of course, our own.

1 Introduction

Nonlinear pricing in competitive markets indicates a fundamental information-asymmetry between firms and consumers. Typically, consumers are assumed to exhibit vertical preference diversity and thus have different reservation values for product quality, but firms do not directly observe consumer preferences. As long as firms have some market power, if they observe each consumer's preferences, they can exploit that information to charge each consumer his reservation value.[1] Nonlinear pricing occurs when firms do not observe this information. Instead, firms offer a price schedule and allow consumers to sort themselves according to their preferences.

Alternatively, consumers themselves may lack information regarding firm-specific price-quality tradeoffs. Consumers are assumed to pay a positive cost to learn about firm offerings. With high enough search costs, the benefits of additional search (i.e., finding a similar quality good at a lower price) are offset by the cost of searching itself, so that consumers have no incentive to perform additional search. Price schedules in such an environment tend to reflect monopoly-based pricing, and as search costs fall, prices in elastic market segments become more competitive. If, on the other hand, consumers display considerable heterogeneity in their horizontal preferences across, say brands, falling search costs may actually decrease the level of competition by allowing consumers to more closely identify their most preferred brands. Under this scenario, it is actually possible for prices to increase across vertically differentiated market segments.

We analyze the relationship between market segmentation as characterized by airline price schedules and the size of the Internet population going online to search for airline travel. Our ability to precisely measure the impact of the Internet on airline pricing will determine which of the consumer search models are supported by the data. We acknowledge, however, that airline fares have been intensely studied since deregulation of the industry in 1978, with competing models of price discrimination and yield-management pricing offering relevant and competent

[1] We are assuming that there are no arbitrage opportunities for consumers and that no government regulations prevent price discrimination.

explanations for the pricing variation observed in the data.

We contribute to this literature by incorporating consumer search costs into our study of airline price discrimination. As many of the theoretical models of consumer search lead to different and often contradictory conclusions on the effect of falling search costs, a primary goal of this paper is induction—our prior beliefs on the effect of the Internet on airline price schedules are weak. We wish to infer from the data which of the models are best supported. We do so by investigating the following empirical issues: (1) do markets with more informed consumers exhibit more competitive pricing, (2) what is the effect of market concentration on airline fares, and (3) how does an increase in the number of informed consumers (as a proxy for a decrease in search costs) influence the effect in (2)?

We find that increases in the size of the online population searching for airline travel are associated with higher unrestricted fares, but have a negligible or uncertain effect on restricted ticket fares. Nevertheless, the ratio of unrestricted to restricted ticket fares measurably increases. Increased competition is associated with lower prices for both unrestricted and restricted ticket prices, but since restricted ticket prices fall by significantly more than unrestricted ticket prices, the ratio of prices is increasing as competition increases, consistent with Borenstein and Rose (1994). We also show that the market for air travel is becoming more price-competitive as the size of the Internet airline travel search population grows. Our findings contradict the predictions in the majority of theoretical models on consumer search; instead, they most strongly support Borenstein's (1985) and Salop's (1995) models of brand-intensive competition in differentiated products markets. There is weak support for the Lal and Sarvary (1999) model that incorporates consumer search costs into a model of product differentiation. Our empirical findings suggest that a theoretical model that combines consumer search costs with a quality-based market screening mechanism may best explain the empirical phenomena we observe in airline fares.

Section 2 of this paper explains in more detail the role of search costs in differentiated products markets, building on models in both the consumer search literature and in the price discrimination literature. Section 3 applies the implications discussed in Section 2 to the market for air travel and the appropriate predictions that we will take to the empirical modeling. In Section 4 we explain our data sources and provide a brief summary description of the data. We present our empirical model in Section 5, and discuss our results and their implications in Section 6. We conclude with a summary of our study's findings in Section 7.

2 Nonlinear Pricing & Consumer Search

Consumer search costs are incorporated into models of price discrimination in two ways. The most common conceptualization, as in Katz (1984), is to assume that heterogeneity in search costs creates multiple potential market segments, with firms screening consumers by their search intensity. Alternatively, we can build on a basic multiproduct, multifirm model as in Borenstein (1985) or Stole (1995), and introduce price and/or product quality uncertainty–the intuition in such models is that consumers with high search costs will sample a smaller number of price schedules than consumers with low search costs, increasing each firm's monopoly power relative to the standard (no search cost) models. In what follows, we develop these models more fully.

In their famous "Bargains and Ripoffs" paper, Salop and Stiglitz (1977) incorporate consumer information-uncertainty into a model of monopolistic competition. Consumers endogenously choose to become informed, with the cost of information varying across consumers. They show that there exists a non-uniform pricing equilibrium such that efficient searchers find the lowest prices while costly searchers pay up to monopoly level prices. Moreover, if the share of low-cost consumers exceeds some critical level, the low-cost group exerts a positive externality in the market and induce all firms to offer perfectly competitive prices.

Katz (1984) extends the Salop and Stiglitz model to accommodate nonlinear pricing as a mechanism for sorting consumers by their information costs. Katz assumes that informed consumers purchase large quantities of the good from the lowest price store, while uninformed consumers make small purchases chosen from a randomly selected store.[2] In this model, unin-

[2] Although the model explicitly discusses quantity discounting as the form of nonlinear pricing, a simple redefining of quality for quantity would lead to the same results for quality-based sorting as well. Specifically, we can think of large quantity purchases as analogous to low-quality purchases, while small quantity purchases would thus constitute high-quality purchases. Necessarily, we must also reinterpret the assumption that the average cost curve be U-shaped in product quantity; consistency requires that costs increase at an increasing

formed consumers pay monopoly-level prices while informed consumers pay less than marginal cost. Similar to the Salop and Stiglitz model, as the number of informed consumers increases, the number of firms increases and the price paid by informed consumers falls. By assumption, uninformed consumers continue to pay their reservation value (monopoly price) for the good.

Unfortunately, the standard role played by consumer search costs in most economic models is somewhat limited in the kinds of phenomena we would like to describe. In particular, the Katz assumption that the pricing schedule itself is a valid mechanism for screening consumers by their search costs precludes other motivations for nonlinear pricing. In the spirit of Borenstein (1985) and Stole (1995), we would like to incorporate preference heterogeneity into a model of information uncertainty. In a widely cited paper, Bakos (1997) claims that as consumer search costs fall in a differentiated products market, the market becomes more competitive and prices are driven toward marginal costs. This interpretation of the effect of falling search costs seems somewhat optimistic, however, since standard nonlinear pricing models, which operate in an economy with no consumer search costs, do not predict marginal cost pricing.[3]

A more realistic interpretation of falling search costs is that they may indeed promote a more competitive market, but that this need not lead to zero firm profits. We believe that the basic intuition in Diamond (1971) holds; high search costs confer on each firm a level of market power approaching the monopoly level regardless of the number of firms. When search costs prevent consumers from comparing the price schedules of competing firms, each firm is potentially a monopoly for the subset of consumers that sample its price schedule. As search costs fall and consumers sample more price schedules, each firm's price-quality tradeoffs become more transparent, and competition, say at a given quality level, increases.

Lal and Sarvary (1999) offer an extreme point-of-view relative to the Bakos argument by examining under what conditions lower search costs actually decrease the level of competition. Their underlying argument is that falling search costs better enable firms to differentiate their products relative to their competitors', because consumers can more easily compare products.

rate as quality increases.

[3]Bakos' argument may be partially justified by any of the homogeneous goods models that find the market moving from monopoly level pricing toward competitive pricing as search costs fall. See, for example, Stahl's (1989) model of sequential consumer search in oligopoly markets.

Moreover, they argue that in reducing search costs, the Internet can actually discourage consumer search by allowing firms to "leverage their brand loyalty." An interesting implication of this model is that prices may actually increase as search costs fall, although it relies on a questionable assumption for the level of brand rivalry in the market.

Maintaining the assumption that falling search costs induce more competition, we can also build on the work of Borenstein (1985), Stole (1995) and Rochet and Stole (2002) to discuss competing implications of lower consumer search costs. Consider Borenstein's example of consumers self-sorting into either a high- or low-quality group, but that under zero vertical preference uncertainty firms could rank consumers by their reservation values and target them directly. When vertical preference diversity is the most relevant information-asymmetry between firms and consumers (relative to horizontal or brand-based heterogeneity in preferences), we expect prices to fall for both groups with additional firm entry. Following Stole (1995), the argument goes that high reservation-value consumers benefit the most from more competition in this environment. While prices for both market segments fall, the high-quality group experiences the largest price decreases. But when horizontal preference uncertainty dominates, low-quality consumers benefit most from added competition, since they are the more elastic group and the most likely to switch brands.

Rochet and Stole (2002) introduce the potential for all consumers along the preference distribution to drop out of the market for a given change in the price schedule, versus the standard assumption that only the lowest-type consumers will drop out. Allowing for random market participation leads to the prediction that high-valuation consumers receive the largest price declines for the same quality level. As opposed to the Stole model where this result occurs when vertical preference uncertainty dominates, the Rochet and Stole prediction arises because high-valuation consumers are the most likely to look for substitutes outside the market when competition increases.

3 Application to Airline Pricing

Examples abound for cases in which we think quality-striated price schedules reflect price discrimination—the most common being the variation in fares paid by different airline travelers across restriction types. Unrestricted-ticket travelers typically pay a premium for added flexibility, while restricted tickets are generally offered at a discount in exchange for giving up some flexibility in travel preferences. If travelers' preferences can be sorted by those who are willing to pay extra for the contingency of less-costly rescheduling or cancellation, and those who are not, then this form of nonlinear pricing by airlines is potentially an effective form of market segmentation. The effectiveness of ticket restrictions depends on demand elasticities between the two groups and across firms, and how willing unrestricted-class travelers are to switch to restricted tickets as they collect more information on the price schedules offered by different firms. If the restricted ticket market is price-elastic, and cross-price elasticities across firms are positive, then restricted ticket prices should fall as search costs fall. At the same time, if the cross-price elasticity between unrestricted and restricted tickets is positive, then we may also expect that some marginal consumers will switch from unrestricted to restricted tickets, shrinking the high-quality market.

We contend that an increasing population of Internet users searching for airline travel online is a valid indicator for either falling search costs or an increase in the share of low search-cost consumers. From 1998 to 2002, the number of consumers using the Internet to search for airline travel increased from approximately 2 million people at the beginning of the time period to just over 65 million people in 2002. Over the same period the share of unrestricted, coach-class tickets fell from about 21% to 7% as a percentage of all coach-class tickets. Peculiarly, prices during this time period do not readily comply with the arguments above. Figure 1 presents average fares by restriction type on the vertical axis with number of consumers using the Internet to search for airline travel on the horizontal axis. Unrestricted fares rise as more people use the Internet to research airline tickets, consistent with a shrinking and possibly less-competitive market for the high-quality good. Yet average restricted-ticket fares are also measurably greater when the number of people using the Internet to research airline tickets is greater, contrary to the claim that low-quality prices should fall toward marginal-costs. Additionally, the disparity between fares is growing with the increase in online searchers, since unrestricted-ticket fares are rising much faster than restricted-ticket fares.

3.1 Empirical Implications for Airline Fares

We determine if our empirical results are supported by any particular model by examining the predictions of the models described in Section 2. Due to the nature of our data, we are limited in the kinds of phenomena we can hope to capture. Our price data is a sample of average prices by fare type (restricted versus unrestricted travel) across markets and across time. In order to extend the implications of the standard consumer search models to our study, we must rely on the claim that airline fare-types screen consumers effectively based on their search costs, which suggests that restricted tickets are directed at low search-cost consumers and unrestricted tickets at high search-cost consumers. No such argument is required in order to extend the predictions of the nonlinear pricing models to our data, as there is a clear quality differentiation between unrestricted and restricted tickets. In the following sections, we describe the theoretical predictive effects of both increasing Internet usage and firm entry. For expositional ease, a summary of these effects is offered in Table 1 as well.

3.1.1 Search & the Internet

The predicted effects on prices of an increase in the online population researching airline travel tend to vary widely across the consumer search cost models. The Salop and Stiglitz (1977) model implies that as the share of informed consumers increases, unrestricted ticket prices will eventually collapse to the (implied) marginal cost pricing. Since restricted tickets are already priced at marginal cost, the price ratio should fall when Internet usage reaches the critical take-up point. Katz (1984) argues exactly the opposite, with the implication that we would expect unrestricted ticket prices to remain fixed at monopoly pricing levels as the number of online searchers increases, while restricted ticket prices fall; thus the price ratio should increase with more Internet airline research activity. Following Stahl (1989), as the share of costless searchers increases, we should see a decrease in the price of both fare types, and no definite conclusion on what should happen to the price ratio.

The Bakos (1997) model can be interpreted to suggest that prices should be falling for all ticket types toward marginal costs as search costs fall, which in turn suggests that price dispersion itself should be falling. In contrast, the Lal and Sarvary (1999) model would imply that prices for both ticket types should increase, with an ambiguous result for the effect on the ratio of fares.

3.1.2 Implications of Entry

Much of the earlier research into airline fares was concerned with the effect of market structure on fares and different measures of price discrimination. We see this as an indicator that any empirical model which does not account for market structure will suffer a potentially severe omitted variables bias in its estimation. In what follows, we outline the predictions for market structure effects of the models considered in this paper. Under the Salop and Stiglitz model, we would expect that unrestricted ticket prices should fall as the number of firms increases, while restricted ticket prices remain fixed at marginal cost. The Katz model would suggest that unrestricted ticket prices remain fixed at monopoly levels, while restricted ticket prices fall. The Stahl model would imply that unrestricted ticket prices increase and restricted ticket prices decrease as competitiveness increases. Neither Bakos (1997) nor Lal and Sarvary (1999) provide predictions for the effects of entry.

Within the price discrimination models, Borenstein's (1985) "competitive-type" (brand-intensive) price discrimination would imply that both unrestricted and restricted ticket prices fall in more competitive markets, but that restricted tickets fall by more than unrestricted. As noted above, Stole's (1995) nonlinear pricing model suggests the same implications for unrestricted and restricted ticket fares as Borenstein's. Rochet and Stole's (2002) model, on the other hand, arrives at the opposite conclusion; unrestricted ticket fares should fall by more than restricted ticket fares when random market participation is introduced.

3.2 Previous Empirical Work

The most important empirical work relative to our own is the study by Borenstein and Rose (1994). They investigate whether price dispersion in airline fares is consistent with a model of price discrimination, and argue that since their measure of price dispersion is greater on more competitive routes, the data suggest that airline fares are in line with third-degree price discrimination à la Borenstein (1985). Liu (2003) uncovers a U-shaped relationship between market concentration and price dispersion in airline fares, which he interprets as consistent with a model of second-degree price discrimination.

The Borenstein and Rose study also spawned follow-up empirical work on the causes of price dispersion in airline fares. Hayes and Ross (1998) find that price discrimination explains the dispersion in prices for less competitive environments, but that most of the dispersion can be attributed to fare wars and peak-load pricing. Their results are consistent with Dana's (1999a) theoretical explanation that price dispersion arises in a model with stochastic demand and price rigidities, such that firms offer a range of prices to shift demand from peak to off-peak periods when capacity costs are sunk. Stavins (2001) directly extends Borenstein and Rose's study by looking at a finer level of data to determine if the result that more competition leads to a greater divergence in fares across restriction types still holds when comparing specific fencing strategies utilized by airlines. Specifically, she finds that market concentration is associated with lower fares for restricted tickets (travel that requires either a Saturday-night stay-over or advanced-purchasing) and higher fares for unrestricted tickets.

Other papers have looked at nonlinear pricing outside the context of airlines. Notably Shepard (1991) and Borenstein (1991) both find evidence of second-degree price discrimination in retail gasoline. Cohen (2000), Clerides (2002), and McManus (2003) each find empirical evidence that suggests firms in oligopoly markets use product quality or quantity to screen customers. Lastly, Busse and Rysman (2001) test some of the implications in Stole's model for the effect of competition on prices at high versus low ends of the quality distribution, finding evidence that increasing competition causes larger yellow-page ad prices to fall more than do smaller ads.

4 Data

The primary data for this study is the DB1B database of the Origin and Destination survey compiled by the U.S. Department of Transportation, which collects a 10% sample of tickets from reporting carriers. We focus on data sampled from Q1/1998 through Q2/2002 to analyze fare levels over the five-year period in which we know the Internet was being used by travelers to book tickets. The price data include average fares by ticket coupon type on 25 markets that were preselected to provide variation in market characteristics based on prior knowledge of the markets. We deviate from earlier analyses in defining a market as the unit of observation as opposed to the route, which allows us to incorporate what we believe are two important phenomena: (1) travelers may be relatively indifferent between airport locations within markets (e.g., both the Oakland and San Francisco airports are nearly equidistant from the primary city center in San Francisco) and (2) airlines may compete in a market by flying through one of the subsidiary airports as opposed to the main airport (e.g., JetBlue recently entered the Los Angeles to New York market by flying out of Long Beach airport instead of LAX). A full listing of the markets and their associated airports is included in Table A1.

As is common in the literature, we restrict attention to coach class, round trip, direct flight tickets, and use the full round trip fare as our dependent variable. Hence the unit of observation is the average fare for the corresponding market, quarter, and restriction type, and so we do not clean the data for miskeys or frequent flyer fares in the manner of Borenstein and Rose (1994), Hayes and Ross (1998), or Stavins (2001).[4] Unfortunately, we lose a level of observation in that we do not observe fares by carrier and thus have fewer total observations overall relative to other studies. This should not present any major difficulties when interpreting our results, since market concentration only varies at the market level anyway. To the extent that market concentration only has an effect on the average price in the market for a given ticket restriction type, and provided the restriction types are appropriately coded in the DB1B database, working with average prices by restriction class at the market level actually *is* the appropriate level of

[4]Note that frequent flier awards may systematically lower the average restricted ticket fare. While our data does not allow us to test for or control for this phenomenon, we feel that so long as the bias is constant across markets and time, it should not substantially affect our results.

data observation. Table 2 lists descriptive statistics for the fare data by fare type, in addition to providing information on the remaining variables in the study.

The Internet data were collected from Nielsen proprietary industry data on use of the Internet for travel purchases overall, as well as the number of consumers using the Internet to search for airline tickets by quarter for the study period. Since the number of users purchasing airline tickets online is anticipated to be highly correlated with the number of users searching for airline travel, and since we feel that the dominant phenomenon with respect to the Internet is how travelers use it to learn about schedule offerings and fares, the analysis here restricts attention to just online searchers. We noted in Section 3 that Figure 1 provides suggestive evidence of the relationship between fares and Internet search activity. Averaging across markets, we see that increasing Internet use is positively correlated with both restricted and unrestricted ticket prices. Moreover, we see preliminary evidence that unrestricted fares seem to trend upward with Internet use much more quickly than restricted fares. Although we have yet to present any results controlling for the variety of factors that may also describe ticket fares, the simple correlative structure of the data suggests that the Internet may be influencing the effectiveness of ticket fencing as a screening mechanism in airline pricing strategy.[5]

Probably the most relevant control variable that we must consider is market structure. In this study we employ the Herfindahl-Hirschman Index (HHI), which is an index of market concentration based on firm market shares. Mathematically it takes the form $HHI_{mt} = \sum_k S_{kmt}^2$, where S_{kmt} is the market share for carrier k in market m at time t. We employ data from the Official Airline Guide (OAG) to gather information on the number of seats offered by carrier per market and quarter to derive the market shares in S_{kmt}. Stavins (2001, working paper version) argues that short-run concentration in any given market can be assumed fixed, and therefore exogenous relative to airline fares. This argument only holds for the airline capacity decision, as it is determined in advance, and would not be valid for HHI determined by seats sold. In order to preserve this exogeneity in our primary measure of competitiveness,

[5]It should be noted that the Nielsen data describe nationwide trends in Internet usage, and do not capture variation in Internet take-up across markets. While market-varying data would provide a natural experimental flavor to our results, we feel it is still interesting to describe the average effect of the Internet on airline fare practices.

we therefore construct HHI based on the number of seats offered as opposed to the number sold.[6,7]

GDP data and oil spot prices were collected by the authors, where GDP is in real 1996 dollars and the oil spot price is the average Cushing FOB spot price for the relevant quarter. Income and population statistics come from median household income and total population by PMSA for the corresponding markets in the 2000 census. Another common flight characteristic assumed to primarily influence costs is the distance between markets, which we derive from the OAG data on the great-circle distances between markets. Differences in temperature are calculated as the average temperature in the destination market minus the average temperature in the origin market for each time period in the sample, where the temperature data are collected from the National Oceanographic and Atmospheric Administration. Four U.S. airports are slot controlled for capacity/safety purposes by the FAA: Washington Reagan, Chicago O'Hare, New York La Guardia and John F. Kennedy. We therefore include in our analysis an indicator variable equal to one if the origin market contains a slot-controlled airport. We also control for the presence of network hubs by including a dummy variable equal to one if the origin market contains a network hub as one of its airports.

5 Empirical Model

We are interested in the relationship between Internet airline search activity and the difference in fares between unrestricted and restricted tickets, and specifically the effect of online research on each fare type individually. Given the existing body of literature on the relationship between price dispersion and market structure, we also estimate the effect of market concentration on fares, and allow the effect to depend on Internet search activity and other time-varying

[6] Other studies, e.g., Borenstein and Rose (1994) and Stavins (2001), use the share of flights at the airport or route level. Since airlines can alter their capacity through the mix of planes with different seat capacities on a route, we instead aggregate the total number of seats to derive a more accurate measure of market shares by airline.

[7] Unlike other studies, we do not include in our regressions dummy variables for markets that are monopoly or duopoly dominated. At the market-, as opposed to the route-, level of observation, there is less concern that the primary option for consumers is dominated by only one or two firms, since alternative carriers offer service at nearby airports. For a market-level analysis, we feel that HHI better describes the potential for a few airlines to still dominate the market by using their overall market share instead of just an indicator that they dominate in potentially only one airport.

characteristics. Due to concern that the predictive distribution for prices must exclude negative values, we use the natural log of fares as our dependent data. This also allows us to describe the marginal effects of our covariates in terms of mean predictive percentage changes in fares, which we believe is an intuitively attractive descriptive statistic.

We are also concerned that variance in the error terms may not be constant across either markets or time, and so include in the model specification both market- and time-level random effects. In addition to providing a potentially broad heteroskedastic characterization of the data, the hierarchical form of our model allows the random effects to have an estimated mean effect, with the benefits that are generally associated with traditional fixed effects modeling—by including market- and time-level mean effects, we are less susceptible to omitted variables problems.[8,9] The exact form of the regression model is thus:

$$y_{mtr} = \ln p_{mtr} \overset{ind}{\sim} N\left(\mu_{mr} + \tau_{tr} + HHI_{mt}\alpha_{tr} + z_{mtr}\beta_r, \sigma_y^2\right), \tag{1}$$

$$\mu_{mr} \overset{ind}{\sim} t\left(d_{mr}\zeta_r, \sigma_\mu^2, \upsilon\right) \text{ for } m = 1, ..., M, \tag{2}$$

$$\tau_{tr} \overset{ind}{\sim} t\left([h_{tr}, Int_t]\,\xi_r, \sigma_\tau^2, \varphi\right) \text{ for } t = 2, ..., T, \tag{3}$$

$$\alpha_{tr} \overset{ind}{\sim} t\left([q_{tr}, Int_t]\,\delta_r, \sigma_\alpha^2, \nu\right) \text{ for } t = 1, ..., T. \tag{4}$$

The first level of the hierarchy describes the dependence of log-fares on market and time means, market concentration, and our remaining market-time varying control variables. The second stage of the hierarchy describes the random effects and coefficients. The μ_{mr} are the market-level random effects, whose conditional expectations depend on the market-varying covariates d_{mr}. The τ_{tr} represent the time-level random effects, whose conditional expectations

[8]Of course, if there are any market-time interactive variables that we have excluded, then we are still subject to a potential omitted variables bias. For example, suppose incomes in L.A. were growing faster than any of our other origin markets during the study period. If firms respond to demand, then prices for flights originating out of L.A. should be higher, ceteris paribus. Failure to account for this effect may cause us to misinterpret the effect of market concentration on fares unless market concentration is uncorrelated with our omitted variables.

[9]Including fixed effects helps alleviate concerns we may have regarding measurement error at the market- or time-level as well. In particular, due to an oversight in the data collection process, we excluded Baltimore airport from the Raleigh-Washington D.C. and St. Louis-Washington D.C. markets. Since Southwest only services these markets via Baltimore, and JetBlue does not service them at all, we have inadvertently excluded the major low-cost carriers from these two markets. So long as this measurement error is time-invariant, the market-level fixed effects should absorb any bias that would adversely affect the main parameters of interest–the effects of the Internet and market concentration.

depend on the time-varying covariates h_{tr}, including Internet usage. Lastly, the α_{tr} are the time-level random coefficients on market concentration (HHI), whose conditional expectations depend on the time-varying covariates q_{tr} and Internet search activity. We feel that specifying the effect of market concentration to vary over time (as opposed to by market) will allow us to reconcile some of the earlier empirical work with our own, and to reconcile some of the theoretical disagreements as well.

In contrast to the standard classical approach to random effects models, a Bayesian hierarchical model does place an explicit distributional assumption on the random effects and coefficients. The traditional Laird and Ware (1982) model considers a Normal density for these parameters. While the normality assumption may at first appear restrictive, it is also appealing in that it implies the marginal distribution of $\ln p_{mtr}$ is also normal (compared to the conditional distribution above), and thus no more restrictive than the classical normal linear regression model with interaction terms and a specific form of heteroskedasticity. Still, out of concern that we may want to allow for a fatter tailed distribution, we have instead included the possibility that the random effects and coefficients are t-distributed with an a priori uncertain degrees of freedom to be estimated in the model. We also estimate separate coefficients for each restriction type, and specify that the market-time level covariates z_{mtr} are not random and do not interact with any market-varying or time-varying covariates.

The control variables that are included in d_{mr} are the entire set of market-varying covariates described above: log of the distance between the two cities in the market, the sizes of the origin and destination markets in terms of population and median household income, each in logs, and whether or not the origin market includes either a slot-controlled airport or a hub for one of the network carriers. The covariates in h_{tr} and q_{tr} are the same (though the data matrix H_r necessarily excludes the first time period), and include log of GDP and log of the average oil spot price for the quarter. Lastly, there is only one variable, difference in temperature between the two cities, that is specified to have a nonrandom coefficient and no interaction terms.

6 Results

The primary results of the study can be found in Table 3, which lists summary information for the posterior distributions of the marginal effects for each covariate. Specifically, Table 3 summarizes the posterior distributions of the marginal effects, which are best interpreted as the means of the posterior predictive densities conditional on a change in the given covariate, holding the remaining control variables constant at a specified value. The estimation output for all of the parameters in the model is provided in Table A3.[10]

Variation in the number of consumers using the Internet to research airline travel occurs solely over time in our data, which raises the concern that the marginal effects ascribed to our Internet variable may actually be some other economic time trend such as GDP. We address this concern by comparing Bayes Factors for the four different models that correspond to including or excluding Internet search activity and GDP from the regression equations. We calculate from the log-marginal likelihoods in Table 4 that the log-Bayes factor for the model which includes Internet search activity and excludes GDP versus the other three models is approximately 30, suggesting that any systematic trends in airline pricing are better described by variation in consumer adoption of Internet technology for fare research than by overall trends in the national economy. While other economic trends may coincide with Internet usage during our sample period, we feel that the most relevant ones would be highly correlated with GDP, and thus also strongly dominated by the effect of Internet usage.[11]

We find that increases in Internet search activity seem to be associated with higher *unrestricted* ticket fares with posterior odds in favor of a positive effect of 2.75:1. Additionally,

[10] It should be noted to what degree learning has occurred in the regression equation. Many of the covariates are associated with mean effects that exhibit low precision. Comparing Table 3 to Table A2 does suggest that significant learning has taken place with sharp reductions in posterior versus prior standard deviations. However, while it seems we can confidently say something about posterior mean effects, we are limited in what we learn about how the posterior predictive densities change for a given increase. To a certain extant, this is not surprising, since the posterior densities of the mean effects exhibit considerable uncertainty themselves, which when combined with the variance in the regression equation lead to posterior predictives that are in some cases very flat. We feel that with either a larger dataset that included more markets and/or time periods, or were we to include carrier level data and actual transacted (versus average) prices, the precision in our predictive densities would improve greatly.

[11] Marginal likelihoods are calculated using the method of Chib (1995). While not presented here for brevity concerns, inspection of the output from the model including both covariates yielded results for Internet search activity that are qualitatively similar to those presented here with GDP excluded.

we find that, while restricted ticket prices also seem to rise with more Internet searches, we are less sure about this effect, and in fact the probability that the effect is actually negative is 0.54. These two results combined suggest that as more people use the Internet to search for airline tickets, the divergence in fares actually increases. It is interesting that our findings do not coincide with any of the predictions in the consumer search literature, which may be due to low precision and not an actual economic finding. If we accept the evidence that unrestricted prices are rising, then to some extent the data support Lal and Sarvary's model that firms are able to exercise more monopoly power and better exploit consumer preferences for their brand as the share of informed consumers increases. It is also possible that what we are really picking up is the high correlation of Internet search activity with Internet purchasing activity. If anecdotal evidence in the industry is correct, then airlines are enjoying large cost savings as more travelers purchase their tickets online, and thus cost-based arguments, derived perhaps from the stochastic peak-load pricing literature, might better justify our findings.

Relative to our prior belief that unrestricted prices should be increasing with more competition on the route, we actually find strong evidence that increasing competition leads to *lower* unrestricted ticket fares. This is in sharp contrast to the finding by Stavins (2001) that unrestricted ticket prices are higher with more competition on the route, suggesting that in environments where a new competitor enters and is able to acquire significant market share from the incumbent firms, he does so by competing for all ticket levels and not just restricted tickets.[12] Importantly, we also find that restricted ticket fares are also falling with the introduction of more competition on routes, and furthermore that restricted ticket prices fall much further as a percentage than unrestricted tickets. This finding corroborates earlier work on the relationship between price dispersion and market concentration. We show that the posterior mean of the effect of HHI falling from 1/2 to 1/3 is for the ratio of prices to increase by 5%; additionally we are confident in the sign of this result with a posterior probability that the effect is positive of 0.99. Combining all three elements of the effect of competition on fares, we see that the evidence favors the predictions in the brand-intensive nonlinear pricing models of

[12] This seems like a reasonable finding since even the low fare carriers like Southwest and JetBlue offer unrestricted tickets at a discount relative to the network carriers.

Borenstein (1985) and Stole (1995). We further note that the consumer search models that provide a predicted effect for an increase in competitiveness each state that unrestricted prices should be increasing or at least not changing, contrasting sharply with our strong findings to the contrary.

To fix ideas explicitly on how Internet search activity has affected the mean effect of competition on prices, we consider a setting where oil spot prices are fixed at their median value of $22 and look at the mean effect of going from 2 to 3 competitors conditioned on the actual size of the online population researching airline travel in the second quarter for each of 1998, 2000, and 2002. Figure 2 provides kernel density estimates of the posterior distributions for this marginal effect in each time period. The most striking result is how Internet search activity has affected the price ratio for 1998 relative to later time periods. From Q2/1998 to Q2/2000, the number of consumers researching airline travel on the Internet grew from 8.3 million to 49.8 million, or 500%, and is associated with a fall in the effect of increasing competition on price dispersion of 5%. Thus, it would seem that prior to the Internet becoming a major search tool for travelers, a fall in HHI from 1/2 to 1/3 was associated with an estimated 10% increase in the ratio of unrestricted to restricted ticket prices. Since then, and we argue almost solely due to the increasing role of the Internet in searching for airline tickets, this same fall in market concentration is now associated with an estimated 5% increase in the ratio of prices. It would seem clear that at least in this respect, the Internet has been associated with a clear change in the competitive environment for airlines.

Table 5 provides more detail on how increasing Internet search activity is associated with the change in fares for each fare type from more competition in a market. Specifically, we see that with more people searching for airline travel online, unrestricted ticket prices have become more sensitive to competition, while restricted ticket prices have not measurably changed in their sensitivity to competition. If we are willing to extrapolate beyond our sample, a back-of-the-envelope calculation using the mean effects in Table A3 predicts that if there were approximately 140 million users researching airline travel online per quarter, the unrestricted market segment would actually be more responsive to firm entry than the restricted ticket market.[13] Conversely, we can also say something about the competitive environment in the very early stages of the Internet as well, as a similar calculation suggests that at levels of Internet search activity below approximately 3.5 million Internet airline travel researchers per quarter, unrestricted ticket prices are predicted to increase with more competition in a market. This finding is actually consistent with that of Stavins (2001), who looked at airline fares in 1995 when the Internet was a negligible factor in airline pricing.

These findings present something of a quandary for describing the economic mechanism behind falling search costs. Overall, Internet search activity is associated with higher prices for unrestricted tickets, yet the results in Table 5 suggest that Internet search activity does induce greater brand rivalry. The only consumer search model that predicts an increase in prices for the high-quality good is that of Lal and Sarvary, but the mechanism through which this occurs is through a lessening of brand rivalry via more effective product differentiation. While the argument that falling search costs better enables firms to target their products at specific market segments is appealing, when we look at the cross-effect of the Internet on the effect of market concentration we see a statistically and economically strong finding that the Internet increases the intensity of brand-rivalry. We interpret the evidence to suggest that as search costs fall, the market for air travel will become more and more like the economy described in Rochet and Stole, where entry leads firms to compete more aggressively for the high end of the market than the low end. While search costs continue to be important, however, none of the consumer search models seem to accurately describe the airline industry. On the other hand, the models of both Borenstein and Stole imply that horizontal (brand) uncertainty dominates in the airline industry as we see it throughout most of our study period.

The remainder of our control variables seem to exhibit limited explanatory power for the variation in prices. From the perspective that the price ratio is potentially just a byproduct of differences in marginal costs across markets, we find reasonably strong support that distance, which should be a reasonable proxy for marginal costs, is associated with higher prices for both

[13] As of Winter 2003, the size of the adult Online population was 137.6 million according to Nielsen, although only roughly 50% of this population was searching for airline travel online.

ticket types, and additionally is associated with a higher price ratio. The marginal effects of the presence of slot-controlled airports in the home market provide additional support for a cost-based argument, though peculiarly the predicted increase in prices is symmetric across fare types. These relatively strong findings are contradicted by the finding that oil spot prices are associated with lower levels of price dispersion, where in fact the posterior odds of a negative versus positive effect on the price ratio for a 1% increase in the oil spot price is 2:1. Moreover, we note that network hubs are potential proxies for lower costs and both unrestricted and restricted ticket prices tend to be lower in origin markets with network hubs than in those without.

Origin population, which we initially consider to be a demand shifter, actually seems to be associated with lower ticket prices, especially for unrestricted tickets, which leads us to suspect that airlines may enjoy large cost benefits from large origin markets, possibly as a result of consistently higher levels of operations. As with most of this analysis, however, a more structural approach is necessary to attach any firm interpretation to this result. Origin income does seem to be a positive demand shifter in raising unrestricted and restricted ticket prices, although it is not clear what effect increasing household incomes by 1% has on the ratio of prices since both fare types seem to increase by an equal amount. Interestingly, if we claim that origin income is a valid demand proxy, our finding that it seems to affect both fare types equally would support the idea that some portion of the price difference results from stochastic pricing mechanisms.

7 Summary

This paper takes a fresh look at price discrimination by asking how search costs, and in particular falling search costs, affect the competitive environment in the airline industry. We claim that increasing Internet search activity is an indicator of either decreasing search costs or a rising share of low search-cost consumers, and postulate that the existing literature on consumer search and price discrimination should provide insights into how rising Internet use may affect the price schedules offered by firms. Our primary identification scheme is to interact

Internet search activity with a measure of market concentration to examine the cross effect of the Internet on the effect of competition on fares. We find evidence that the Internet seems to intensify competition, with the airline industry becoming more and more represented by the economic environment described by Rochet and Stole (2002). While search costs remain significant, however, the market segmentation practiced by airlines seems most consistent with models of brand-intensive competition à la Borenstein (1985) and Stole (1995).

Quantifying the effect of the Internet may also help to provide insight into the current spell of financial instability in the airline industry. Hayes and Ross (1998) claim that the financial difficulties experienced by the airlines in the early 1990's may have been caused by an increase in the level of competitiveness in the industry, especially as the share of low-cost carriers like Southwest was aggressively increasing. Currently, however, the roots of poor profitability may lie elsewhere. If indeed the Internet is decreasing search costs for a growing proportion of consumers, then it may be that the airlines' existing fencing strategies will be harder to maintain as a mechanism for extracting extra surplus from consumers. Such an argument would be in line with the Bakos (1997) model, which asserts that in markets characterized by product differentiation, as search costs fall, firms are less able to exploit consumer search costs to charge monopoly prices to different market segments.

Appendix

Our hierarchical regression model takes the form

$$y_{mtr} = \ln p_{mtr} \overset{ind}{\sim} N\left(\mu_{mr} + \tau_{tr} + x_{mtr}\alpha_{tr} + z_{mtr}\beta_r, \sigma_y^2\right), \tag{5}$$

$$\mu_{mr} \overset{ind}{\sim} N\left(d_{mr}\zeta_r, \omega_{mr}\sigma_\mu^2\right) \text{ for } m = 1, ..., M, \tag{6}$$

$$\tau_{tr} \overset{ind}{\sim} N\left(h_{tr}\xi_r, \psi_{tr}\sigma_\tau^2\right) \text{ for } t = 2, ..., T, \tag{7}$$

$$\alpha_{tr} \overset{ind}{\sim} t\left(q_{tr}\delta_r, \sigma_\alpha^2, \nu\right) \text{ for } t = 1, ..., T, \tag{8}$$

$$\omega_{mr}^{-1} \overset{iid}{\sim} G\left(\frac{\upsilon}{2}, \frac{2}{\upsilon}\right), \tag{9}$$

$$\psi_{tr}^{-1} \overset{iid}{\sim} G\left(\frac{\varphi}{2}, \frac{2}{\varphi}\right), \tag{10}$$

and

$$\lambda_{tr}^{-1} \overset{iid}{\sim} G\left(\frac{\nu}{2}, \frac{2}{\nu}\right). \tag{11}$$

We estimate each of the random parameters as independent Student-t distributed with unknown degrees of freedom parameters. Following Geweke (1993), the estimation of the degrees of freedom parameters is accomplished with a scale mixture of normals specification that is easily accommodated within a hierarchical specification in the Gibbs sampling routine.

A.1 Predictive Densities for the Marginal Effects

In order to describe the predictive distributions for parameters of interest, we take advantage of the simple interpretations inherent in the model and define a large set of statistical quantities. Recalling that previous studies have devoted much attention to the effect of market concentration on prices, let $-6\eta_{mtr}^{HHI} = \partial E_{y|\Gamma}\left[y_{mtr}\right]/\partial HHI_{mt}$, where the -6 comes from noting that moving from 2 firms exactly sharing the market to 3 firms exactly sharing the market implies a change in HHI from 1/2 to 1/3, or a difference of $-1/6$. Hence, $\eta_{mtr}^{HHI} = -\alpha_{tr}/6$ and it should be noted that even when we condition on the unknown model parameters Γ, η_{mtr}^{HHI} is a random variable as per the specification above. The simple linear nature of this parameter of interest is useful in describing the price ratio as well. Note that

$\partial E_{y|\Gamma}\left[\ln\left(p_{mtU}/p_{mtR}\right)\right]/\partial HHI_{mt} = \eta_{mtU}^{HHI} - \eta_{mtR}^{HHI}$, which is interpreted as describing the percent change in the ratio of unrestricted to restricted ticket prices for a given change in HHI. Since the specification is linear in logs, we can define the quantity $\eta_{mt}^{HHI} = -\left(\alpha_{tU} - \alpha_{tR}\right)/6$, which is thus distributed $t\left(-\left(q_{tU}\delta_U - q_{tR}\delta_U\right)/6, \sigma_\alpha^2/18, \nu\right)$. This specification for measuring the effect of market concentration on prices allows the effect to vary across time with an observed conditional mean that depends on Internet search activity, and also allows for unobserved variation across time as well through the distributional assumption on the coefficient.

In addition to describing the effect of market concentration on airline fares and allowing the number of Internet airline travel researchers to influence it, our specification allows us to measure the direct impact of Internet search activity on prices as well. To this end, let $\eta_{mtr}^{Int} = \partial E_{y|\Gamma}\left[y_{mtr}\right]/\partial \ln Int_t$, where the log specification for Internet search activity suggests an interpretation for η_{mtr}^{Int} as the expected percent increase in price for ticket type r for a one percent increase in the number of people researching airline travel online. As before, we can also define the effect of Internet search activity on the price ratio as $\eta_{mt}^{Int} = \eta_{mtU}^{Int} - \eta_{mtR}^{Int}$. In contrast to the effect of market concentration on prices, the effect of Internet search activity is assumed to be an unknown fixed quantity. The specification does allow the effect of Internet activity to depend on market concentration, but unlike the effect of market concentration, there is no unobserved variation across either time or market.

A.2 Prior Analysis

We make use of conjugate priors wherever possible to further simplify the estimation algorithm, which can be done for all parameters save the degrees of freedom parameters υ, φ, and ν. Our prior specification is thus

$$\beta_r \overset{ind}{\sim} N\left(\underline{\beta}_r, \underline{V}_\beta^r\right), \tag{12}$$

$$\zeta_r \overset{ind}{\sim} N\left(\underline{\zeta}_r, \underline{V}_\zeta^r\right), \tag{13}$$

$$\xi_r \overset{ind}{\sim} N\left(\underline{\xi}_r, \underline{V}_\xi^r\right), \tag{14}$$

$$\delta_r \overset{ind}{\sim} N\left(\underline{\delta}_r, \underline{V}_\delta^r\right), \tag{15}$$

$$\sigma_u^{-2} \sim G\left(\underline{\kappa}_1, \underline{\kappa}_2\right), \tag{16}$$

$$\sigma_\mu^{-2} \sim G\left(\underline{\pi}_1, \underline{\pi}_2\right), \tag{17}$$

$$\sigma_\tau^{-2} \sim G\left(\underline{\gamma}_1, \underline{\gamma}_2\right), \tag{18}$$

$$\sigma_\alpha^{-2} \sim G\left(\underline{\chi}_1, \underline{\chi}_2\right), \tag{19}$$

$$\upsilon \sim p\left(\upsilon\right), \tag{20}$$

$$\varphi \sim p\left(\varphi\right), \tag{21}$$

and

$$\nu \sim p\left(\nu\right). \tag{22}$$

Due to the wide variation in predicted theoretical outcomes for the effect of consumer search costs on prices, we choose to be noninformative and nondogmatic in our prior beliefs about the effects of Internet search activity. Hence $\eta_{mt\cdot}^{Int}$ and each of the components in Γ^{Int} are centered at zero in the prior specification and variances are chosen so that the direct effect of the Internet on prices is mostly limited to the interval $[-.5, 1]$, or that the effect of the Internet is a priori uncertain enough to allow for anywhere from a halving of prices up to a doubling of prices. We feel that this specification is in line with the inductive nature of our research on the effect of the Internet—we want to test whether or not Internet search activity has had any measurable impact, and conditional that it has, what do the data inform us about the effect?

Given the vast literature on the effect of market concentration on price dispersion, we have strong prior information on how to parameterize the effect of increasing competition on prices. Stavins (2001) found that increasing competition, as measured by a falling Herfindahl index, was associated with a decrease in discounted tickets (advanced purchases and Saturday-night stayover restrictions), while unrestricted tickets were found to correspond with higher prices for a given decrease in market concentration. Taken together, these imply the result consistent with Borenstein and Rose (1994) and others that price dispersion is increasing with more competition on routes. We incorporate these results into our prior specification so that, conditional on likely values for the time-varying coefficients, an increase in the number of competitors in a market is associated with a 10% decrease in restricted ticket prices and a 10% increase in unrestricted ticket prices, and specify our degree of belief in the signs of these effects as 0.63 for the effect on restricted tickets and 0.66 for the effect on unrestricted tickets. Moreover, our specification is consistent with the literature in that the ratio of unrestricted to restricted ticket prices is centered at an increase of 20% as the number of competitors goes from two to three, though we have been reasonably nondogmatic about this effect with a standard deviation of 54% and a prior degree of belief in the effect on the ratio of prices being positive of 0.64.

Table A2 summarizes the prior distributions of the expected marginal effects, and illustrates the informative yet nondogmatic nature of our prior specification. In summary, we feel that GDP, origin and destination market population sizes, origin and destination market median household incomes, and positive differences in destination minus origin temperatures will all increase equilibrium prices through an outward demand shift. We also feel that higher oil spot prices, longer distances between markets, and slot controlled airports increase equilibrium prices through an increase in marginal costs for airlines, while the presence of a network hub probably lowers costs for airlines and thus leads to lower equilibrium fares.[14]

The variance parameter associated with the first level of the hierarchy has parameters $(\underline{\kappa}_1, \underline{\kappa}_2) = (2.25, 1.25)$, which corresponds to a mean expectation for $\sigma_y^2 = 0.8^2$ and a prior belief in the fit of the regression model corresponding to a R^2 centered at 0.25. We wish to be nondogmatic about the presence of unidentified heterogeneity in the random effects and coefficients, which leads us to favor values near zero for the prior distribution of the corresponding variance terms. For the variances of the market-level random coefficients, $(\underline{\pi}_1, \underline{\pi}_2) = (2.25, 80)$, which corresponds to an expected value for σ_μ^2 of 0.01. Both of the variance terms for the time-level random effects and random coefficients have prior means of 0.04, achieved through the following parameterization: $\left(\underline{\gamma}_1, \underline{\gamma}_2\right) = \left(\underline{\chi}_1, \underline{\chi}_2\right) = (2.25, 20)$. For all of the variance parameter specifications, we have assumed that $\sqrt{Var\left(\sigma_\cdot^2\right)} = 2E\left(\sigma_\cdot^2\right)$, which seems to be a reasonably diffuse parameterization. Finally, the prior distribution for the degrees of freedom parameters

[14] Note that we have allowed the prior specification to vary across the different models considered for hypothesis testing in the paper. Specifically, coefficient prior means (especially the constants) were determined jointly, so that the removal of any given covariate from the model implies different prior values for the remaining coefficients' prior specifications. Details for all parameterizations are available upon request from the authors.

are uniform over the interval $[2, 32]$ and discretized with a grid length equal to 0.0001, which is fine enough to allow very precise estimation of the posterior conditional density while still allowing computational tractability through the Gibbs algorithm.[15]

A.3 Estimation of the Posterior Density

Conditional on the scale parameters ω_{mr}, ψ_{tr}, and λ_{tr}, estimation of the hierarchical model proceeds straightforwardly via the Gibbs algorithm described in Chib and Carlin (1999). Moreover, since the distribution of the degrees of freedom parameters is discretized, the entire algorithm is still estimable via the Gibbs algorithm by simply extending it to allow for draws of the scale parameters and subsequently the discrete conditional densities of the degrees of freedom parameters.

Following Geweke (1993), and without loss of generality, consider the posterior distribution for ω_{mr} given draws for ζ_r, σ_μ^2 and υ. We have that

$$\begin{aligned} f\left(\omega_{mr}|y, \Gamma_{-\omega_{mr}}\right) &\propto \phi\left(\mu_{mr}; d_{mr}\zeta_r, \omega_{mr}\sigma_\mu^2\right) f_G\left(\omega_{mr}^{-1}; \upsilon/2, 2/\upsilon\right) \\ &= \omega_{mr}^{-\left(\frac{\upsilon+1}{2}-1\right)} \exp\left(-\frac{1}{2\omega_{mr}}\left[\left(\frac{\mu_{mr}-d_{mr}\zeta_r}{\sigma_\mu}\right)^2 + \upsilon\right]\right) \end{aligned}, \tag{23}$$

which implies that $\omega_{mr}^{-1}|y, \Gamma_{-\omega_{mr}} \sim G\left(\frac{\upsilon+1}{2}, 2/\left[\left(\frac{\mu_{mr}-d_{mr}\zeta_r}{\sigma_\mu}\right)^2 + \upsilon\right]\right)$. Subsequently, the posterior conditional kernel density for υ is

$$f\left(\upsilon|y, \Gamma_{-\upsilon}\right) = p\left(\upsilon\right)\left[\Gamma\left(\upsilon/2\right)\left(2/\upsilon\right)^{(\upsilon/2)}\right]^{-MR} \exp\left(-\frac{\upsilon}{2}\sum_{m,r}\left(\log\omega_{mr} + \omega_{mr}^{-1}\right)\right). \tag{24}$$

Since there is no conjugate distribution for υ, the parameter space is discretized and the Gibbs draw is taken directly by normalizing $f\left(\upsilon|y, \Gamma_{-\upsilon}\right)$ to sum to one. The remaining parameters in the chain are drawn based on the traditional hierarchical linear model results described in Chib and Carlin.

[15] Alternatively, we could specify that the degrees of freedom parameter is itself Gamma distributed (à la Geweke 1993). This would introduce a Metropolis-within-Gibbs step in the estimation algorithm and would also necessitate a more complicated estimation routine for calculating the marginal likelihoods for each of the models considered.

References

Bakos, J. Yannis, "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science* 43 (December 1997), 1976-1993.

Borenstein, Severin., "Price Discrimination in Free-Entry Markets" *RAND Journal of Economics* 16 (Autumn 1985), 380-397.

——, "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins," *RAND Journal of Economics* 22 (Autumn 1991), 354-369.

Borenstein, Severin and Nancy L. Rose, "Competition and Price dispersion in the U.S. Airline Industry," *Journal of Political Economy* 102 (August 1994), 653-683.

Busse, Meghan and Marc Rysman, "Competition and Price Discrimination in Yellow Pages Advertising," *Working Paper*, U.C. Berkeley Haas School of Business and Boston University (June 2001).

Chib, Siddhartha, "Marginal Likelihood from the Gibbs Output" *Journal of the American Statistical Association* 90 (December 1995), 1313-1321.

Chib, Siddhartha and Bradley P. Carlin, "On MCMC Sampling in Hierarchical Longitudinal Models" *Statistics and Computing* 9 (1999), 17-26.

Clerides, Sofronis K., "Book Value: Intertemporal Pricing and Quality Discrimination in the U.S. Market for Books," *International Journal of Industrial Organization* 20 (December 2002), 1385-1408.

Cohen, Andrew, "Package Size and Price Discrimination in Paper Towels," *Mimeo*, University of Virginia (September 2000).

Dana, James D., "Advance-Purchase discounts and Price Discrimination in Competitive Markets," *Journal of Political Economy* 106 (April 1998), 395-422.

——, "Using Yield Management to Shift Demand When Peak time is Unknown," *RAND Journal of Economics* 30 (Autumn 1999a), 456-474.

——, "Equilibrium Price Dispersion Under Demand Uncertainty: The Roles of Costly Capacity and Market Structure," *Rand Journal of Economics* 30 (Winter 1999b), 632-660.

Diamond, Peter, "A Model of Price Adjustment," *Journal of Economic Theory* 3 (June 1971), 156-168.

Gale, Ian L. and Thomas J. Holmes, "The Efficiency of Advance-Purchase Discounts in the Presence of Aggregate Demand Uncertainty" *International Journal of Industrial Organization* 10 (September 1992), 413-437.

——, "Advance-Purchase Discounts and Monopoly Allocation of Capacity," *The American Economic Review* 83 (May 1993), 135-146.

Geweke, John, "Bayesian Treatment of the Independent Student-t Linear Model," *Journal of Applied Econometrics* 8 (December 1993), S19-S40.

Hayes, Kathy J. and Leola B. Ross, "Is Airline Price Dispersion the Result of Careful Planning or Competitive Forces?" *Review of Industrial Organization* 13 (October 1998), 523-541.

Katz, Michael L., "Price Discrimination and Monopolistic Competition," *Econometrica* 52 (November 1984), 1453-1472.

Laird, N.M. and J.H. Ware, "Random Effects Models for Longitudinal Data," *Biometrics* 38 (1982), 963-974.

Lal, Rajiv and Miklos Sarvary, "When and How is the Internet Likely to Decrease Price Competition?" *Marketing Science* 18 (1999), 485-503.

Lane, Leonard D., "An Investigation of the Direct and Ripple Effect of an Environmental Jolt on the Airline Industry," *Mimeo*, Case Western Reserve University (2002).

Lane, Leonard D., "Price Discrimination in the U.S. Domestic Airline Industry: The Effect of the Internet," *Mimeo*, Case Western Reserve University (2003).

Li, Michael Z.F., "Pricing Non-Storable Perishable Goods by Using a Purchase Restriction with an Application to Airline Fare Pricing" *European Journal of Operational Research* 134 (November 2001), 631-647.

Liu, Qihong, "The Effect of Market Structure on Price Dispersion: An Analysis of the U.S. Airline Industry," *Working Paper*, S.U.N.Y. Stony Brook (April 2003).

McManus, Brian, "Nonlinear Pricing in an Oligopoly Market: the Case of Specialty Coffee," *Mimeo*, Washington University Olin School of Business (May 2003).

Rochet, Jean-Charles and Lars A. Stole, "Nonlinear Pricing with Random Participation," *Review of Economic Studies* 69 (January 2002), 277-311.

Salop, Steven, "The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination," *The Review of Economic Studies* 44 (October 1977), 393-406.

Salop, Steven and Joseph Stiglitz, "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *The Review of Economic Studies* 44 (October 1977), 493-510.

Shepard, Andrea, "Price Discrimination and Retail Configuration," *The Journal of Political Economy* 99 (February 1991), 30-53.

Stahl, Dale O. II, "Oligopolistic Pricing with Sequential Consumer Search," *The American Economic Review* 79 (September 1989), 700-712.

Stavins, Joanna, "Price Discrimination in the Airline Market: The Effect of Market Concentration," *Review of Economics and Statistics* 83 (February 2001), 200-202.

Stole, Lars A., "Nonlinear Pricing and Oligopoly" *Journal of Economics & Management Strategy* 4 (1995), 529-562.

Varian, Hal R. "Price Discrimination," in *Handbook of Industrial Organization* 1 edited by Richard Schmalense and Robert Willig (Amsterdam: North Holland, 1989), 598-654.

Tables and figures

Table 1

Theoretical Predictive Effects of Increasing Internet Use and Entry on Airline Fares

	Internet Usage			Entry		
Model	Unrestricted	Resticted	Ratio	Unrestricted	Restricted	Ratio
Consumer Search						
Salop & Stiglitz	-	0	-	+	0	+
Katz	0	-	+	0	-	+
Stahl	-	-	?	+	-	+
Bakos	-	-	-	?	?	?
Lal & Sarvary	+	+	?	?	?	?
Price Discrimination						
Borenstein	-	-	-/+	-	-	-/+
Stole	-/0	-/-	-/+	-/0	-/-	-/+
Rochet & Stole	-	-	-	-	-	-

Table 2

Descriptive Statistics for Regression Analysis

Variable	Mean	Std. Dev.	Min	Max
Unrestricted Fare ($)	422	195	87	1,029
Restricted Fare ($)	168	57	51	401
HHI	0.59	0.27	0.21	1.00
Internet Airline Search Activity (millions of users)	45	19	2	65
GDP (1996, in $trillions)	9.0	0.3	8.4	9.4
Oil Spot Price ($/BBL)	22	6	13	32
Distance (000's miles)	1.1	0.7	0.2	2.7
Origin Income (median, in $000's)	45	12	29	69
Destination Income (median, in $000's)	44	11	25	69
Origin Population (millions)	3.2	3.2	0.3	9.5
Destination Population (millions)	1.7	1.9	0.0	9.5
Difference in Temperature (Fahrenheit)	-4	15	-44	43
Slot Controlled Airport in Origin Market	0.44	-	0	1
Network Hub in Origin Market	0.12	-	0	1

Table 3

*Posterior Marginal Effect of Covariates on Fares**

	%Δ Unrestricted Ticket Fares			%Δ Restricted Ticket Fares			%Δ Ratio of Fares		
Covariate[†]	Mean	Std. Dev.	Pr(• > 0 | D)	Mean	Std. Dev.	Pr(• > 0 | D)	Mean	Std. Dev.	Pr(• > 0 | D)
2 to 3 competitors	-4.41%	1.64%	< 0.01	-9.47%	1.49%	< 0.01	5.06%	2.23%	0.99
Internet Searches	4.80%	6.61%	0.77	0.83%	6.82%	0.46	3.97%	9.43%	0.73
Oil Spot Price	-0.17%	9.82%	0.50	5.80%	10.06%	0.72	-5.98%	13.89%	0.33
Distance	41.68%	11.79%	1.00	27.62%	10.42%	0.99	14.06%	15.48%	0.82
Origin Population	-10.15%	7.70%	0.10	-0.66%	7.09%	0.46	-9.50%	10.44%	0.18
Destination Population	7.29%	6.65%	0.86	4.21%	5.92%	0.77	3.08%	8.73%	0.65
Origin Income	16.20%	20.03%	0.79	22.40%	20.02%	0.87	-6.20%	28.17%	0.42
Destination Income	-2.72%	21.28%	0.44	-11.21%	20.13%	0.29	8.49%	28.95%	0.62
Slot Controlled Airport in Origin Market	12.03%	18.56%	0.74	11.71%	18.36%	0.74	0.32%	26.16%	0.50
Network Hub in Origin Market	-13.70%	12.22%	0.13	-15.00%	11.73%	0.10	1.30%	17.02%	0.53
Difference in Temperature	-0.04%	0.11%	0.37	-0.02%	0.09%	0.40	-0.01%	0.15%	0.46

* Results reflect the posterior distribution of the mean of the marginal effect on the predictive density of fares.

† Results for HHI are conditioned on median values for Internet Searches and Oil Spot Price. Results for Internet Searches and Oil Spot Price are conditioned on HHI = 0.5. The marginal effects of the remaining covariates are independent of any other covariate. Most covariates are in logs, so that results are reflective of a 1% increase in the corresponding covariate. HHI is in base units, the Difference in Temperature is in degrees Fahrenheit, and both the Slot and Hub variables are dummy variables.

Table 4

Model Comparisons

Model	Excluded Variables	Log-Marginal Likelihood	$Pr(M_i \mid Data)$
1	*None*	406.85	< 0.01
2	*GDP*	441.26	> 0.99
3	*Internet Search Activity*	431.27	< 0.01
4	*GDP, Internet Search Activity*	418.38	< 0.01

Table 5

*Mean effect of going from 2 to 3 firms**

	$Int_2 = 8.3\,million$	$Int_{20} = 64.9\,million$	$\Pr(\eta_{20} - \eta_2 > 0)$
$\%\Delta$ *Unrestricted Fares*	-1.39% (2.59%)	-4.74% (2.81%)	0.115
$\%\Delta$ *Restricted Fares*	-10.00% (3.05%)	-9.34% (2.92%)	0.578
$\%\Delta$ *Ratio of Fares*	8.61% (3.87%)	4.60% (4.06%)	0.171

* The first number is the posterior mean. The number in parentheses is the posterior standard deviation.

Table A1

Markets and Airports

ID	Origin Market	Destination Market	Origin Airports	Destination Airports
1	Atlanta	Portland	ATL	PDX
2	Boston	San Francisco	BOS	SFO,SJC
3	Boston	Cleveland	BOS	CLE
4	Boston	St. Louis	BOS	STL
5	Dallas	Memphis	DFW	MEM
6	Denver	Houston	DEN	IAH
7	Detroit	Minneapolis	DTW	MSP
8	Houston	Chicago	IAH	MDW,ORD
9	Los Angeles	Dallas	BUR,LAX,LGB,ONT,SNA	DAL,DFW
10	New York	Los Angeles	EWR,JFK	LAX,LGB,ONT,SNA
11	Philadelphia	Chicago	PHL	MDW,ORD
12	Pittsburg	Orlando	PIT	MCO
13	St. Louis	Washington D.C.	STL	DCA,IAD
14	San Francisco	Seattle	OAK,SFO	SEA
15	Seattle	Chicago	SEA	MDW,ORD
16	Raleigh	Washington D.C.	RDU	DCA,IAD
17	Los Angeles	San Francisco	BUR,LAX,ONT,SNA	OAK,SFO,SJC
18	San Francisco	Denver	OAK,SFO,SJC	DEN
19	Seattle	Anchorage	SEA	ANC
20	New York	Chicago	EWR,JFK,LGA	MDW,ORD
21	New York	Miami	EWR,LGA	MIA
22	Detroit	Orlando	DTW	MCO
23	St. Louis	San Diego	STL	SAN
24	Chicago	Phoenix	MDW,ORD	PHX
25	Los Angeles	Honolulu	BUR,LAX,ONT,SNA	HNL

Table A2

*Prior Marginal Effect of Covariates on Fares**

Covariate†	%Δ Unrestricted Ticket Fares Mean	St. Dev.	Pr(• > 0 | D)	%Δ Restricted Ticket Fares Mean	St. Dev.	Pr(• > 0 | D)	%Δ Ratio of Fares Mean	St. Dev.	Pr(• > 0 | D)
2 to 3 competitors	10.00%	24.11%	0.66	-10.00%	29.47%	0.37	20.00%	53.58%	0.64
Internet Population	0.00%	26.89%	0.50	0.00%	32.36%	0.50	0.00%	59.25%	0.50
GDP	2.00%	27.43%	0.53	1.00%	32.68%	0.51	1.00%	60.11%	0.51
Oil Spot Price	1.00%	27.16%	0.51	1.00%	32.68%	0.51	0.00%	59.84%	0.50
Distance	1.00%	27.16%	0.51	1.00%	32.68%	0.51	0.00%	59.84%	0.50
Origin Population	1.00%	27.16%	0.52	2.00%	33.01%	0.52	-1.00%	60.17%	0.49
Destination Population	1.00%	27.16%	0.52	2.00%	33.01%	0.52	-1.00%	60.17%	0.49
Origin Income	1.00%	27.16%	0.52	2.00%	33.01%	0.52	-1.00%	60.17%	0.49
Destination Income	1.00%	27.16%	0.52	2.00%	33.01%	0.52	-1.00%	60.17%	0.49
Slot Controlled Airport in Origin Market	2.00%	27.43%	0.53	2.00%	33.01%	0.52	0.00%	60.44%	0.50
Network Hub in Origin Market	-2.00%	26.35%	0.47	-2.00%	31.71%	0.47	0.00%	58.06%	0.50
Difference in Temperature	0.10%	26.92%	0.50	0.20%	32.43%	0.50	-0.10%	59.35%	0.50

* Results reflect the prior distribution of the mean of the marginal effect on the predictive density of fares.

† Results for HHI are conditioned on prior conceived typical values for Internet Search Activity, GDP, and Oil Spot Price. Results for Internet Search Activity, GDP, and Oil Spot Price are conditioned on HHI = 0.5. The marginal effects of the remaining covariates are independent of any other covariate. Most covariates are in logs, so that results are reflective of a 1% increase in the corresponding covariate. HHI is in base units, the Difference in Temperature is in degrees Fahrenheit, and both the Slot and Hub variables are dummy variables.

Table A3

*Posterior Distribution of Regression Parameters**

	Variable	Posterior Distribution Mean	Std. Dev.	Median	Pr(• > 0 | D)
First Stage					
Unrestricted Fares	*DTemp*	0.00038	0.00114	0.00036	0.6263
Restricted Fares	*DTemp*	0.00023	0.00094	0.00024	0.5990
	Variance	0.02804	0.00140	0.02800	1.0000
Second Stage - Market-level Random Effects					
Unrestricted Fares	*Constant*	5.5390	1.1343	5.5577	> 0.9999
	Distance	0.4168	0.1179	0.4248	0.9996
	Origin Pop.	-0.1015	0.0770	-0.1028	0.0957
	Destination Pop.	0.0729	0.0665	0.0736	0.8646
	Origin Inc.	0.1620	0.2003	0.1628	0.7903
	Destination Inc.	-0.0272	0.2128	-0.0313	0.4414
	Origin Slot	0.1203	0.1856	0.1164	0.7382
	Origin Hub	-0.1370	0.1222	-0.1360	0.1320
Restricted Fares	*Constant*	4.3126	1.0218	4.3137	> 0.9999
	Distance	0.2762	0.1042	0.2794	0.9933
	Origin Pop.	-0.0066	0.0709	-0.0065	0.4613
	Destination Pop.	0.0421	0.0592	0.0432	0.7701
	Origin Inc.	0.2240	0.2002	0.2242	0.8668
	Destination Inc.	-0.1121	0.2013	-0.1095	0.2917
	Origin Slot	0.1171	0.1836	0.1162	0.7364
	Origin Hub	-0.1500	0.1173	-0.1520	0.1002
	Variance	0.0687	0.0284	0.0674	1.0000
	Degrees of Freedom	10.7380	8.7266	7.2218	1.0000

* Dependent variable is ln(fare).

Table A3 Continued

*Posterior Distribution of Regression Parameters**

	Variable	Posterior Distribution Mean	Std. Dev.	Median	Pr(• > 0 | D)
Second Stage - Time-level Random Effects					
Unrestricted Fares	*Constant*	-0.0308	0.1849	-0.0349	0.4298
	Internet Search Activity	-0.0008	0.0666	-0.0011	0.4941
	Oil Spot Price	-0.0194	0.0964	-0.0180	0.4263
Restricted Fares	*Constant*	-0.1191	0.1890	-0.1242	0.2636
	Internet Search Activity	0.0033	0.0688	0.0047	0.5297
	Oil Spot Price	0.0546	0.0990	0.0552	0.7085
	Variance	0.0074	0.0025	0.0070	1.0000
	Degrees of Freedom	12.3350	8.4067	9.8794	1.0000
Second Stage - Time-level Random Coefficients on Market Concentration					
Unrestricted Fares	*Constant*	-0.2290	0.2238	-0.2321	0.1498
	Internet Search Activity	0.0975	0.0454	0.0968	0.9851
	Oil Spot Price	0.0353	0.0685	0.0351	0.6972
Restricted Fares	*Constant*	0.6384	0.2257	0.6368	0.9973
	Internet Search Activity	-0.0233	0.0442	-0.0235	0.2890
	Oil Spot Price	0.0069	0.0694	0.0074	0.5395
	Variance	0.0086	0.0033	0.0080	1.0000
	Degrees of Freedom	9.4713	7.6225	6.6224	1.0000

* Dependent variable is ln(fare).

Figure 1: Plot of average prices (across markets) versus the number of consumers researching airline travel online.

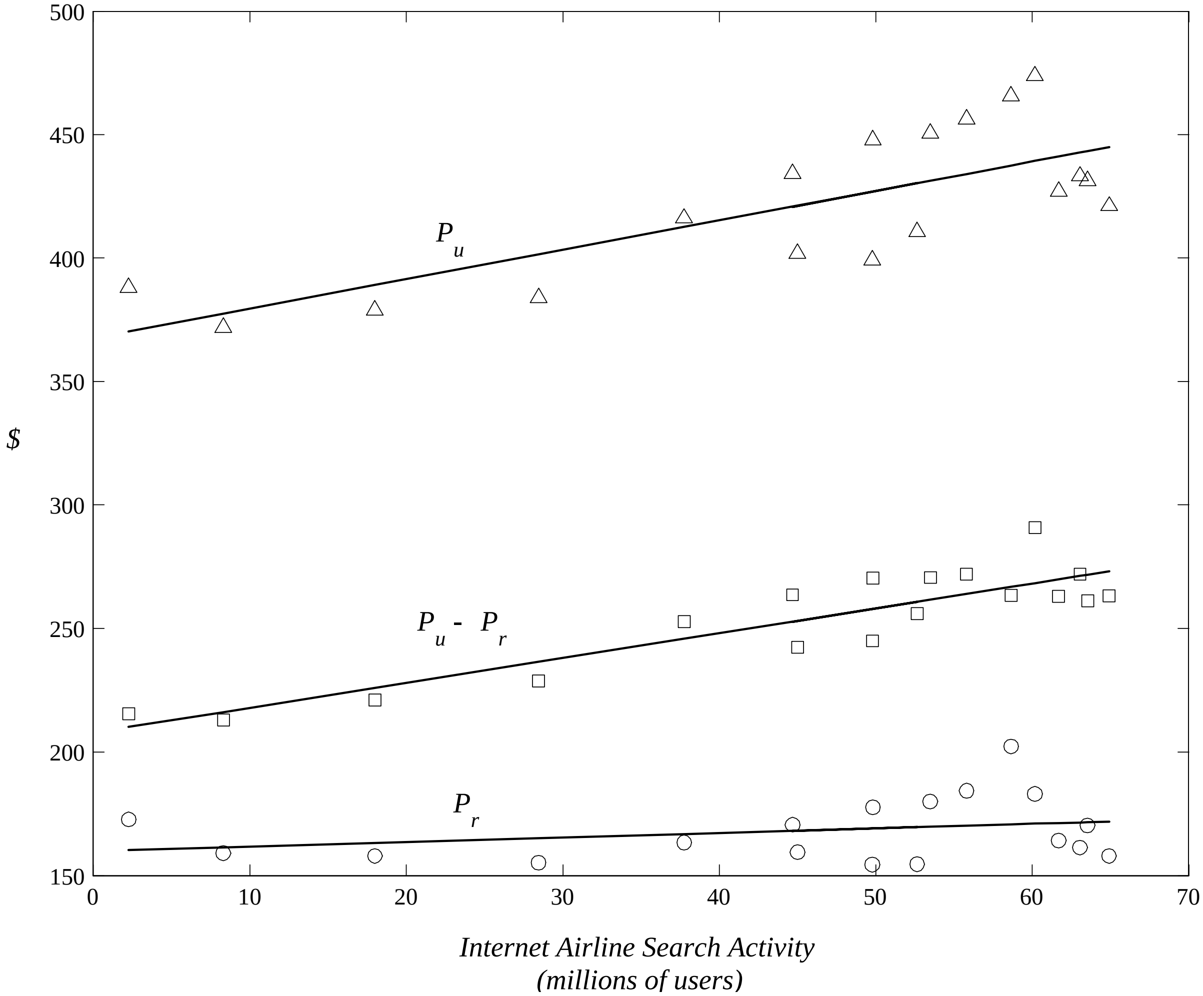

Figure 2: Posterior distribution of the expected change in the ratio of fares as the number of competitors increases from 2 to 3, for three levels of actual Internet airline travel research population sizes: Q2/1998 - 8.3 million; Q2/2000 - 49.8 million; Q2/2002 - 64.9 million.