

Price Discrimination and Business-Cycle Risk*

Marco Cornia,[†] Kristopher S. Gerardi,[‡] Adam Hale Shapiro[§]

February 2, 2011

Abstract

A parsimonious theoretical model of second degree price discrimination suggests that the business cycle will affect the degree to which firms are able to price discriminate between different consumer types. We analyze price dispersion in the airline industry to assess how price discrimination can expose airlines to aggregate-demand fluctuations. Performing a panel analysis on 17 years of data covering two business cycles, we find that price dispersion is highly pro-cyclical. Estimates show that a rise in the output gap of one percentage point is associated with a 1.9 percent increase in the interquartile range of the price distribution in a market. These results suggest that markups move pro-cyclically in the airline industry, such that during booms in the cycle, firms can significantly raise the markup charged to those with a high willingness-to-pay. The analysis suggests that this impact on firms' ability to price discriminate results in additional profit risk, over and above the risk that comes from variations in cost.

*We thank Jim Adams, Ana Aizcorbe, Ben Bridgman, Abe Dunn, Matt Osborne, and Kyle Hood for helpful comments and suggestions. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Bureau of Economic Analysis, the U.S. Department of Commerce, the Federal Reserve Bank of Atlanta, or the Federal Reserve System. The research presented here was primarily conducted while at the Federal Reserve Bank of Boston.

[†]Graduate Student, Johns Hopkins University

[‡]Federal Reserve Bank of Atlanta

[§]Corresponding Author: Bureau of Economic Analysis; e-mail: adam.shapiro@bea.gov

1 Introduction

Firms can expose themselves to fluctuations in demand by choosing specific pricing strategies. In this study, our aim is to better understand how price discriminatory tactics can make airlines vulnerable to aggregate demand risk. Given the large number of bankruptcies in this industry, there is reason to believe that airlines are particularly sensitive to movements in the business cycle. However, while many of the traditional, legacy airlines, including United Airlines, U.S. Airways, and Delta Airlines, have been forced to file bankruptcy, a select group, known as low-cost carriers (LCCs) have been able to stay profitable. The success of these LCCs has been impressive considering the extreme profit volatility that most airlines have had to endure during the post-regulation era.

While differences in cost likely play an important role in the variation in profits across airlines, we focus on a different potential explanation for the volatile nature of airline profits. We argue that price discrimination is a pricing strategy that accentuates profit variation over the business cycle. Price discrimination is a risky tactic in the sense that it exposes the firm to aggregate-demand movements and therefore induces greater revenue fluctuations. We find that relative to LCCs, legacy carriers price discriminate to a large degree, inducing high volatility in revenues. As legacy airlines generally have higher cost levels than LCCs, our results imply that legacy carriers' large revenue variation has been a major contributor to the large number of bankruptcies.

Price discrimination is a well-known strategy that airlines use to try to increase profits. It involves charging higher prices to those consumers with a lower price elasticity of demand, or equivalently, with a higher willingness-to-pay for an airline ticket. In this manner, airlines are able to increase the average markup of prices to marginal cost, and thus, increase their profits. In order to price discriminate, however, airlines must be able to identify and separate consumers with different willingness-to-pay, which they do by offering tickets with various types of restrictions so that consumers separate themselves through the ticket choices that they make (this type of self-selection is referred to as second-degree price discrimination). In this paper we present empirical evidence that suggests price discrimination in the airline industry is highly pro-cyclical. During booms it is easier for airlines to price discriminate between consumers, while in business cycle

troughs it is much more difficult. This results in pro-cyclical average markups, and hence pro-cyclical airline profits. Furthermore, it suggests that airlines which rely more heavily on price discrimination strategies should expect more volatile profits over the business cycle.

To help understand this aspect of airline pricing behavior, we develop a simple theoretical model of second-degree price discrimination based on Mussa and Rosen (1978). The model shows that as long as consumers have diminishing marginal utility of income, profit-maximizing behavior on the part of firms will result in the degree of price discrimination covarying positively with aggregate income, and thus the business cycle. Specifically, during booms in the cycle, the difference between incentive compatible prices charged to high- and low-income consumers is expected to increase, while during troughs in the cycle the difference is expected to decrease.

One significant obstacle to overcome in the empirical analysis is the fact that price discrimination is difficult to identify and measure. We do not have access to data with enough information about consumer and ticket characteristics to identify precise instances of price discrimination and to directly study how price discrimination evolves over time. Instead, we adopt the strategy of the previous literature and use price dispersion at the airline-route level to proxy for price discrimination. In particular, we look at how various measures of price dispersion are correlated with the business cycle, while controlling for variation in price dispersion that is likely due to other factors, such as competition and cost.

Using a relatively long panel dataset on ticket prices, which spans almost two full business cycles, we find evidence that price dispersion is highly procyclical, corroborating the mechanism delineated in the simple pricing model. We interpret this finding as evidence that during booms in the business cycle airlines are able to exploit larger differences in willingness-to-pay between different types of consumers (i.e. business and leisure travelers), which allows them to raise the markup charged to more price inelastic consumers. This story is consistent with the highly procyclical nature of profits in the airline industry as seen in Figure 1.

In addition, this finding suggests that the business cycle induces significant fluctuations in airline profits, over and above variations in cost. As markups are tied to the

degree to which an airline can successfully price discriminate, recessions significantly hinder this profitable pricing mechanism, while booms accentuate it. This story is consistent with the observation that the airlines with the highest within-route variance in price dispersion are the same airlines that have the highest variability in profits as seen in Figure 2. Thus, our analysis suggests that those airlines that rely more on price discrimination techniques will face greater profit risk stemming from the business cycle. It is apparent from Figure 2 that the four largest LCCs (Southwest, JetBlue, Frontier, and AirTran) are characterized by low price dispersion as well as low profit volatility in our sample. This suggests that firms relying more on price discrimination techniques have had more variable profits.

This study is related to numerous microeconomic studies on pricing strategies and business cycle conditions. For instance, Rotemberg and Saloner (1986) theorize that, during booms firms may be less likely to collude since the benefits of cheating are higher, causing firms to cut prices. Another set of theories, based on switching costs and brand loyalty, show that during booms new customers may enter the market causing demand to become more elastic and firms to lower prices (see Bilal (1989), Klemperer (1995), and Stiglitz (1984)). A third theory, put forth by Greenwald, Stiglitz, and Weiss (1984) and analyzed by Chevalier and Scharfstein (1996), shows that during recessions cash-strapped firms may forego offering low prices to attract new customers in order to generate a higher cash flow.

This study also provides potentially interesting insights to macroeconomists who have had a longstanding interest in the cyclical nature of markups. In macroeconomic models of perfect competition, fluctuations in aggregate demand do not shift the labor demand schedule. Thus, in order to generate movements in employment, the models rely on outward shifts in the labor supply curve (see Hall 1980). To generate labor demand shifts, models of imperfect competition (for example, Rotemberg and Woodford (1992), Goodfriend and King (1997)) were developed which introduced an endogenous variable, the markup of price over marginal cost.¹ In particular, a countercyclical markup will act in a similar manner to a positive productivity shock during booms in the business cycle,

¹Including labor hoarding in the framework as in Burnside, Eichenbaum, and Rebelo (1993) has also been a way to explain procyclical labor productivity.

shifting the labor demand curve out and raising real wages.² Our estimates provide some evidence against this hypothesis, at least as an exhaustive possibility.

Although our empirical analysis is confined to one industry, the model indicates that the procyclicality of markups is likely to occur in many other industries as well. Specifically, as diminishing marginal utility of income is inclusive to a broad class of utility specifications, procyclical markup variation should occur in industries in which firms have market power and can successfully price discriminate. Researchers have found many other industries in which firms have been successful at price discrimination such as hotels, stadiums, restaurants, theaters (Leslie 2004), yellow-page advertising (Busse and Rysman 2005), cement (Miller and Osborne 2010) and personal computers (Aizcorbe and Shapiro 2010).

The paper is structured as follows: Section 2 reviews the theoretical framework by which price discrimination can lead to procyclical markup variation. Section 3 discusses possible effects of cost on price dispersion. Section 4 contains a detailed discussion of the data. In Section 5 we perform a fixed-effects, panel estimation of the relationship between price dispersion and various proxies for the business cycle. Finally, Section 6 contains concluding remarks.

2 Theoretical Framework

Price dispersion at the flight level is a common characteristic of the airline industry, and one of the leading explanations for such dispersion is the practice of price discriminating among different types of consumers. It is well known that airlines implement second- and third-degree price discrimination in which they segment heterogeneous groups of consumers and charge them distinct prices for essentially the same product. Advance-purchase requirements, non-refundable tickets, and Saturday-night layovers are a few examples of restrictions that airlines use to identify passengers with different price elasticities of demand. Since high-income or business consumers tend to place a high value on their time, they are more likely to purchase more expensive tickets without such re-

²Other research in this area includes Gali (1994), Edmunds and Veldkamp (2006), and Jaimovich (2007).

restrictions. By making use of these techniques, airlines are able to separate price-sensitive travelers from price-insensitive travelers.

According to a 1993 Gallup Survey, 49% of air travel in the U.S. was for business, and those respondents cited a rise in new business activity or improved financial conditions in their respective firms as the primary reason for increased business travel (Busse, 2002).³ This suggests that when aggregate demand conditions are good, high-income or business consumers have higher demand for air travel, and are also more likely to pay higher prices for tickets characterized by less restrictions and preferred flight times. Thus, during periods of high aggregate demand, carriers will likely experience both a higher proportion of price inelastic travelers demanding tickets, and an increase in the willingness-to-pay for those price inelastic consumers. If the fraction of tickets sold to price inelastic consumers increases and the difference in ticket prices between price inelastic consumers and price elastic travelers increases in periods of high aggregate demand, then both the average markup and price dispersion will follow the business cycle.

We illustrate this relationship between price discrimination and the business cycle with a simple model of monopoly pricing. With the model we derive the precise condition on preferences that ensures a positive correlation between price discrimination and the business cycle. While an oligopolistic model would prove more realistic, it also complicates the analysis and moves away from the scope of the empirical study—the effect of aggregate conditions on airline pricing patterns. Furthermore, in the empirical analysis we control for variations in competition over time, and thus any effect of the business cycle on prices, in a dynamic sense, will be measured in a setting where the degree of competition is held fixed. One could extend our theoretical framework to an oligopolistic game by adding a horizontal dimension on the consumer’s valuation of the product (e.g. brand) as in spatial models such as Schmalensee (1978), Salop (1979) and Brito (2003). Such brand differentiation would allow firms to compete in price and also have some degree of market power. One should therefore think of the theoretical model below as a market in which brand preference is strong enough to allow firms to price above marginal cost, and that consumers’ preferences over brand are distributed

³On the flip side, respondents who indicated less business travel attributed it to decreasing business activity or bad financial conditions in their respective firms.

independently of income.

2.1 Consumers

We consider a model with multiple types of consumers, where consumers differ in their level of income, y_i . A type i consumer with income, y_i , solves the following maximization problem:

$$\max_{d \in \{0,1\}} d \cdot x_v + u(m) \quad (1)$$

subject to:

$$y_i = m + d \cdot p$$

where x_v is valuation of the ticket where $v \in 1, 2$ indicates its attributes. For instance, $v = 2$ indicates a ticket that has an advance purchase requirement or saturday-night stayover requirement, while $v = 1$ indicates a less restrictive ticket. It follows that with positive time costs, the net quality of $v = 1$ will be higher than that of $v = 2$ such that $x_1 > x_2$. The variable m represents the numeraire commodity and d represents the consumer's decision to buy or not buy the good. Note that $u(\cdot)$ is the functional form representing the manner in which the consumer values the numeraire commodity relative to the discrete good. It follows that the indirect utility function for the case in which the consumer purchases the discrete good ($d = 1$) is given by:

$$U = x_v + u(y_i - p). \quad (2)$$

As in Tirole (1988), we make the assumption that a consumer's income is very large relative to the valuation, x_v , and subsequently to the equilibrium price charged. This allows us to take a Taylor expansion around $p^* = 0$ which, under the assumption that $y_i - p \approx y_i$, yields:

$$U = x_v + u(y_i) - u'(y_i)p. \quad (3)$$

It follows that for a given consumer to be better off consuming the good, it must be the case that $x_v + u(y_i) - u'(y_i)p \geq u(y_i)$, which means demand for the good is:

$$d_i(p_v) = \begin{cases} 1 & \text{if } p_v \leq \frac{x_v}{u'(y_i)} \\ 0 & \text{if } p_v > \frac{x_v}{u'(y_i)} \end{cases} \quad (4)$$

2.2 Firm behavior

To simplify the firm's problem we assume two types of consumers: a high income consumer with income y_h and a low income consumer with income y_l . The firm's problem in the two-consumer-type case is to maximize profits given consumer demand derived above. The firm has the option to separate the market by offering different types of tickets. To obtain a separating equilibrium, the firm must be able to separate the market and also find it profit-maximizing to do so. It follows from Mussa and Rosen (1978) that prices must satisfy:

$$p_1^* = b_h x_1 - (b_h - b_l)x_2 \quad (5)$$

$$p_2^* = b_l x_2. \quad (6)$$

where $b_h = \frac{1}{u'(y_h)}$ and $b_l = \frac{1}{u'(y_l)}$. As the high-income consumer values x_2 more than the low-income consumer, the firm must lower the price of x_1 to dissuade the high-income consumer from deviating and purchasing x_2 . Specifically, the bound is lowered by the extra utility the high-income consumer would have received over the low-income consumer by consuming x_2 , $(b_h - b_l)x_2$. This lowering of the price ensures that the high-income consumer does not purchase x_2 instead of x_1 (this ensures that the equilibrium is incentive compatible). In either case, the firm will find it profit-maximizing to separate the market as long as the number of low-income consumers is sufficiently large (see Tirole (1988)).

Note that the price range between the high and low price ticket (a simple measure of price dispersion) will be:

$$D = p_1 - p_2 = b_h(x_1 - x_2). \quad (7)$$

An interesting result is that the level of dispersion does not depend on the low-income consumer's marginal utility of income; only the high-income consumer. This outcome is

attributable to the incentive compatibility constraint. In particular, if the firm is able to raise the price charged to a low-income consumer (without exceeding the low-income consumer's reservation price), then the firm can safely transfer this price rise on to the high-income consumer (without providing an incentive to the high-income consumer to choose the low-quality ticket). It follows that the elasticity of price dispersion relative to a change in aggregate income is:

$$\varepsilon_{D,y} = \frac{\partial D}{\partial y} \frac{y}{D} = \left[-\frac{u''(y_h)}{u'(y_h)} y_h \right] \quad (8)$$

which is simply the coefficient of relative risk aversion (CRA). As long as the CRA is positive (i.e. diminishing marginal utility of income), price dispersion will widen with an increase in aggregate income. Thus, under very general conditions on the consumer's utility of income, price dispersion will follow aggregate income.⁴

To show the linkage between price dispersion and markups we first denote the average markup as:

$$\bar{\mu} = \alpha_h * \mu_1 + (1 - \alpha_h) * \mu_2 \quad (9)$$

where α_h is the proportion of high-income consumers purchasing tickets and $\mu_v = p_v - c_v$ is the markup charged on ticket of type v with marginal cost c_v . Assuming that marginal costs do not covary with aggregate demand, it follows from (8) and (9) that the ratio of the markups charged to each type of consumer, $\frac{\mu_1}{\mu_2}$, will rise with aggregate demand, and so will the average markup.

Overall, the model indicates that, *ceteris paribus*, procyclical price dispersion is indicative of a procyclical average markup. It is this relationship between price dispersion and aggregate demand that we will explore in the airline ticket price data. In particular, we will study the cyclicity of markups in the airline industry using price dispersion as a proxy for average markups. If the model assumptions are reasonable, then we expect to find a positive correlation between price dispersion and aggregate demand. This linkage

⁴It is also relevant to note from (7) that dispersion will also be procyclical if $x_1 - x_2$ follows aggregate income. Thus, price dispersion will be cyclical if time costs are also cyclical. This implies that the effect of aggregate income on price dispersion will be pronounced if nonrestrictive tickets become more valuable during booms in the cycle.

of price dispersion and markup variation, however, is dependent on the assumption that marginal cost does not vary with the business cycle. We address this issue in the following section.

3 Capacity and Marginal Cost

There are two issues we need to address before proceeding. First, if the carrier is constrained by capacity, then as more flights reach full capacity, the expense of an additional passenger becomes very large as either a bigger aircraft or an extra flight is needed to supply the extra seat-mile. Eden (1990) shows that effect can induce price dispersion to rise in periods of peak demand when full capacity is reached. Second, wages of crew and maintenance workers are apt to rise during booms in the cycle, which would cause marginal costs to rise. Thus, even if price dispersion is procyclical, markups may not necessarily be. We address these two possibilities in this section.

3.1 Effective Capacity Cost

In discussing the effect of capacity constraints on pricing, it is useful to decompose marginal cost into its two primary components, which we refer to as the passenger cost and the capacity cost. If the aircraft is not operating at full capacity, then marginal cost is simply equal to the passenger cost; the cost of adding an additional passenger to the airplane. This cost is mostly made up of the extra fuel required to transport the additional weight of the passenger, while other, lesser components include the in-flight costs of serving the additional passenger (i.e. meals, snacks, etc.). However, if the airplane is operating at full capacity, then marginal cost is equal to the direct cost of an additional passenger as well as the more substantial cost of an additional flight. This cost is incurred regardless of whether or not seats on the airplane are filled with passengers, while the passenger cost is only incurred on seats that are sold. This implies that marginal cost at the route level is given by,

$$c_{ij} = \begin{cases} \beta_{ij}, & \text{if capacity is not reached} \\ \beta_{ij} + \lambda_{ij}, & \text{if capacity is reached} \end{cases}$$

where β_{ij} is the passenger cost of serving an additional passenger one mile on route j by carrier i , and λ_{ij} is the cost of an additional flight (in seat-miles).

If airlines price according to stochastic demand concerns, then aggregate demand could alter the firm's expected probability of selling a ticket, and subsequently alter the "effective" capacity cost. In particular, if ex-ante the carrier is uncertain about the level demand for a flight, then under price-setting commitments and costly capacity, profit-maximizing behavior induces a distribution of prices rather than a single price. The intuition is that if the firm were allowed to change price after the realization of the state, then it would set a low price in the low-demand state and a high price in the high-demand state. However, because the firm must commit to a menu of prices ex-ante, its profit maximizing strategy is to assign multiple prices to specified quantities of the good.⁵ That is, if a firm must pay costs irrespective of whether or not its output is sold, then it has a large incentive to set higher prices on goods that are less likely to be sold.⁶

Eden (1990) formalized a model in a setting of perfect competition where there is uncertainty regarding the number of agents who will show up to exchange goods in the marketplace. In such a setting, goods are characterized by the probability that they will be sold, and in equilibrium firms face a tradeoff between price and the probability of sale. In the model, equilibrium prices are given by the condition,

⁵It should be pointed out that more a subtle factor is also needed to induce such pricing behavior by firms. The firm's ex-post optimal price must be positively correlated with the level of demand (Dana 1999, 2001). That is, if the firm's optimal price in the high-demand state is larger than it's optimal price in the low-demand state, then it should use multiple prices to achieve some of the benefits that it would have if it were able to adjust price in response to demand. If instead, the ex-post optimal price is constant or decreasing in demand, the firm's optimal strategy will be to set one price.

⁶Prescott (1975), was the first study to address this issue in the economics literature. His paper was more focussed on the theory of unemployment, and his model was more of an example rather than a formal model of pricing.

$$p_s = \beta + \underbrace{\frac{\lambda}{\text{prob}(\text{sale})_s}}_{\lambda_s^{\text{eff}}} \quad (10)$$

where p_s is the price of the sth good, β is an operating cost that the firm must pay for each good that it sells, λ is the unit capacity cost, and $\text{prob}(\text{sale})$ is the probability that good s is sold. The second term on the right-hand side of the equation can be interpreted as an “effective” capacity cost of good s , λ_s^{eff} . In competitive equilibrium, it is the case that firms are indifferent between selling a high-priced good with low probability and selling a low-priced good with high probability. Dana (1999) extended Eden’s model to monopoly and oligopoly market structures. In this setting, the monopolist sets a higher price for a good that sells only in high demand states since its effective cost is higher.⁷

Under this setting, where the carrier commits to prices ex-ante, the highest priced tickets—tickets with the highest effective capacity cost—are not purchased until demand rises sufficiently high to purchase all of the low priced tickets. Thus, if the carrier is pricing solely with stochastic demand concerns, then peaks in aggregate demand will induce higher price dispersion through the higher effective capacity cost of the remaining seats on crowded aircrafts. In this case, it could be argued that procyclical price dispersion is not evidence of a procyclical markup, it is evidence that the effective capacity cost rises during periods of peak demand due to low available capacity.

While we believe that carriers do take effective capacity cost into account, we discount its prominence in affecting price dispersion at the business cycle frequency. If stochastic demand pricing was causing price dispersion to follow the business cycle, then aircraft capacity utilization should covary positively with the business cycle. Figure 5 shows that the mean aircraft capacity utilization rate has been steadily increasing over the course of the sample period, and is not necessarily positively correlated with the business cycle.⁸

⁷Indirect evidence found by Stavins (1996) is consistent with stochastic demand pricing. Stavins found that even controlling for advance purchase requirements, ticket prices increase as the number of days to departure decreases.

⁸As a robustness check, we estimate an instrumental-variable fixed-effects regression, similar in fashion to the estimator discussed in the next section. We estimated two simple regressions. First we estimated the univariate regression: $\ln(\text{util})_{ijt} = \alpha + \beta_1 YGAP_t + \gamma_{ij} + \varepsilon_{ijt}$, where util_{ijt} is the utilization rate (fraction of occupied seats) for carrier i on route j in quarter t . Second we estimated a

Rather, there is significant seasonal variation in this variable, indicating that effective capacity concerns are likely important at high frequency levels, but not necessarily at middle and low frequency levels. Overall, this figure makes it apparent that, at the business cycle frequency, carriers are more likely altering capacity as opposed to letting aircraft utilization vary. This is consistent with the findings by Puller, Sengupta and Wiggins (2010).⁹

3.2 Marginal Cost

Marginal cost may vary over the business cycle for many reasons. For instance, wages of pilots and flight attendants may rise during booms, as may the price of fuel. For this reason, in our empirical analysis below we proxy for variations in marginal cost using a measure of the carrier’s average variable cost. Numerous studies, such as Caves, Christensen, and Tretheway (1984) and Gillen, Oum, and Tretheway (1985, 1990), have found that the carriers’ passenger output displays constant-returns-to-scale in firm size. This finding suggests that average variable cost may be a valid approximation to marginal cost in this context. The BTS’s P-52 database defines a measure called the “total aircraft operating cost,” which includes fuel, crew wages, maintenance, aircraft leasing, and depreciation.¹⁰ This variable measures the tangible cost to the carrier of operating in a given quarter.

Figure 4 plots this variable as a proportion of total seat-miles for four carriers in our sample. The figure shows that cost per seat-mile is correlated across firms, and has generally increased through the course of the sample period. As expected, Southwest and JetBlue, the two largest LCCs in our sample, have lower cost levels than do the two legacy

regression controlling for competition: $\ln(util)_{ijt} = \alpha + \beta_1 YGAP_t + \beta_2 \ln \widehat{HHI}_{ij} + \gamma_{ij} + \varepsilon_{ijt}$, where HHI_{jt} is the Herfindahl Index associated with route j in quarter t (measured using passenger shares). The coefficient β_1 is estimated to be -0.44 with a standard error of 0.04 in the first specification and -0.43 with standard error 0.03 in the second. This negative coefficient indicates that utilization moves countercyclically.

⁹For robustness purposes, in our empirical analysis that follows we include aircraft utilization as a control for effective capacity cost.

¹⁰As described in the next section, in the fixed-effects specification we decompose this variable between its fuel component and its other components because fuel is relatively much more volatile.

carriers, US Airways and United. This differentiation in cost between legacies and LCCs is ubiquitous across the entire airline industry. The large rise and fall in costs in 2008 can be attributed to the spike in oil prices that occurred during that summer. Overall, including a proxy for marginal cost in the empirical specification removes any variation in price dispersion induced by tangible cost variations. This leaves the estimation of the effect of the business cycle on prices in line with that of the simple model where marginal cost is assumed to be constant.

4 Data

Our study focuses on domestic, direct, coach-class airline tickets over the period 1993q1 to 2009q4. Our sample includes nine major domestic airlines, often referred to as “legacy carriers,”¹¹ as well as a number of low-cost carriers¹² (LCCs) and regional carriers. Ticket prices represent 10-percent of all domestic tickets issued by airlines and are obtained from the DB1B database. In addition to ticket prices, the DB1B includes other quarterly itinerary information, such as origin and destination airports, passenger quantities, number of stops (plane changes), and fare class.¹³ Tickets less than \$20 are believed to be frequent-flyer tickets and are eliminated.

The data is a panel, where an observation is a flight conducted by a specific airline, between an origin and destination airport (route), in a specific time period (year and quarter). For example, an American Airlines direct, coach-class ticket, from Dallas (DFW) to San Francisco (SFO) in the first quarter of 1999 is considered an observation in our data. The direct ticket data include both one-way flights and round-trip flights. The DB1B contains numerous itineraries and fares for the same flight by the same carrier,

¹¹The legacy carriers in our sample include United, US Airways, Delta, American, Alaskan, TWA, Continental, Northwest, and America West.

¹²The list of LCCs, obtained from Ito and Lee (2003), includes Air South, Access Air, AirTran, American Trans Air, Eastwind, Frontier, JetBlue, Kiwi, Morris Air, National, Pro Air, Reno, Southwest, Spirit, Sun Country, ValuJet, Vanguard, and Western Pacific. For a more detailed discussion of LCCs see Goolsbee and Syverson (2008).

¹³There are three different sub-components to the DB1B data set. They are market data, coupon data, and ticket data; and we combine variables from all three. For further reference, see the BTS’s website <http://www.transtats.bts.gov>.

reflecting the quarterly frequency of the data, as well as the many different fares found within the same fare class, on the same flight, at a given point in time. Thus, the data comprise distributions of prices for carrier-route itineraries.¹⁴ Price dispersion is measured using three separate proxies: the interquartile range, the Gini coefficient, and the 90th and 10th price percentiles estimated separately. The interquartile range and Gini coefficient are advantageous in that they summarize dispersion with one statistic, while the price percentiles have the advantage that they provide more detailed information about the tails of the distribution.

The median Gini coefficient in our entire sample is 0.225, and is 0.217 for legacy carriers and 0.171 for LCCs. To get a better handle on the dynamics of price dispersion, we present a few graphical examples of the pricing patterns seen in the data. Figure 3 plots price percentiles of three routes along with a plot of the output gap.¹⁵ The top two panels correspond to routes operated by two legacy carriers, American Airlines and Delta Airlines, while the bottom panel consists of a route operated by Southwest Airlines. It is noteworthy that in the legacy carrier panels, the higher price percentiles seem to closely follow the output gap. The top portion of the price distribution rises and falls with the boom in the late 1990s and then begins to gradually fall as aggregate demand deteriorates. In contrast, we do not see the same relationship in the Southwest panel, as none of the percentiles of the price distribution seem correlated with the output gap. In the next section we conduct a more systematic analysis using panel data methods in an effort to confirm these observations.

5 Estimation

We utilize the same empirical strategy as in Gerardi and Shapiro (2009) to study the relationship between price dispersion and the business cycle. Since the data is a panel of airline-route observations, it is possible to assess the effects of business cycle variations on price dispersion while holding fixed time-invariant, route-specific factors,

¹⁴See Appendix B for more details on the construction of the dataset, and Gerardi and Shapiro (2009) for an even more detailed description.

¹⁵The output gap is defined as the log difference between the actual nominal GDP and the CBO's measure of potential output.

as well as any route-specific variation in the degree of competition and carrier-specific variation in fuel and other operating costs. We use a fixed-effects panel estimator, which exploits the times-series variation along a specific route in the estimation routine. We use two different approaches in measuring the effect of business cycle variations on price dispersion.

The model reviewed in Section 2 predicts that high-income consumers' price elasticity of demand will be more sensitive to business cycle variations than that of low-income consumers. Specifically, equation (8) predicts that a rise in aggregate income will induce a rise in price dispersion. This relationship can be formulated by the following empirical specification:

$$DISP_{ijt} = \theta_0 + \beta_1 * YGAP_t + \beta_3 * \ln \widehat{HHI}_{jt} + \beta_4 * \ln UTIL_{ijt} + \gamma' * X_{it} + \gamma_{ij} + \varepsilon_{ijt}. \quad (11)$$

where i indexes the carrier, j the route, and t the time period. In this specification, the output gap, $YGAP_t$ is used to proxy for the business cycle, as measured by the Congressional Budget Office (CBO). Carrier-route fixed effects are represented as γ_{ij} , and \widehat{HHI}_{jt} is the market concentration of the route as measured by the Herfindahl index and is instrumented using the same variables as in Borenstein and Rose (1994) and Gerardi and Shapiro (2009).¹⁶ We control for costs on a specific carrier i with the vector X_{it} , which includes the logarithm of the carrier's average fuel cost per gallon, $\ln FUEL_{it}$, as well as the remaining operating cost per seat-mile, $\ln COST_{it}$, measured by the BTS for a specific carrier. To control for effective capacity cost variations, we include the logarithm of the utilization rate of carrier i on route j in period t .¹⁷

The second specification takes the form:

$$DISP_{ijt} = \theta_0 + \beta_1 * UR_{jt} + \beta_2 * \ln \widehat{HHI}_{jt} + \beta_4 * \ln UTIL_{ijt} + \gamma' * X_{it} + \gamma_{ij} + \varepsilon_{ijt}. \quad (12)$$

¹⁶See the appendix for a description of the instruments .

¹⁷It is important to note that since this variable is endogenous, the estimate on β_4 will be biased downwards. For this reason, we do not report the estimate of β_4 . No results changed when we did not include the utilization rate in the regression.

where the average unemployment rate of the two endpoint states on the route, UR_{jt} , is used as an alternative proxy for the business cycle. In both specifications, price dispersion, $DISP_{ijt}$, is measured in three different ways: the logarithm of the interquartile range, the Gini log-odds ratio,¹⁸ and the 90th and 10th percentiles, each estimated in separate regressions. Analyzing the top and bottom of the price distribution separately provides additional information regarding the source of the change in price dispersion. Observations are weighted by the total number of passengers on the route over the entire sample period and standard errors are clustered by route in order to control for autocorrelation as well as correlation between carriers on the same route.

6 Estimation Results

Table 1 contains estimation results for both specifications, using the logarithm of the interquartile range and the Gini log-odds ratio as the dependent variable. We report results for all direct routes in our 17-year sample.¹⁹ The effect of a rise in the business cycle—as measured by the output gap—on price dispersion is positive and significant at the 1-percent significance level. The estimate indicates that a one percentage point rise in the output gap (i.e. from 0.01 to 0.02) is associated with an increase in the interquartile range by 1.88 percent and the Gini log odds by 0.012. The results from the second specification are similar to the first, indicating that a decrease in the unemployment rate induces an increase in the amount of price dispersion on a given route.²⁰ A one percentage point fall in the unemployment rate is associated with a 2.5 percent increase in the interquartile range. These estimated coefficients are suggestive of the existence of

¹⁸We measure price dispersion using the Gini log-odds ratio given by given by $G_{ij}^{lodd} = \ln \frac{G_{ij}}{1-G_{ij}}$, which produces an unbounded statistic. No results change when the log of the Gini coefficient is used instead. See Hayes and Ross (1998) for further discussion.

¹⁹This sample includes 154,333 carrier-route observations when using $\ln(IQR)$ as the dependent variable and 155,961 carrier-route observations using the Gini log-odds ratio. The reason for the fewer observations in the first specification is that observations in which the interquartile range was equal to zero were necessarily dropped.

²⁰This sample includes 153,706 carrier-route observations when using $\ln(IQR)$ as the dependent variable and 155,331 carrier-route observations using the Gini log-odds ratio. We have fewer observations in this specification because we do not have unemployment information for American Samoa or St. Thomas.

procyclical markups in the airline industry.

A look at the estimates from the percentile regressions in Table 2 sheds further light on the manner in which price dispersion follows the business cycle. The estimates show that an increase in the output gap raises the 90th-percentile price level but has no significant effect on the 10th-percentile price level. An increase in the output gap by one percentage point is associated with a 1.35 percent increase in the 90th percentile price, but is not correlated with the 10th percentile price. Similarly, a fall in the unemployment rate by 1 percentage point is associated with a 1.6 percent increase in the 90th percentile price, while there is statistically significant, but small -0.25 percent negative correlation between the unemployment rate and the 10th percentile price. Thus, the estimates are consistent with the interpretation that during booms in the business cycle, airlines increase the extent to which they price discriminate by raising prices to those consumers with the highest willingness-to-pay for air travel.²¹

As in Gerardi and Shapiro (2009), we find that the effect of a decrease in competition—as measured by an increase in market concentration $\ln \widehat{HHI}$ —on price dispersion is positive and significant at the 1-percent significance level.²² There also appear to be interesting dynamics occurring on the cost side. Fuel costs seem to filter into the 10th percentile prices and 90th percentile prices equally, while other operating costs filter more into the 90th percentile prices. There are many plausible stories that could explain this result. One possibility may be that carriers simply find it easier to pass costs on to the more price inelastic consumers as they are more likely to lose the less elastic consumers to competition.

Overall, the fixed-effects, panel estimates provide evidence of a positive relationship between the business cycle and price dispersion in the airline industry. Furthermore, the results show that variations in the business cycle have a large positive effect on the upper portion of the price distribution. These results are consistent with the mechanism outlined by the model delineated in Section 2, whereby carriers raise the markup charged to those consumers who are least price elastic during booms in the business cycle.

²¹We also divided the sample into two periods based on the National Bureau of Economic Research (NBER) business cycle dates. We found positive and significant coefficients in each sample period, although a larger effect in the first.

²²All instruments were relevant at the 1 percent level as measured by the Cragg-Donald statistic.

6.1 Carrier and Consumer Heterogeneity

As discussed earlier, legacy carriers, sometimes called “hub-and-spoke carriers” or network carriers, tend to implement different pricing strategies compared to the LCCs. For instance, some legacy carriers offer “economy-plus,” which offers passengers more leg room, separate access through security, and/or early boarding. There is also evidence from Figure 2 which suggests that LCCs have lower variability in price dispersion over the cycle. To determine whether these different types of carriers actually price differently over the business cycle, we divide the sample between legacy carriers and LCCs and re-estimate the main econometric specification for each sample separately.

The estimates divided by carrier type are reported in the top panel of Table 3 and show that most of the effects from the business cycle on price dispersion in the full sample of routes stem from the legacy carriers.²³ The effect of the output gap on the interquartile range is roughly two to three times the magnitude in the sample of legacy carriers ($\hat{\beta}_1 = 2.77$ compared to the estimated effect in the sample of LCCs $\hat{\beta}_1 = 1.03$). This is consistent with the interpretation that the pricing considerations of legacy carriers are more sensitive to the movements in willingness-to-pay of more price inelastic consumers.

We further explore the relationship between price dispersion and business cycle conditions using a subsample analysis from Gerardi and Shapiro (2009). In that paper, we attempted to decompose the full sample of routes into those with a heterogeneous consumer base and those with a more homogeneous consumer base. Since routes between large cities tend to attract both business and leisure travelers, they tend to have a bimodal distribution of prices while routes to largely leisure destinations, like islands and famous beaches, tend to have unimodal price distributions and lower median prices.²⁴ Thus, airlines may have more opportunities to implement price discrimination strategies on these big-city routes since they include relatively more high-income, business consumers. Furthermore, note that equation (8) implies that dispersion on big-city routes will be more sensitive to the cycle if the utility over income displays increasing relative

²³For ease of view, we do not report estimates on $HERF_{jt}$ or cost variables. No results significantly changed from the preceding section.

²⁴For a full list of the cities in each sample as well as a detailed description of how these subsamples are created see Gerardi and Shapiro (2009).

risk aversion and less sensitive to the cycle if it displays decreasing relative risk aversion.²⁵

The bottom two panels of Table 3 contain estimates of the correlation between the interquartile range of the price distribution for a carrier-route observation and the business cycle for big-city routes versus leisure routes. The estimates show that price dispersion is more closely tied with the output and the city-wide unemployment rate for the big-city route sample than the leisure sample. For instance, in the big-city, legacy carrier sample, a one percentage point rise in the output gap is associated with a 3.6 percent increase in the interquartile range, while it is associated with a 1.8 percent increase in the leisure route sample. For low-cost carriers, the effect is 1.3 percent in the big-city sample and 0.79 percent in the leisure route sample.

7 Conclusion

Our analysis has shown that price dispersion is significantly procyclical in the airline industry. This result suggests that markups of price over marginal cost are also procyclical, in line with the recent findings of Nekarda and Ramey (2010). The results are consistent with a parsimonious theoretical model of second-degree price discrimination, which shows that as long as firms have the ability to price discriminate and consumers have diminishing marginal utility of income, price dispersion will follow aggregate demand. Thus, we believe our results can be extrapolated to other industries in which price discrimination plays an important role. Overall, the analysis suggests price discrimination is a risky strategy in the sense that it accentuates the degree to which the firm is exposed to fluctuations in aggregate demand. We posit that this may be an additional reason why low-cost carriers, such as Southwest and JetBlue, which do not price discriminate to the same extent as legacy carriers, have experienced more profit stability over the last two decades.

Future work on the issue could involve examining whether other microeconomic mechanisms are at work in determining markup cyclicalities, as well as looking at industries with different types of firm pricing strategies. For instance, Miller and Osborne (2010) find evidence of price discrimination in the cement industry to be based on the spatial

²⁵This can be seen by taking the derivative of 8 with respect to y_h .

location of the consumer. In the healthcare industry, players are involved in a bargaining game over price. Markup cyclicalities are therefore likely distinct for certain classes of industries.

A Variable Definitions

- $\ln P(k)_{ijt}$ - The logarithm of the k th price percentile of carrier i on route j in period t , obtained from the DB1B.
- $\ln IQR_{ijt}$ - The logarithm of the interquartile range, given by $P(75)_{ijt} - P(25)_{ijt}$, where $P(k)_{ijt}$ is the price percentile of carrier i on route j in period t , obtained from the DB1B.
- G_{ijt}^{lodd} - The Gini log-odds ratio, given by $G_{ijt}^{lodd} = \ln(\frac{G_{ij}}{1-G_{ij}})$, where G_{ijt} is the Gini coefficient of carrier i 's price distribution on route j in period t , calculated using data from DB1B.
- $\ln HHI_{jt}$ - The logarithm of the Herfindahl index of route j in period t , calculated using passenger shares obtained from the DB1B.
- $YGAP_t$ - The log of nominal GDP in period t minus the log of nominal potential GDP in period t , as measured by the Congressional Budget Office (CBO).
- UR_{jt} - The average metropolitan unemployment rate in period t of the origin and destination state of route j , obtained from Bureau of Labor Statistics (BLS).
- $\ln FUEL_{it}$ - The average cost per gallon fuel by carrier i in period t , obtained from the BTS P-52 database.
- $\ln COST_{it}$ - Total operating costs minus total fuel costs divided by total seat-miles for carrier i in period t , obtained from the BTS P-52 database.

Instruments

- $\ln PASSRTE_{jt}$ - The logarithm of total enplaned passengers on route j in period t from the T-100 Domestic Segment Databank.
- $IRUTHERF$ - This instrument is identical to one used by Borenstein and Rose (1994). This variable is the square of the fitted value for $MKTSHARE_{ijt}$ from its first-stage regression, plus the rescaled sum of the squares of all other carrier's

shares. See Borenstein and Rose (1994) for a more detailed explanation. It is equal to $\widehat{MKTSHARE}_{ijt}^2 + \frac{HERF_{jt} - \widehat{MKTSHARE}_{ijt}^2}{(1 - \widehat{MKTSHARE}_{ijt})^2} * (1 - \widehat{MKTSHARE}_{ijt})^2$.

- $GENSP - \frac{\sqrt{ENP_{j1} * ENP_{j2}}}{\sum_k \sqrt{ENP_{k1} * ENP_{k2}}}$, where k indexes all airlines, j is the observed airline, and ENP_{k1} and ENP_{k2} are airline k 's average quarterly enplanements at the two endpoint airports. This instrument is similar to one used by Borenstein and Rose (1994), with the difference being that Borenstein and Rose use average daily enplanements, while we use average quarterly enplanements, as a result of data availability. Data on enplanements were obtained from the T-100 Domestic Segment Databank.

B Data Construction

In this appendix, we discuss our methods and assumptions involved in constructing our panel of airline-route ticket observations from the DB1B and T-100 Domestic Segment databases maintained by the BTS at their online website, Transtats. There are three different sub-components to the DB1B data set. They are market data, coupon data, and ticket data, and we combine variables from all three sources.²⁶

We use only domestic, coach-class itineraries and keep only tickets containing direct flights.²⁷ Direct flights typically account for 30 percent of the itineraries in the DB1B over the course of our sample, with no apparent trend.

The BTS includes a variable that describes the reliability of each ticket price (“dollar cred”). The variable takes on a value of 0 if the fare is of questionable magnitude, based on a set of limits defined by the BTS, and it takes a value of 1 if it is credible. We drop all tickets for which this variable takes a value of zero.

The DB1B also provides limited information regarding the fare class of each ticket. Each ticket is labeled as either coach-class, business-class, or first-class, and we eliminated all first-class and business-class itineraries. Unfortunately, the DB1B does not have any

²⁶For further reference, see the BTS’s website <http://www.transtats.bts.gov>.

²⁷The sample of direct flights encompasses both non-stop flights and flights in which there is a stop but no change of plane.

direct way of identifying frequent-flyer tickets, but there are indirect methods that have been used in the previous literature, and we follow these in our analysis. First, we drop all fares coded as 0. Next, we dropped all fares that are less than or equal to \$20 (\$10 for one-way tickets).

In addition to eliminating frequent-flyer tickets and higher-class tickets, we also eliminate tickets in which the operating and ticketing carriers are different due to code sharing arrangements. Code sharing is a practice where a flight operated by an airline is jointly marketed as a flight for one or more other airlines. Due to the uncertainty regarding the actual airline who is setting the price schedule in such an arrangement, we decided to eliminate these itineraries. Code sharing first appears in the data in 1998:Q1. On average, approximately 80 percent of the original number of direct tickets in the DB1B is retained in the analysis.

After filtering the ticket data for each quarter of the DB1B, we combined tickets from all 55 quarters and collapsed the data into airline-route observations. For example, if we had 10,000 United Airline tickets between Boston and Los Angeles in 1993:Q1, we calculated summary statistics (such as the Gini coefficient), and collapsed the data into a single observation corresponding to a United Airlines flight between Boston and Los Angeles in 1993:Q1.

The merge between the DB1B and T-100 Segment databases was not exact (around 45 percent matched). First, since the DB1B does not provide complete coverage for all airlines and routes, there are a number of direct routes in the T-100 data that we do not find in the DB1B (especially low-volume routes). Second, the DB1B does not allow us to distinguish between a non-stop, direct ticket and a ticket that involves a stop without a plane change. For example, if a passenger takes a flight from Boston to Orlando that stops in Atlanta, but does not involve a plane change, his itinerary will look identical to that of a passenger who flies from Boston to Orlando without any stops. For this reason, we identified some airline routes as direct in the DB1B, that are not non-stop, and therefore do not have segment information in the T-100 data. While we lose many airline-route observations during the merge as a result, we believe that this merge actually provides a nice filter, since we would ideally like to use only non-stop, direct flights. Thus, by merging data between the DB1B and the T-100, we likely eliminate a

large proportion of flights that are direct, but not non-stop due to a plane change.

In an effort to eliminate possible coding errors, we drop certain airline-route observations from the data that we believe do not have adequate coverage to calculate reliable price dispersion statistics. We drop any airline-route observation that does not have at least 100 passengers in the DB1B. Furthermore, for each airline route observation, we calculate the average number of passengers over time in both the DB1B and the T-100 Segment databases. If the number of passengers on an airline route in a given quarter falls below 25 percent of its mean over time in one of the databases, but not in the other, then we drop the observation from our data, on the basis that its value is most likely measurement error. However, if the number of passengers on an airline route in a given quarter falls below 25 percent of its mean in both the DB1B and the T-100 Segment databases, then we keep the observation in our data.

Finally, we addressed the issue of “double counting.” Since we defined a route as a directional trip in our data, any round-trip ticket would count twice. For example, a round-trip fare from Boston to San Francisco would appear twice in the data — once as BOS-SFO and once as SFO-BOS. Since this would have no effect on the consistency of our estimates, but a significant effect on the size of our standard errors, we chose to drop one of the directions. Of course, the drawback of this assumption is that some one-way fares were dropped from the data as a result. In our judgment, the first issue outweighed the second issue.

References

- [1] Borenstein, S., Rose, N., 1994. Competition and Price Dispersion in the U.S. Airline Industry. *The Journal of Political Economy* Vol. 102(4), 653-683.
- [2] Bills, M., 1989. "Pricing in a Customer Market." *Quarterly Journal of Economics*, 104(4), pp. 699-718.
- [3] Burnside, C, Eichenbaum, M., and Rebelo, S. 1993. "Labor Hoarding and the Business Cycle." *The Journal of Political Economy* Vol. 101, No. 2 pp. 245-273.
- [4] Busse M., 2002. Firm Financial Condition and Airline Price Wars. *The RAND Journal of Economics* 33(2): 298-318.
- [5] Busse, M., and Rysman, M., 2005. Competition and Price Discrimination in Yellow Pages Advertising. *The RAND Journal of Economics* 36(2): 378-390.
- [6] Chevalier, J. and Scharfstein, D., 1996. "Capital-Market Imperfections and Countercyclical Markups: Theory and Evidence," *American Economic Review*, vol. 86(4), pages 703-25, September.
- [7] Caves, D., L. R. Christensen, and M. W. Tretheway, 1984. "Economies of Density versus Economies of Scale: Why Trunk and Local Service Airline Costs Differ, *Rand Journal of Economics*, 15 (4), Winter, 471-489.
- [8] Dana J., 1999. "Equilibrium Price Dispersion Under Demand Uncertainty: The Roles of Costly Capacity and Market Structure." *The RAND Journal of Economics* Vol. 30(4), 632-660.
- [9] Dana J., 2001. "Monopoly Price Dispersion Under Demand Uncertainty." *International Economic Review* Vol. 42(3), 649-670.
- [10] Eden, B., 1990. "Marginal Cost Pricing When Spot Markets Are Complete." *The Journal of Political Economy* Vol. 98(6), 1293-1306.
- [11] Edmunds, C., and L. Veldkamp. 2006. "Income dispersion and countercyclical markups, *Journal of Monetary Economics* vol. 56(6), pages 791-804.

- [12] Gali, J., 1994. “Monopolistic Competition, Business Cycles, and the Composition of Aggregate Demand,” *Journal of Economic Theory*, 63(1), 7396
- [13] Gerardi, K., and Shapiro, A.H., 2009. “Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry.” *Journal of Political Economy*. vol. 117(1), pages 1-37, 02.
- [14] Gillen, D.W., T.H. Oum, and M.W. Tretheway, 1990. “Airline Cost Structure and Policy Implications: A Multi-Product Approach for Canadian Airlines, *Journal of Transport Economics and Policy*.
- [15] Goodfriend, M. and King, R., 1997. The New Neoclassical Synthesis and the Role of Monetary Policy. In *NBER Macroeconomics Annual*, pp. 23183. MIT Press.
- [16] Goolsbee, A. and Syverson, C. 2008. “How do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines,” *The Quarterly Journal of Economics*, vol. 123(4), pages 1611-1633.
- [17] Greenwald, B., Stiglitz, J. and Weiss, A. 1984. “Informational Imperfections in the Capital Market and Macroeconomic Fluctuations.” *American Economic Review, Papers and Proceedings*, 74(2), pp. 194-99.
- [18] Hall, R., 1980. Employment Fluctuations and Wage Rigidity. Brookings Papers on Economic Activity.
- [19] Hayes, K., Ross, L., 1998. “Is Airline Price Dispersion the Result of Careful Planning or Competitive Forces?” *Review of Industrial Organization* 13: 523–542.
- [20] Jaimovich, N., 2008. “Firm Dynamics, Markup Variations, and the Business Cycle.” *Journal of Monetary Economics* vol. 55(7), pages 1238-1252.
- [21] Miller, N. and Osborne, M., 2010. “Competition among Spatially Differentiated Firms: An Empirical Model with an Application to Cement”, EAG Working Paper 10-2.

- [22] Mussa, M. and Rosen, S. 1978. "Monopoly and product quality", *Journal of Economic Theory* 18, 301-317.
- [23] Nekarda, C. and Ramey, V.,2009. "The Cyclical Behavior of the Price-Cost Markup," UCSD Working paper.
- [24] Prescott, E., 1975. "Efficiency of the Natural Rate." *The Journal of Political Economy* Vol. 83(6), 1229-1236.
- [25] Puller, S., Sengupta, A., and Wiggins, S. 2009. "Testing Theories of Scarcity Pricing in the Airline Industry," NBER Working Papers 15555.
- [26] Rotemberg, J. and Saloner, G. 1986. "A Super-game Theoretic Model of Business Cycles and Price Wars During Booms." *American Economic Review*, 76(3), pp. 390-407.
- [27] Rotemberg, J., J., and Woodford,M., 1991. "Markups and the Business Cycle," in *O.J. Blanchard, and S. Fischer (eds.), NBER Macroeconomics Annual 1991 Cambridge, MIT Press*
- [28] Rotemberg, J. and Woodford,M., 1992. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, 100(6), 1153-1207.
- [29] Shapiro, A.H. and Aizcorbe, A., 2010. "Implications of Consumer Heterogeneity on Price Measures for Technology Goods" *BEA Working Paper*.
- [30] Stavins, J., 1996. Price Discrimination in the Airline Market: The Effect of Market Concentration. *Review of Economics and Statistics* Vol. 83(1).
- [31] Stiglitz, J. 1984. "Price Rigidities and Market Structure." *American Economic Review*, 74(2), pp. 350-55.

Table 1: Full Sample Estimates

	$\ln(IQR)$		G^{lodd}	
YGAP	1.882*** (0.210)		1.219*** (0.104)	
-UR		2.535*** (0.279)		1.714*** (0.152)
$\ln \widehat{HERF}$	0.232*** (0.040)	0.227*** (0.040)	0.069*** (0.023)	0.066*** (0.023)
$\ln FUEL$	0.103*** (0.012)	0.087*** (0.013)	-0.030*** (0.008)	-0.040*** (0.008)
$\ln COST$	0.194*** (0.059)	0.168*** (0.055)	0.465*** (0.029)	0.449*** (0.028)
Observations	154333	153632	155961	155254

Notes: All regressions include carrier-route-specific dummies. Standard errors are in parentheses and are clustered by route to account for both autocorrelation and correlation between carriers on the same route. One, two, and three asterisks indicate significance at the 10-percent, 5-percent, or 1-percent significance level, respectively.

Table 2: Full Sample Estimates: Percentiles

	ln(90)		ln(10)	
YGAP	1.349*** (0.108)		0.033 (0.062)	
- UR	1.590*** (0.148)		-0.246** (0.100)	
ln \widehat{HERF}	0.279*** (0.027)	0.276*** (0.027)	0.211*** (0.014)	0.211*** (0.014)
ln <i>FUEL</i>	0.073*** (0.009)	0.060*** (0.010)	0.112*** (0.005)	0.110*** (0.005)
ln <i>COST</i>	0.306*** (0.032)	0.286*** (0.031)	-0.017 (0.015)	-0.019 (0.014)
Observations	155961	155254	155961	155254

Notes: All regressions include carrier-route-specific dummies. Standard errors are in parentheses and are clustered by route to account for both autocorrelation and correlation between carriers on the same route. One, two, and three asterisks indicate significance at the 10-percent, 5-percent, or 1-percent significance level, respectively.

Table 3: Panel Estimates by Carrier and Route Type

	Legacy Carriers		Low-Cost Carriers	
Full Sample				
YGAP	2.772*** (0.222)		1.030*** (0.173)	
- UR	2.961*** (0.309)		1.867*** (0.268)	
Observations	105562	104920	40941	40926
Big-City Routes				
YGAP	3.616*** (0.394)		1.312*** (0.332)	
- UR	3.684*** (0.550)		2.112*** (0.468)	
Observations	34807	34807	8276	8276
Leisure Routes				
YGAP	1.839*** (0.367)		0.791* (0.408)	
- UR	1.996*** (0.517)		1.361** (0.560)	
Observations	20877	20235	10244	10229

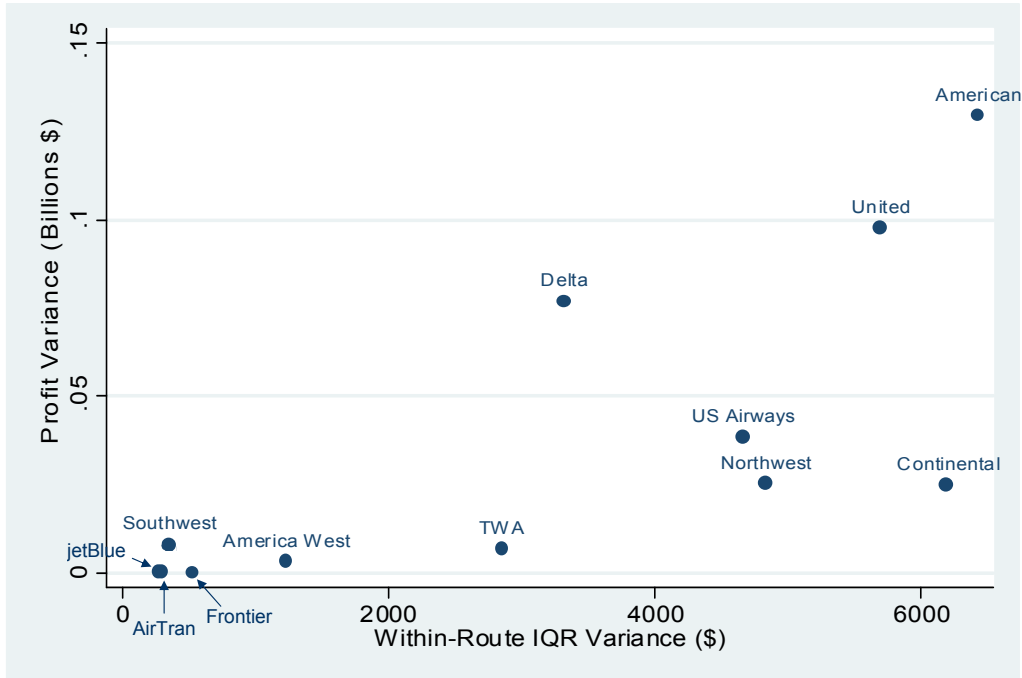
Notes: The dependent variable is the logarithm of the interquartile range. All regressions include carrier-route-specific dummies. Standard errors are in parentheses and are clustered by route to account for both autocorrelation and correlation between carriers on the same route. One, two, and three asterisks indicate significance at the 10-percent, 5-percent, or 1-percent significance level, respectively.

Figure 1: Operating Profits



Notes: Total operating profits are taken from the BTS's P-12 database and includes profits for carriers with annual operating revenues of \$20 million or more. The output gap is the difference between actual and potential output as measured by the Congressional Budget Office (CBO) and the Bureau of Economics Analysis.

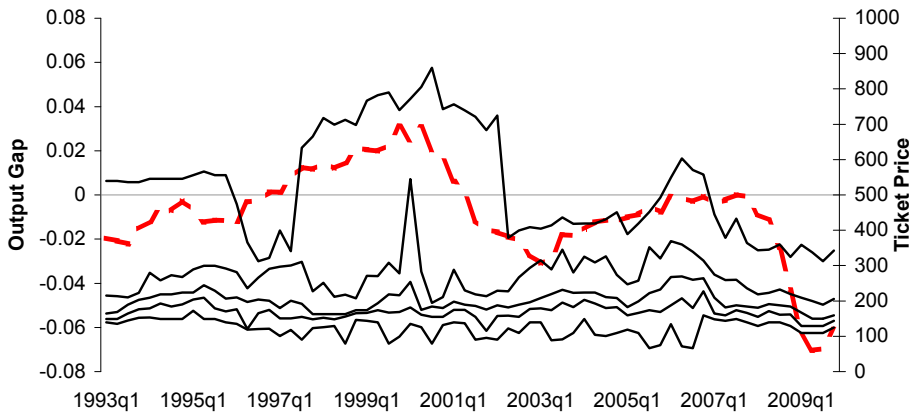
Figure 2: Profit Variance and Price Dispersion Volatility



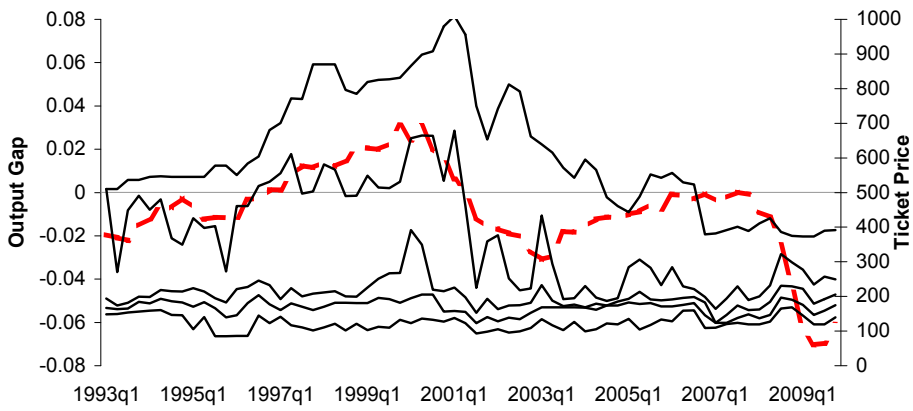
Notes: Depicted are the ten largest carriers in terms of passengers per quarter. We also included two large LCCs, Frontier and AirTran. The within-route interquartile range (IQR) of prices was calculated as the average (weighted by passengers for each carrier), of which we calculated the variance for each carrier. This figure uses domestic operating profits from the BTS's P-12 database.

Figure 3: Pricing Dynamics

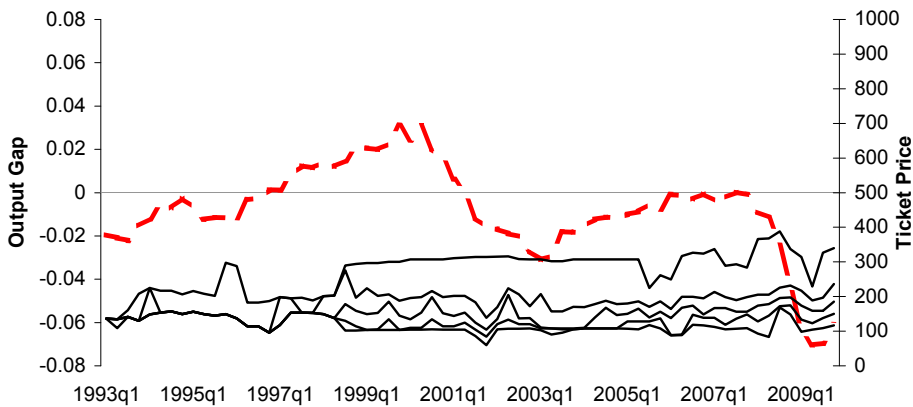
Atlanta (ATL) to Phoenix (PHX) - Delta Airlines



Dallas (DFW) to San Francisco (SFO) - American Airlines



Nashville (BNA) to Phoenix (PHX) - Southwest Airlines



Notes: Depicted are 10th, 30th, 50th, 70th and 90th price percentiles for three airline-route observations. The output gap is also depicted as a dashed line.

Figure 4: Aircraft Operating Costs

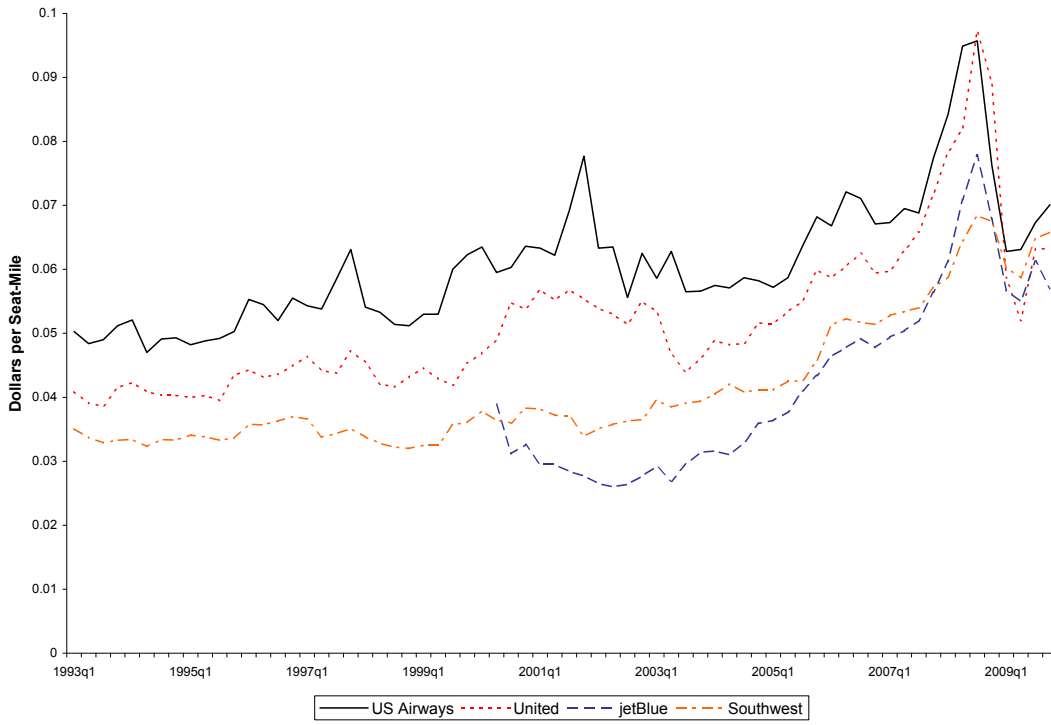


Figure 5: Aircraft Utilization Rate

