

Estimating the Form of Adjustment Costs with Data on Heterogeneous Capital Goods

by

Austan Goolsbee
University of Chicago, Graduate School of Business
American Bar Foundation and NBER

and

David B. Gross
Lexecon
Chicago, Illinois

October, 2000

Abstract

This paper estimates the form of the micro-level marginal costs of adjusting capital using detailed data on investment decisions in the US airline industry. The data include the capital stock, investment, retirement, market values, operating costs, and utilization rates of 16 different types of capital goods for each airline. These data on heterogeneous capital goods allows us to estimate the desired stock of capital for each type of plane while controlling for unobserved changes in airline profitability. The results show significant evidence of non-convex adjustment costs. The adjustment costs seem to combine a region of inaction for capital investment with approximately quadratic adjustment costs conditional on positive or negative investment. The adjustment costs for utilization show similar non-convexities but with smaller adjustment costs. Aggregating to the firm level or using accounting data obscures the non-convexities and biases the estimated adjustment costs upward.

We would like to thank Andrew Abel, Jason Abreveya, Daron Acemoglu, Susanto Basu, Ricardo Caballero, Steve Davis, Daniel Hamermesh, Wei Jiang, Dimitrios Kavvathas, Anil Kayshap, John Leahy, and Matthew Shapiro for helpful comments. Austan Goolsbee would

also like to acknowledge the financial support of the Alfred P. Sloan Foundation, the National Science Foundation and the American Bar Foundation.

I. Introduction

Neo-classical investment models have generally not done well in explaining firm investment, despite numerous attempts and refinements. Recent research has argued that the standard assumption of convex adjustment costs may be one source of the problem and has explored how fixed costs or other non-convex costs of adjustment can alter both firm and aggregate investment dynamics, making investment a non-linear function of shocks.¹ A critical element in the debates over investment theories is empirical knowledge about what form the micro-level costs of adjustment actually take.^{2,3}

One of the basic difficulties with estimating the form of such adjustment costs is the difficulty of finding appropriate data. Many studies acknowledge, but do not solve such standard problems as poorly measuring the true market value of capital, not observing important intangible assets, or failing to capture the forward looking nature of the investment process.

Less frequently discussed is the equally important issue of capital heterogeneity.⁴ Firms own many varieties of capital goods with different purposes and different adjustment costs and these goods are often not very substitutable. Existing studies which focus on the dollar value of net investment are forced to assume unrealistic specifications for adjustment costs. Unless the capital goods are perfect substitutes, the dollar value of net investment will not be a sufficient statistic for the adjustment process. Buying one type of capital and selling another of equal value, for example, even if it leaves market value unchanged, may not result in zero adjustment costs. This fact has driven researchers to study firm or plant level data, but even at the plant level, capital is not homogenous. Once capital heterogeneity is allowed for, it is possible to estimate a more realistic adjustment cost process.

¹ Some examples from this large literature include Dixit and Pindyck (1994), Abel and Eberly (1994, 1996, 1999) and Caballero and Engel (1994). Another explanation for the poor performance of the neo-classical model focuses on financial market imperfections. See, for example, Gross (1997) and Hubbard (1997).

² See Caballero, Engel, and Haltiwanger (1995), Cooper, Haltiwanger, and Power (1995), or Barnett and Sakellaris (1998).

³ Micro-level adjustment costs for output have been analyzed in Bresnahan and Ramey (1994) for the automobile industry and adjustment costs for labor have been explored by Hamermesh (1992) and Caballero, Engel, and Haltiwanger (1997). A general survey can be found in Hamermesh and Pfann (1996).

⁴ Some exceptions include Blackorby and Schworm (1983), Wildasin (1984), Chirinko (1993) and Hayashi and Inoue (1991) and the empirical work of Goolsbee (1999).

This paper deals with the problems of unobserved capital, unobserved market value, and capital heterogeneity by studying an industry with data on capital at a level of detail unavailable in previous work. These data include actual market prices, physical stocks and flows, utilization rates, and operating costs in every year for 16 different types of capital held by a panel of firms in the airline industry. With multiple capital goods, we observe multiple optimization decisions by each firm in each period so we can properly control for the unobserved components of firm profitability which affect investment demand. Intuitively, what matters for investment is the value of an additional seat on a plane. After controlling for difference in prices and costs of different types of planes, we can estimate this quantity because the firm internally optimizes its investment decision over each type of capital.

Following this approach, we can estimate the form of the micro-level costs of adjustment for both investment and utilization using a two-step procedure. First, we estimate the desired capital stock (and utilization) for each capital variety in each firm-year. Next, we non-parametrically estimate the investment function for the firm as a function of the gap between the firm's desired capital stock and its actual capital stock. The shapes of the estimated investment and utilization functions enable us to characterize the micro-level costs of adjustment on both of these margins.

The estimated marginal cost of adjustment cost functions show significant non-linearities. When we account for interactions between utilization and capital decisions, the non-convexities appear to take the form of a traditional region of inaction. Our results place bounds on the size of the inaction region. An airline only changes its output from a capital variety when desired output from that variety is at least 10% greater or at most 40% less than actual output. Conditional on positive or negative net investment, the estimated adjustment costs seem to be approximately linear in the amount of investment. The results on utilization (hours flown per plane) show smaller adjustment costs. Finally, we show that aggregating up to firm-level data or using more traditional accounting data obscures the true adjustment cost function and the evidence of non-convexities and biases the estimated adjustment costs upward.

The paper has seven sections. Section II describes some basic features of the airline industry and the data which we use in the paper. Section III discusses the estimation of the

“frictionless” stock of capital. Section IV uses the estimates from III to calculate the gap between actual and desired capital and to non-parametrically estimate the firm’s investment function and its implied costs of adjustment. These results are also compared to estimation results using only firm-level data without any capital heterogeneity. Section V describes the estimation of the firm’s utilization decision and dynamic interactions across both margins. Section VI concludes and an appendix describes the data in greater detail.

II. Airlines and Decisions About Capital

A. Airlines

Several factors make airlines an excellent place to examine hypotheses about investment decisions. The first is the long tradition of detailed record keeping on every aspect of the business including data on the stock, the purchases, and the retirements of capital for every type of plane for every major airline from 1978-94. There are also comprehensive annual data on the costs of operation for each plane in every airline-year, new and used price data for the planes, and data on the utilization (hours flown per day) of each aircraft. Second, purchasing and retiring capital goods at the right time is critical to carrier performance and the costs involved force them to take the decisions very seriously. Airline economics textbooks give substantial coverage to issues of investment.⁵ Finally, the cost structure of the airlines makes them responsive to observable shocks to capital and fuel prices.

Using these data, this paper presents some of the first direct evidence on the nature of micro-level adjustment costs on capital. While the availability of detailed data makes airlines an attractive place to estimate the form of adjustment costs and to search for evidence of non-convexities, it is important to note that airlines may not be representative of the broader economy. That said, the airline industry is important in its own right and widely followed by both the general public and economists.

⁵ Examples include James (1982) or O’Connor (1989).

Our study begins with the deregulation of the airline industry in 1978 and analyzes the capital stock decisions of the historically “major” airlines until 1994.⁶ In practice, the majors included any national carrier which started the sample with more than 50 jets: American, Continental, Delta, Eastern, Northwest, Pan-Am, TWA, United, and USAir.⁷ Eastern and Pan-Am went bankrupt in our sample.

The airline industry as a whole grew substantially in the 17 year period we examine and a substantial part of that growth was among the majors. Figure 1 shows a clear upward trend in the combined number of jet aircraft owned by the nine majors in our sample. The total number of jets almost doubles. The picture looks very similar if the sum is weighted by the number of seats or the prices of the planes.

In this paper, we are forced to abstract away from many of the complex decisions which go into the actual operation of an airline. Airlines jointly determine route structures, gate access, flight schedules, capital purchases and leases, utilization, personnel staffing, advertising, and prices for thousands of different flights.⁸ For tractability, we assume that airlines sell a single, homogenous good called travel-miles, which they use labor, energy, and many types of capital goods (planes) to produce. From this assumption, we explore how the airlines make capital decisions in response to shocks of various sizes.

B. Overview of Airline Capital Decisions

A common way to investigate non-convex adjustment costs in the existing literature, given the limited information about shocks, is to look at the “lumpiness” of investment.⁹ If firms do the majority of their investing in discrete bursts, this is often used

⁶ The term “major” (or “trunk”) is used in the data supplied by the F.A.A. and the C.A.B. and distinguishes these airlines from airlines classified as regional, commuter, or local.

⁷ Braniff and Southwest Airlines had enough jets to qualify but our empirical method requires variation between plane types for the same carrier. Southwest’s fleet is entirely 737s and Braniff’s was almost entirely 727s. Some airlines including Western, Republic, and Piedmont merged with airlines in our sample so they are generally included from the start of the sample. The appendix explains this in more detail.

⁸ Surveys of the extensive literature on the economics of airlines include Borenstein (1992) and Morrison and Winston (1986).

⁹ See the work of Doms and Dunne (1997) and Cooper, Haltiwanger, and Power (1995).

as evidence suggesting there are non-convex adjustment costs. If firms spread their investments out over time, this suggests more traditional, convex adjustment costs.

To get a preliminary look at lumpiness among airlines, if we examine either the number of planes or the dollar value of planes, there is no year in the sample for any airline in which investment is zero. Total investment is fairly smooth. This might suggest convexity. Looking at investment separately by aircraft type, however, shows a much lumpier pattern. For more than one third of the plane-years, there is no change in the number of planes of a given variety held by the airlines. This number even excludes the planes which an airline doesn't own which are always zero by definition.¹⁰ Simple descriptive statistics for the investment to capital ratio indicate heavy positive skewness and kurtosis, and that, on average, about 60% of the total acquisitions of an aircraft type by a given airline take place in the largest two-year investment episode. This seems to suggest non-convexity.

The micro data do not, however, consistently show either lumpiness or smoothness. A sample of the ownership and investment patterns for three of our airlines in selected years shows this in Table 1. For example, United increased the number of 727-200s by a large amount in a single year at the start of the sample and reduced the number by a large amount in a single year at the end of the sample. For United's 747s, however, the investments are spread out over the 12 years from 1982 to 1994. This diversity is typical of the various plane-years in the sample.

Furthermore, while the observed lumpiness or smoothness may be suggestive of the nature of the adjustment costs, it is by no means definitive. If investment is lumpy but the shocks to investment demand are even lumpier, investment will have high skewness and kurtosis but still face convex adjustment costs. To understand the nature of the airlines' adjustment costs, we need to look more formally at detailed data on capital and pay particular attention to the shocks. Identifying shocks to desired capital has been a key stumbling block in the previous literature because standard measures of shocks to investment demand like marginal Tobin's q are not

¹⁰ Many airlines hold no planes of various types throughout the sample. This is indicative of fixed costs associated with holding capital varieties, but we do not pursue such issues here. The prevalence of zeros also prohibits using some simple functional forms for the profit function below, such as Cobb-Douglas.

correct in the presence of many forms of non-convex adjustment costs.¹¹ They are also difficult to find for heterogeneous capital types.

Our airline data contain comprehensive cost and price information on many types of capital and are compiled by the Civil Aviation Board and AVMARK. These sources give us the operating expenses associated with every plane type for each airline in every year. Data at this level of detail is important and this paper is the first to combine data on observable shocks to the relative productivity of heterogeneous individual capital goods with observed investment in those same goods in order to identify the nature of the adjustment costs facing the firms.

In addition to choosing the number of planes of each variety, we also explore the airlines decisions concerning capital utilization, measured as the number of hours flown per plane-day. If we think of the airline as providing a good called travel-miles (and assuming a constant average speed) then output from a given variety of capital actually equals the number of planes times the number of hours flown per plane. Other researchers have studied how adjustment costs differ across various margins for other industries and we will be able to investigate this for the airline industry.¹²

C. What Kinds of Adjustment Costs do Airlines Face?

Since there is an active lease market for airplanes it is relevant to ask whether airlines face any investment adjustment costs at all and why airlines do not adjust to shocks immediately. Although it is possible for airlines to lease their capital every period, there are many potential sources of adjustment costs which apply regardless of whether an airline leases or buys its capital. Some examples include basic set-up costs such as re-configuring the cockpit and interior cabin of a plane to match the rest of the fleet, or the traditional expenses of hiring and training additional pilots and ground personnel to fly and maintain it. The overall convexity of these costs is not known *a priori* (large investments may have smaller or larger costs per unit) and they need not be symmetric (hiring and firing a pilot may have different costs, for example).

¹¹ See Caballero and Leahy (1996). They do show that under some forms of adjustment costs, average q can be a good approximation for the investment state variable. Abel and Eberly (1994) also examine the conditions under which q will be the appropriate measure.

¹² See for example Shapiro (1986) or Bresnahan and Ramey (1994).

Another important set of potential adjustment costs for an airline are the costs that new capital can have in disrupting a carrier's overall system operation. These adjustment costs include the need to change flight times and route structures or to acquire gate and hangar space. There may also be some carrier specific value in gates or landing slots which cause the purchase price and sale price of existing gates to differ. The hub system or the concentration of carriers at an airport can lead a gate in the United terminal at Chicago O'Hare (United's hub) to be more valuable to United than to some other airline, for example, creating a cost of adjusting the capital stock.

Many of the adjustment costs on system operations apply primarily to *net* investment rather than gross investment. An airline might incur little adjustment cost if it buys a plane that it uses to replace an old plane on some existing route. To buy a plane and expand by one route means changing the route structure and buying new gate space. Because of these potential net investment costs, it is also important to think about the impact of airport capacity (total departures expanded about 2.9% annually over this sample) as a potential reducer of the adjustment costs for net investment. We will address this issue later in the paper.

All of these adjustment costs imply that while, in theory, airlines could lease and return planes at annual frequencies in response to shocks, in practice the additional adjustment costs will induce them not to. When we look at the data in this sample, we see that among the major airlines in this time period, ownership of planes is the norm and the leases that do occur usually last many years. It appears that the decision to buy or lease a plane is primarily a financing decision and is somewhat tangential to our effort of estimating adjustment costs from adjusting the capital stock.

We will not spend time in this paper trying to identify the existence of external adjustment costs, despite frequent discussions about backlogs at Boeing and the evidence in Goolsbee (1998) of an upward sloping supply curve for aircraft.¹³ Instead, we rely on the micro data and assume that individual airlines are price takers allowing us to identify the firm level

¹³ The information on backlogs may also be of questionable value since at several points in the last fifteen years, Boeing has reported substantial backlogged orders while simultaneously laying off thousands of workers.

adjustment costs. So long as the impact of external adjustment costs affects all airlines, it will appear as a nuisance parameter in our results.

III. Estimating The Demand For Capital With No Adjustment Costs

A. Deriving the “Frictionless” Capital Stock

The basic approach of the paper is that with data on the gap between a firm’s desired and actual capital stocks, we can estimate the form of adjustment costs based on how much the firm’s actual investment differs from its desired investment. The key element in this formulation is calculating the desired capital stock. One approach would be to fully specify a structural model with a specific functional form for the adjustment cost function. Then the model, or a first order condition from the model, could be estimated and the relative magnitude of adjustment costs can be determined. This is a relatively simple exercise with quadratic or other convex adjustment cost models and it leads to both the q -theory of investment and the investment Euler equation. Unfortunately, with non-convex adjustment costs, most models cannot be solved analytically. While it is possible to use numerical techniques in conjunction with econometric estimation, such approaches often need to impose so much structure that they would lose the benefit of having the detailed, heterogeneous data on capital available in this paper.

Instead, we follow a semi-structural, two-stage approach similar to Caballero, Engel, and Haltiwanger (1995). In the first stage we estimate the “frictionless” stock of capital, K^f , by solving the firm’s problem as if it never faces adjustment costs. We assume that the frictionless capital stock is proportional to the desired capital stock, K^* , which is the capital stock which the firm would choose if adjustment costs were temporarily removed. This is not as strong an assumption as it may sound. Essentially any generalized (S,s) modeled can be expressed in this way, as long as the bands are constant functions of the models parameters.¹⁴ It does require, however, that there be no temporary shocks.¹⁵

¹⁴ This assumption is almost universal, either implicitly or explicitly, in the empirical (S,s) literature. For an investment example see Caballero, Engel, and Haltiwanger (1995), for consumer durables see Bertola and Cabello (1990), Caballero (1993), Eberly (1994), or Attanasio (1995), for employment see Caballero, Engel and Haltiwanger (1997).

¹⁵ Bertola and Caballero (1994) show that this relationship holds exactly in the case of irreversible investment, with an isoelastic production function, and a shock process with independent increments. The

The percent deviation of the firm's desired stock of capital from its actual capital, $\frac{K^*}{K}$, represents a measure of the firm's desired investment. In many ways the desired investment ratio is similar to a Tobin's q -type measure in that a value greater than one implies that the firm wants to invest and a value less than one implies that the firm wants to disinvest. In the second stage of the procedure, which is described in the next section, we non-parametrically estimate the firm's investment response as a function of its desired investment. The shape of this function will provide information about the adjustment costs facing the firm.

To solve for the frictionless capital stock, $K_{i,j,t}^f$, for aircraft type i , airline j , in year t , the firm maximizes profits as if there were no adjustment costs. This is equivalent to assuming that there is a perfect rental market for capital and having the firm solve

$$V_{j,t} = \max_{K_{i,j,t}} \Pi(K_{1,j,t}, \dots, K_{n,j,t}, G_{j,t}; c_1, \dots, c_n) - \sum_{i=1}^n cost_{i,j,t} h^* K_{i,j,t} - p_{i,t} (r_t + \mathbf{d}) K_{i,j,t}. \quad (1)$$

where $cost_{i,j,t}$ is the operating expense per hour for an airline's plane in a given year, h^* is the optimal hours per year, $p_{i,t}$ is the sale price of the aircraft type, r_t is the interest rate and \mathbf{d} is the depreciation rate. All flexible factors are assumed to have been maximized out. $G_{j,t}$ represents a composite index of all of the unobserved fixed factors and includes any demand or productivity shocks which affect the profitability of the entire airline. The c_i 's are fixed aircraft effects representing persistent differences in the profitability of different types of planes and are assumed to be the same across airlines. The maximization is performed for each K , conditional on the observed value of all of the other aircraft types for that airline.¹⁶

We assume that the profit function takes the form

profit function in our specification is isoelastic and we cannot reject that our shock process is a random walk. The assumption provides at least a reasonable approximation for many other adjustment cost functions.

¹⁶ An alternative would be to solve for the unconditional frictionless capital stock by maximizing the firm's profits for all aircraft types jointly, resulting in an expression for the frictionless capital stock of each plane as a function of all of that airlines costs and prices and fixed effects. We solve for the conditional K^f instead, because firms do not usually adjust all of their capital types simultaneously. Footnote 23 presents evidence in support of this assumption.

$$\Pi(K_{1,j,t}, \dots, K_{n,j,t}, G_{j,t}; c_1, \dots, c_n) = \left[\sum_{i=1}^n (c_i K_{i,j,t})^a \right]^{\frac{r}{a}} G_{j,t}^b \quad (2)$$

which has a constant elasticity of substitution $\mathbf{s} = \frac{1}{1-\mathbf{a}}$ between different types of capital, where $\mathbf{a} < 1$.¹⁷ When $\mathbf{s} = 0$, the capital aggregate is Leontieff, when $\mathbf{s} = 1$, the capital aggregate is Cobb-Douglas and when $\mathbf{s} = \infty$, the capital goods are perfect substitutes. As long as \mathbf{s} is not too large, different varieties of capital will be imperfect substitutes and capital heterogeneity will be important. Hours are assumed to be equal to their unconditional desired values since they are adjusted more frequently than capital so the value of h^* is contained in the fixed effects in the profit function.

The first order condition for $K_{i,j,t}$ sets the marginal product of capital equal to the neoclassical user cost.

$$r c_i^a K_{i,j,t}^{a-1} \left[\sum_{i=1}^n (c_i K_{i,j,t})^a \right]^{\frac{r}{a}-1} G_{j,t}^b = \text{cost}_{i,j,t} h^* + p_{i,t} (r_t + \mathbf{d}). \quad (3)$$

Taking logs implies that

$$(\mathbf{a} - 1) \ln K_{i,j,t} + \mathbf{a} \ln c_i + \left\{ \ln r + \left(\frac{r}{\mathbf{a}} - 1 \right) \ln \left[\sum_{i=1}^n (c_i K_{i,j,t})^a \right] + \mathbf{b} \ln G_{j,t} \right\} = \ln(\text{cost}_{i,j,t} h^* + p_{i,t} (r_t + \mathbf{d})). \quad (4)$$

Note that the entire term within the curly brackets is independent of the capital variety, i , and constant for each airline-year. The first two terms contain the effect of the capital aggregator, and the last term contains the unobservable fixed factors of production and stochastic shocks to airline profitability which are embodied in G . Since this term is independent of i and constant by airline-year, having data on multiple capital goods for the same airline-year allows us to estimate the bracketed term as a fixed airline-year effect and avoid the conventional unobserved variable problems.

¹⁷ Note that because of the G term, the profit function does not need to show overall constant returns to scale, though nothing rules it out. Kumbhakar (1990), Antoniou (1991) discuss the evidence on returns to scale in the airline industry. This formulation also imposes little restriction on the elasticity of substitution

Solving for $\ln K$, abbreviating the non-varying terms as fixed effects, and adding a residual gives us a frictionless capital stock equation which we can estimate directly:

$$\ln K_{i,j,t} = \tilde{c}_i + \tilde{\mathbf{h}}_{j,t} + \frac{1}{(\mathbf{a} - 1)} \ln(\text{cost}_{i,j,t} h^* + p_{i,t}(r_t + \mathbf{d})) + \mathbf{e}_{i,j,t}. \quad (5)$$

This is the first stage regression which yields the frictionless capital stock for each airline-plane-year, $K_{i,j,t}^f$, as the fitted value from equation (5) for each observation.

When we take into account utilization, the firm optimizes over hours as well as capital. Adding hours to the profit function yields an expression in total hours flown

$$\Pi(K_{1,j,t}, \dots, K_{n,j,t}, h_{1,j,t}, \dots, h_{n,j,t}, G_{j,t}; c_1, \dots, c_n) = \left[\sum_{i=1}^n (c_i h_{i,j,t} K_{i,j,t})^{\mathbf{a}} \right]^{\frac{r}{\mathbf{a}}} G_{j,t}^b. \quad (6)$$

Replacing it into equation (1), taking the first order condition with respect to h , and renormalizing the fixed effects yields a frictionless total hours equation for each type of capital.

$$\ln(h_{i,j,t} K_{i,j,t}) = \hat{c}_i + \hat{\mathbf{h}}_{j,t} + \frac{1}{(\mathbf{a} - 1)} \ln(\text{cost}_{i,j,t}) + \mathbf{m}_{j,t}. \quad (7)$$

We estimate equations (5) and (7) in the next section.

Two specification issues arise with regard to estimation of the frictionless equations. The residual of equation (5) represents deviations of the actual capital stock from the estimated frictionless capital stock. It arises because of adjustment costs, and is obviously correlated with the independent variables and their lags. However, the frictionless capital stock equation can still be consistently estimated with OLS because the dependent and independent variables are cointegrated. This must be true because if there is a large deviation between the frictionless and actual capital stocks, it is worthwhile for the firm to pay the adjustment cost and close the gap.^{18,19} By the same argument, the desired capital stock is also cointegrated with both the frictionless and actual capital stocks.

between capital and non-capital inputs. In particular, it is valid under the (plausible) assumption that the capital goods are Leontief in pilots and flight attendants.

¹⁸ The time dimension of the data is not especially long but we did conduct an augmented Dickey-Fuller test and it did not reject that the capital stocks and costs variables have a unit root and that equation (5) is cointegrated. See Hamilton (1994) for a more complete discussion. Footnote 23 also presents evidence on the cross-sectional independence of the observations.

¹⁹ See Bertola and Caballero (1990) for a further discussion.

Second, most airlines do not own all types of planes. The obvious explanation for this fact is the presence of fixed costs associated with holding a capital variety. Since we are interested in adjustment costs associated with investment, we only consider observations in year t when the airline had a positive capital stock during both years $t-1$ and $t+1$. Years when the firm decides to acquire a new type of plane for the first time or sell all of its remaining planes will depend on both investment adjustment costs and the fixed costs of holding a capital variety, which we want to avoid. Any bias which results from this sample selection criteria will tend to bias us away from finding non-convexities.

B. Results

The results of estimating equation (5), the frictionless capital stock regressions, are listed in table 2. Column (1) measures the capital stock as the number of planes at the end of the year as compiled by the International Civil Aviation Organization (ICAO). The cost term is the log of the sum of yearly operating expenses and interest and depreciation times the price. The estimated coefficient on the cost variable corresponds to $-\sigma = \frac{1}{\alpha - 1}$, the negative elasticity of substitution. It is of the predicted sign and highly significant. The elasticity of substitution is 1.5 implying an α of 0.33. This coefficient is estimated precisely and suggests that there is substitution across varieties but that heterogeneity between planes is fairly important.²⁰ The gaps between the desired capital stock generated from this regression and the actual capital stock show considerable variation both across time and across planes. The standard deviation of the capital gap across years for the median airline-plane is .272 while the standard deviation of the capital gap across planes for the median airline-year is .613.

²⁰ This elasticity of substitution between capital varieties is slightly smaller than the elasticity of substitution between capital varieties in the machinery industries studied in Goolsbee (1999). This is partially due to the restriction that the elasticity be the same across all model types. We also investigated modifying the profit function to allow separate elasticities of substitution between different types of planes. The cost was making the empirical specification much less tractable. Allowing different elasticities of substitution between McDonnell-Douglas planes and between Boeing planes, we found that the within group substitution elasticities were larger (around 3) but also with larger standard errors. This specification did not change our qualitative findings in the second stage.

Since the fitted values from this regression determine the frictionless capital stock used throughout the paper, the accuracy of the first stage specification is important for our estimation of adjustment costs.²¹ To show that the equation performs well, we first look at the plane and airline-year effects in more detail. F tests on the plane and airline-year effects show that each set of coefficients is jointly significant at the 0.1% level. More importantly, the fixed effects estimates appear plausible given the specification of the profit function. Since the plane effects measure the persistent component of the marginal product of that capital variety, larger planes—which hold more passengers per unit of capital—should have larger plane effects. Figure 2 plots the log of the number of seats against the estimated plane effect and shows a strong positive relationship. Similarly, the airline-year effects measure the relative productivity of a given airline in a given year. Figures 3(a) and 3(b) compare the airline year effects for the two most successful airlines in the sample, American and United, with two of the biggest money losers, Eastern and TWA. The airline-year effects also seem to be doing a good job picking up carrier performance.

A second test of the specification comes from columns (2)–(4) where we investigate various modifications of the data to check robustness. In column (2), we use the data from AVMARK and the C.A.B. on capital assigned, rather than the ending capital of the ICAO. We find a similar coefficient and therefore a similar α . Column (3) repeats the analysis using the ending capital stock but adds taxes to the user cost and column (4) repeats the analysis only for airlines which did not participate in a merger (5 of the 9 airlines). None of the estimated coefficients are significantly different from one another and they all imply an alpha parameter between 0.20 and 0.45. The average correlation among the residuals is 0.96.²²

Finally, in column (5), we present a more stringent specification check by estimating equation (7) for total hours flown rather than for capital, as in the previous equation. As equation (7) dictates, the log cost variable for the hours specification does not include the price term, only the operating costs. Since it comes from the same profit function, the elasticity of

²¹ Note that the very high R^2 values in table 2 are partially due to the differences in capital stock levels across airlines and capital varieties and the presence of plane and airline-year fixed effects.

²² We also ran a specification with used prices rather than new prices and a specification incorporating expected future capital gains and the substitution elasticity of both were almost identical to those presented.

substitution in the hours equation should, in principle, be the same as in the capital equations of columns (1)-(4). When we do this estimation, the coefficient is indeed very similar and implies an alpha of about 0.25 which is not significantly different from any of the coefficients on capital. We will use residuals from column (5) when we jointly examine the capital and utilization decisions.

The results from the first stage show that the frictionless regressions fit the data well, are precisely estimated, seem to give reasonable conclusions for the plane and airline-year effects, give a self-consistent elasticity estimate for the hours and capital equations, and are robust across various specifications. For these reasons, we feel confident using them as a basis for determining the frictionless capital stock.

IV. Estimating the Form of the Costs of Adjustment

A. Methodology

In the second stage, we investigate adjustment costs by comparing the firm's actual investment, which is observable, to its desired investment which comes from the first stage regression.²³ We take a non-parametric approach, allowing the qualitative nature of the adjustment process to be different at different levels of desired investment.²⁴ We will estimate functions of the form

$$\frac{I_{i,j,t+1}}{K_{i,j,t}} = f\left(\frac{K_{i,j,t}^*}{K_{i,j,t}}\right) + I_{i,j,t} \quad (8)$$

²³ We are treating the investment decisions for each of the plane types as if they are independent of one another. To check this assumption we ran linear regressions of the investment for plane type a on its own desired investment and on the desired investment of some other type b (where data were only included for airlines that had both types of the planes in their fleets simultaneously). If the decisions are not independent, we would expect that investment would also depend on the desired investment of other planes. Of the 180 different pairwise investment regressions, only 8 percent of them showed a significant coefficient on the other plane type at the 5 percent level (14 percent at the 10 percent level). The exceptions were mainly in cases one might expect such as investment in 737-200s also depending on the firm's desired investment in 737-300s. Overall, the independence assumption does not seem too inaccurate and it greatly simplifies the estimation.

²⁴ All regressions use the Epanechnikov (quadratic) kernel. Further details on the method as well as the mechanism for choosing the bandwidth are described in the appendix. An explanation of kernel regression can be found in Hardle (1990). Gross (1997) applies such methods to investment with financial market imperfections.

using a Nadarya-Watson kernel estimator which puts almost no restrictions on the shape of the function f . Based upon the theory, this function is the inverse of the marginal adjustment cost function. The shape of the resulting function gives us direct evidence on the form of the firm's adjustment costs.

The percent deviation of the firm's desired level of capital from its actual capital, $\frac{K_{i,j,t}^*}{K_{i,j,t}}$, represents a measure of the firm's desired investment. Values greater or less than one imply that if adjustment costs were temporarily removed the firm would invest or disinvest. The fitted values from equation (5) provide estimates of the frictionless capital stock. Multiplying by a constant scale factor, \mathbf{q} , and dividing by the actual capital stock gives the desired investment ratio.

$$\frac{K_{i,j,t}^*}{K_{i,j,t}} = \mathbf{q} \frac{K_{i,j,t}^f}{K_{i,j,t}} = \mathbf{q} \exp(-\mathbf{e}_{i,j,t}). \quad (9)$$

Other authors have made strong identifying assumptions to determine an actual value of \mathbf{q} , in some cases even allowing it to vary by firm.²⁵ These assumptions will clearly influence the results and can directly alter the qualitative nature of the estimated adjustment cost function. We choose not to specify \mathbf{q} and instead present results as if $\mathbf{q} = 1$ with the understanding that the scale of the horizontal axis may be off by a constant factor. Since the airline industry is trending upward, it is likely that desired capital is greater than frictionless capital and $\mathbf{q} > 1$. Regardless of the value of \mathbf{q} , however, since our conclusions about the existence of non-convexities depend only on the shape of the function, they will not change with \mathbf{q} .

B. Theory and the Shape of the Adjustment Cost Function

Figure 4 illustrates the shape of the desired investment function, f , in equation (8) for different specifications of adjustment costs.²⁶ With no adjustment costs at all, f will cross the horizontal axis when $\frac{K^*}{K} = 1$, with a slope equal to one. In other words, any deviation of desired capital from actual capital leads to an immediate increase of exactly that amount. With

²⁵ See Caballero, Engel, and Haltiwanger (1995).

²⁶ See Dixit and Pindyck (1994) and Abel and Eberly (1994), for example, for more complete descriptions.

quadratic adjustment costs, the relationship will also be linear but with a slope of less than 1. This is the partial adjustment model, with the firm closing a constant part of the gap between the desired and actual capital stocks each period. This is the key assumption of the conventional investment literature.

Irreversibility or other large downward adjustment costs appear as a flat portion for small values of $\frac{K^*}{K}$, indicating that although desired capital is low, firms do not disinvest. Partial irreversibility, i.e., a wedge between the purchase and sale price of capital will generate a flat “region of inaction” with zero investment around $\frac{K^*}{K} = 1$.²⁷ The wedge creates an option value of waiting when the desired investment ratio is near one. A similar wedge will also arise from asymmetric adjustment costs for positive versus negative investment.

More generally, arbitrary non-convexities in adjustment costs will appear as convexities (concavities) in f for desired capital greater (less) than actual capital, indicating that large absolute values of desired investment lead to proportionately larger changes in actual investment than small values. This general pattern may arise empirically if the cutoff points in a region of inaction differ by firm or by capital type or simply because the smoothing of our non-parametric estimator makes it difficult to identify a “kink” in the adjustment cost function. This pattern may also arise if adjustment costs are less convex than quadratic but still imply that small changes in fundamentals lead to proportionately even smaller changes in investment.²⁸ For this reason, we are primarily concerned with evidence of non-linearity in the function f .

C. Results

²⁷ Fixed costs will result in a similar region of inaction, but the investment function will be discontinuous.

²⁸ Specifically, this would be the case of an adjustment cost function which is less convex than quadratic but still convex, such as $C(I) = |I|^{1+\epsilon}$, $0 < \epsilon < 1$. In this case the adjustment cost function is differentiable when investment is 0, so formally, there will not be any lumpiness or region of inaction. Since empirically, it will be difficult to significantly distinguish between a true region of complete inaction and a highly non-linear response to shocks, we will simply treat this case the same as our generalized non-convex adjustment costs case and test for the non-linearity. The economic implications are quite similar since in either case small changes in fundamentals will lead to less than proportional changes in investment and investment will not appear smooth since its reaction to shocks will be non-linear and the basic assumption of the standard partial adjustment model will be violated.

Our basic non-parametric regression of equation (8) is presented in figure 5 for net investment (purchases minus retirements) using the residuals from the frictionless regression. The pattern is notable. Although we cannot pin down the position of the horizontal axis because of the scale factor, the shape of the estimated function suggests that there is a basically flat portion and a positively sloped, linear portion. This implies several interesting things about the airlines' capital adjustment costs.

First, there is an apparent non-linearity in the marginal adjustment cost function. When the desired capital stock is less than the actual capital stock, the firm does not reduce its capital stock. This seems to be evidence of some type of irreversibility, or at least of large adjustment costs associated with disinvestment. The observed negative shocks are not large enough to cause reductions in the capital stock on average.²⁹

A well known problem with kernel estimators is the large standard error bands they generate. We bootstrap the standard errors and show the 95 percent, pointwise confidence bands in this figure. This bootstrapping takes desired investment as if it were measured precisely when it is, in fact, the residual from the first stage regression. However, since the rate of convergence for the kernel estimator is slower than the rate of convergence for the first stage linear regression, asymptotically, the error bands around the kernel need not account for the first stage error. That said, the precision of the first stage estimates is likely to imply that the correction would be small. To illustrate this we also display the 95% confidence bands from bootstrapping the entire two step procedure (sampling the original data with replacement, running the frictionless regression on each sample and calculating the desired investment ratio, and then estimating the kernel regression on each sample). They are quite close to each other.

The shape of the function suggests non-linearity, and we cannot reject that the region of inaction is at zero, as predicted by models with non-convex adjustment costs. However, the imprecision of the standard errors bands, implies that we cannot reject alternative shapes for the form of the adjustment costs including linearity. Conducting a Chow test for a structural break at the capital ratio of 0.8 does reject the null of no structural break at the 5 percent level, although

without a correction for the randomness induced in the first stage. (The coefficient for desired investment is zero for values less than 0.8 and about 0.11 for values above 0.8.).

Since our primary goal, however, is not to test the location of the adjustment cost function but rather the shape, the confidence intervals and the Chow test do not really answer the question about the form of the adjustment costs. We are primarily concerned with evidence for non-linearities in the marginal adjustment cost function. If the estimated curve from each sample has a non-linear shape but each is shifted up or down by some amount, this will generate a wide confidence interval even if there were clear evidence of non-linearity in every case. To explicitly test whether the relationship behind figure 5 is convex. Abrevaya and Jiang (1999) provide an explicit, non-parametric test of his hypothesis.³⁰ If the value of their statistic significantly exceeds 0.5 then one can reject the null of linearity in favor of convexity. Note this test is conducted using the full data set and not the estimated function. For the data generating the function in figure 5, we find the convexity statistic to be 0.543 with a standard error of 0.008. This yields a z statistic of 5.3, strongly rejecting the null of linearity at any conventional level of significance. We also bootstrapped the entire procedure, to take into account the error induced by the first stage estimation.³¹ The average statistic was 0.540 with a standard error of 0.009, corresponding to a z statistic of 4.55, again highly significant evidence of non-linearity. In other words, the data indicate significant evidence of non-linearity in the marginal adjustment cost function, and the kernel estimates suggest a more specific form, namely a region of inaction.

It is important to note that the dependent variable in figure 5 is *net* investment so the notion of irreversibility implied by the shape of the function is somewhat different than the conventional notion. Clearly there is not complete irreversibility of gross investment since we observe firms disposing of used airplanes in existing secondary markets. Figure 6 presents a

²⁹ Abel and Eberly (1996) show how even modest price wedges can generate conditions looking very much like complete irreversibility. Later in the paper, we will construct a more informative measure of desired investment which will show downward adjustment following large negative shocks.

³⁰ The intuition behind the Abrevaya and Jiang test is to examine all possible collections of three points from the data, and count how many times the slope between the pair with the higher “x” values is smaller than the slope between the pair with the lower x values. A convex function will have a higher slope, the larger the x value of the points.

non-parametric regression of aircraft purchases and retirements separately on the desired investment ratio.³²

The gross investment function in figure 6 is flat not at zero investment (as in a gross irreversibility model) but rather at a positive level, approximately equal to the level of capital retirements. This suggests that airline adjustment costs are function of *net* investment (like in figure 5) rather than gross investment. As we discussed above, it is intuitive to think that the many capital adjustment costs for airlines actually do apply to net rather than gross investment (*i.e.*, selling a plane does not require changing your route structure, so long as it is replaced with a new plane) so this finding is not especially surprising. The retirements series is basically constant at all levels of desired investment, perhaps consistent with regulation induced retirement or with one-hoss-shay depreciation.

The second thing that we can say about the form of adjustment costs in figure 5 is that conditional on positive net investment, the desired investment function seems basically linear. This is consistent with quadratic adjustment costs in the region of positive investment. The slope of the linear portion is about 0.1. This is steeper (meaning smaller adjustment costs) than in much of the literature on q but, unfortunately, both the position of the linear portion and its slope will be affected by the value of q so we can only be confident only about the general shape of f . This is also a problem for any attempt to quantify how much greater than 1 the shock must be to commence positive investment. The graph indicates that the breakpoint is below 1 but a value of q greater than one can shift that kink point to the right of $\frac{K^*}{K} = 1$ (as predicted by most non-convex adjustment cost specifications).

Figure 7 verifies that the previous results are not due to problems in the price or cost variables. If we simply assume that prices and costs are the same for each airline in each year, then we can replace the price and operating expense data in the first stage regressions with fixed plane-year effects. Using the residuals from this regression, we repeat the non-parametric

³¹ The argument about the rates of convergence for Kernel estimators does not apply to the Abrevaya and Jiang test since it is root-n consistent. However, since the results are so similar given the accuracy of the first stage, we do not bootstrap the later tests.

estimation in figure 7. Although the results are somewhat noisier, the pattern remains exactly the same—non-convexity downward and basically linear upward. The Abrevaya-Jiang test strongly rejected linearity.

D. Heterogeneity vs. Firm Level Data

Next we show how estimates of the adjustment costs which use only more traditional, firm-level data fail to capture the nature of the actual adjustment costs. To explore this question we look at two different measures of capital for each airline—market value and book value. We create the market value of the entire stock of planes for each airline in each year by summing the market values of each individual aircraft. This market value measure eliminates all information on different capital varieties but it is still a much better measure of capital than is typically available since we observe true market prices of capital. We fit the first stage regression for frictionless capital using this market value of capital series and the average operating expenses and aircraft prices for the airline in that year.³³ Using the residuals from this regression, we repeat the non-parametric analysis on the aggregated data and the results are shown in figure 8. Even at this level of aggregation, all evidence for non-convexity has disappeared and the slope of the estimated function is 0.085, is biased downward by about 20% compared with the linear portion of the function in figure 5.

When we take the further step of using the conventional perpetual-inventory method to construct the capital stock from COMPUSTAT accounting data, the results are even more biased.³⁴ For this regression, we do not use any detailed data on operating costs and instead fit the first stage with the standard user cost of capital. When we repeat the non-parametric analysis, shown in figure 9, there is again no evidence for non-convexity and the estimated function is even flatter. The slope is now only 0.05, less than half of the estimated value using heterogeneous capital data.

³² Using Abrevaya and Jiang test for convexity, we can strongly reject linearity for the gross investment series, but as expected, cannot reject linearity for retirements.

³³ Of course we cannot include airline-year effects in either the market value or book value specifications. Instead we include airline dummies and a time trend. The results were similar when we use fixed time effects instead of the trend and drop the price and cost variables.

³⁴ See the appendix for a more complete description.

In other words, looking exclusively at firm level data, we can only find evidence of convex adjustment costs and even these are rather seriously biased. Since we have shown that capital goods in the airline industry are not perfect substitutes, this result is incorrect and only the micro-level, heterogeneous capital data can give us an accurate description.³⁵

V. Joint Determination of Capital and Utilization

A. Utilization

Airlines have the ability to adjust their utilization rates (hours flown per day) as well as their capital stocks. In this section we investigate the size and type of adjustment costs that airlines face when we look at both of these decisions.³⁶ We first look at hours alone. Equation (7) provided estimates for the frictionless value of total hours flown by plane type, $h_{i,j,t}^f K_{i,j,t}^f$. Combining these estimates with the frictionless capital regression yields estimates of the desired hours ratio.

$$\frac{h_{i,j,t}^*}{h_{i,j,t}} = \mathbf{q}_h \frac{\exp(-\mathbf{m}_{i,j,t})}{\exp(-\mathbf{e}_{i,j,t})}. \quad (10)$$

We then non-parametrically estimate equation (11) relating the percent change in hours to the desired hours ratio:

$$\frac{\Delta h_{i,j,t+1}}{h_{i,j,t}} = f_2 \left(\frac{h_{i,j,t}^*}{h_{i,j,t}} \right) + \mathbf{n}_{i,j,t}. \quad (11)$$

This is exactly analogous to equation (8) for capital investment. This non-parametric regression is graphed in figure 10.

³⁵ In order to investigate why aggregation to the firm level produced these results we performed a simple Monte Carlo study by aggregating randomly generated data using the estimated heterogeneous capital investment function. As the amount of aggregation increased, the non-convexity disappeared and the estimated slope declined. This is a basic property of convex functions and is not surprising since, in some sense, we are averaging over the kink point the more we aggregate. It is also consistent with previous empirical observations. At the level of the entire economy, investment is sufficiently smooth that, until very recently, quadratic adjustment costs were taken by researchers to be a reasonable approximation.

³⁶ In an earlier version of this paper we estimated the adjustment costs for a second utilization margin, the share of seats filled (load factor). Not surprisingly, given the annual frequency of our data, we found almost no adjustment costs on this margin.

Like the case of investment, the adjustment function for hours is convex. There does appear to be a greater slope to the function when desired hours are high than when they are low. This indicates some type of non-convexity in the adjustment cost function for changes in hours. Unlike investment, repeating an F test with two linear line segments does not detect a structural break. This is not surprising given the absence of a kink in the estimated function. The Abrevaya and Jiang (1999) test for convexity results in an estimated statistic of 0.516 with a standard error of 0.008 or a z statistic of 1.94. Therefore in this one-sided test, we can reject linearity at the five percent level.

The most interesting thing to note about the response of hours is that the slope of the upward part is much steeper than the corresponding slope in the investment function. This indicates that smaller costs of increasing hours cause firms to close an hours gap about four times as fast as a capital gap. When desired hours are high, airlines close about 40% of the gap per year, compared to only 10% for investment.

B. Total Output and Interactions across Adjustment Margins

The previous sections have documented how an airline adjusts its capital and utilization rates in response to shocks. We will now examine interactions between the two margins by investigating the response of total output, hK . Does an airline that wants to increase its output increase the number of planes in operation, the number of hours per plane, or both? In order to explore these interactions we construct the desired output ratio, $\frac{h^*K^*}{hK}$, and see how it is related to firms' actual output, investment, and hours decisions. Since the desired output ratio incorporates information on both margins of adjustment, it should be a more informative signal than either desired hours or desired capital alone.

We use the same methodology as in previous sections. The desired output ratio is constructed from the residual of equation (7) adjusted by the scale factor, \mathbf{q}_{hk} :

$$\frac{h_{i,j,t}^*K_{i,j,t}^*}{h_{i,j,t}K_{i,j,t}} = \mathbf{q}_{hk} \exp(-\mathbf{m}_{i,j,t}). \quad (13)$$

We then non-parametrically estimate equation (14) to compare actual output to the desired output ratio.

$$\frac{\Delta h_{i,j,t+1} K_{i,j,t+1}}{h_{i,j,t} K_{i,j,t}} = f_3 \left(\frac{h_{i,j,t}^* K_{i,j,t}^*}{h_{i,j,t} K_{i,j,t}} \right) + \mathbf{x}_{i,j,t}. \quad (14)$$

This is analogous to the previous investment and hours specifications. Figure 11 shows the estimate of equation (14).

This picture for total output is very much in line with the theoretical discussion presented above. As in the case of net investment, there is a clear upward sloping, linear region for large positive shocks. However, there is now also a downward sloping, linear-type region when desired output is low. In between, the middle portion of the figure shows an apparent region of inaction. The shape of the figure seems to correspond rather closely to a model with non-convex adjustment costs when desired output is near actual output and quadratic adjustment costs conditional on large deviations. The Abrevaya-Jiang test indicates that the upward portion of the graph to the right of 0.8 is significantly convex (statistic of 0.524 with a standard error of 0.01 resulting in a z statistic of 2.4). For the downward portion of the curve, although the point estimate of the test suggests concavity, the standard error indicates that this difference is not significant (statistic of 0.494 with a standard error of 0.008).

Note that the value of the change in output in the inaction region is approximately 3% instead of zero as standard theory would predict. However, as mentioned before, during this sample period airport capacity has expanded by 2.9%. As we will discuss in the next section, one of the leading explanations for the estimated non-convex adjustment costs is a wedge between the purchase and sale prices of existing gates and landing slots arising from the hub system and economies to density. The value of an existing gate or landing slot is carrier specific, creating a wedge in its price. If an airline does not face this wedge if expands at the growth rate of airport capacity, this would shift the estimated function downward by 0.029, causing the flat region in figure 11 to be at the theoretically correct value of zero.³⁷

³⁷ This is analogous to previous studies which assume no adjustment costs on replacement investment and therefore calculate adjustment costs as a function of gross investment minus the depreciation rate.

The size of the region of inaction is substantial. $q_{hk} = 1$ provides a lower bound for the size of the inaction region in the upward direction, and an upper bound in the downward direction. Desired output from a capital variety must exceed actual output by at least 10%, before a firm increases output more than the trend growth in capacity. Similarly, firms reduce output compared to trend when output exceeds desired output by at most 40%. In the quadratic adjustment cost regions, the estimated slope for positive shocks is approximately 0.2, and for negative shocks, 0.1, indicating adjustment costs are twice as high for reductions in output than for increases.

Equations 15 and 16 investigate how changes in output are distributed across the investment and hours margins.

$$\frac{I_{i,j,t+1}}{K_{i,j,t}} - g = f_4 \left(\frac{h_{i,j}^* K_{i,j,t}^*}{h_{i,j,t} K_{i,j,t}} \right) + \mathbf{x}_{i,j,t}. \quad (15)$$

$$\frac{\Delta h_{i,j,t+1}}{h_{i,j,t}} = f_5 \left(\frac{h_{i,j}^* K_{i,j,t}^*}{h_{i,j,t} K_{i,j,t}} \right) + \mathbf{z}_{i,j,t}. \quad (16)$$

The specifications are similar to previous regressions, however we use the more informative desired output ratio instead of desired investment or hours alone. In addition, we correct for the growth in airline capacity of 2.9% in equation (15).³⁸ The results are graphed in figure 12. The dotted line represents investment and the solid line hours.

Unlike the previous investment and hours results, using the more informative desired output ratio now seems to indicate regions of positive and negative adjustment for both investment and hours. Estimated adjustment costs for investment are substantially larger than for hours. If we interpret the point estimates as indicating regions of inaction, the width of the inaction region for investment is approximately four times as large for hours. For both hours and investment, in the quadratic adjustment cost regions, increases in capital or hours seem less costly than reductions but the slopes are not significantly different.

³⁸ The correction is employed for investment and not for hours because hours are stationary while capital can trend in the long run. If the correction were used in the initial results in figure 5, the flat region shifts from 0.01 to -0.02.

C. *What are the Sources of Adjustment Costs in the Airline Industry?*

If we want to go farther than just estimating the form of the adjustment costs facing airlines and actually attempt to conclude something about what the source of those costs might be, we must find candidates that are consistent with the observed evidence. The observed patterns of airline adjustment in this paper, while not specifically proving any one source of adjustment costs, can rule out some types of adjustment costs.

First, any costs that enter into the plane fixed effects cannot explain the shapes of the estimated relationships. This includes anything that increases the cost of buying a plane by some fixed amount such as yearly maintenance, additional safety equipment, and the like. Second, any fixed costs associated with buying the first of some type of aircraft such as establishing maintenance facilities and parts networks for the new plane type cannot explain the observed costs in our data. The fact that every airline has a zero stock of several types of aircraft indicates that such costs may be important but our results condition on positive capital stock.

The two key characteristics of the adjustment costs estimated are: one, the adjustment costs are on net rather than gross investment (because investment and retirements are both positive and approximately equal when the desired investment ratio is less than one) and two, the adjustment seem to entail a region of inaction.

The net versus gross investment distinction rules out many potential sources of adjustment costs. Costs associated with the direct expenses of installing an airplane (the one time costs of reconfiguring the cockpit, painting the plane, *etc.*) would imply that there are adjustment costs associated with each purchase or sale of capital—costs on gross investment. Net adjustment costs also rules out production lags at Boeing or other aggregate convexities as the source of the adjustment costs estimated here. If airlines were rationed on the supply side, they should not have positive investment approximately equal to the rate of retirements when their desired investment ratio is less than one.

A region of inaction can arise for two reasons.³⁹ A first source would be asymmetric adjustment costs, i.e., if the linear portions of the adjustment cost curve have different slopes, this will lead to a region of inaction. The slopes of the two regions in our results are close to one

another and with the standard error bounds being what they are, we are unable to provide a direct test for the impact of this factor. This is, therefore a possible explanation but intuitively we prefer the second.

The second explanation is a difference between the purchase and sale price of capital. Although it is true that there are differences between sale and purchase prices of airplanes, this type of partial irreversibility in gross investment, however, would lead the region of inaction to be at zero investment not at the rate of aircraft retirements.⁴⁰ We have argued that the barrier is not likely to form from gross investment in the planes but rather in net investment because of the gates or routes.

To generate a region of inaction in the adjustment cost function for net investment, an existing route, gate, or landing slot must be worth more to the existing owner than to another airline. Network externalities and the benefit of the hub system provide a natural explanation of this phenomenon. Important work on the airline industry has argued that fares are higher and marginal costs are lower for carriers with a hub in a particular city than for non-hub carriers.⁴¹ Berry, Carnall, and Spiller (1997) find that markups are between 30% and 60% higher (due to both higher prices and lower costs) for flights from a hub than flights not from a hub. If we think of this a wedge between the value of a route to an incumbent and the value to a new entrant, the magnitude can comfortably explain the region of inaction in our estimates and are an intuitive explanation of the estimated form of adjustment costs.

VI. Conclusion

In this paper we have set out to directly estimate the form of the micro-level marginal adjustment costs of investment facing airlines by gathering data on capital at a previously unavailable level of detail. Information on heterogeneous capital goods along with prices and costs for each variety allow us to control for unobservable changes to airline profitability as well

³⁹ See the discussions in Dixit and Pindyck (1994) or Hamermesh and Pfann (1996).

⁴⁰ The difference in prices between new and used versions of the same model are quite large but this is misleading. The issue is whether the identical plane has a different value for current user than for the potential buyer. The price difference between very late model used aircraft and the identical new counterpart averages around 10 percent in the data from the *Airliner Price Guide* (1993).

⁴¹ See Borenstein (1989, 1991), Breuckner, Dyer, and Spiller (1992) and Breuckner and Spiller (1994).

as unobserved components of capital. In a two step procedure, we first estimate the desired stock of capital for each aircraft type and then look at the investment response to the gap between the desired and actual capital stocks.

The results from using such detailed data provide significant evidence of non-linearities in the marginal adjustment cost function. Kernel estimates point to the Airlines having a sizable region of inaction when adjusting their capital stocks. Desired output must differ from actual output by between 10% and 40% before investment takes place. Conditional on investing, however, adjustment costs seem basically quadratic. The adjustment cost function for utilization follows a similar pattern with a much smaller inaction region and much smaller implied quadratic adjustment costs. When we aggregate the data even to the firm level, we lose all evidence of non-convexities and our estimates of the adjustment costs are biased upward. The apparent importance of the region of inaction for these firms provides empirical support to the expanding literature which examines the theory of investment under non-convex adjustment costs.

Data Appendix

1. Airline Data

The data on aircraft come from the International Civil Aviation Organization's *Fleet Statistics* from 1978-1995. For each airline in each year the data track the number of planes at the start of the year, the number acquired, the number disposed of and the number at the end of the year. In years in which the beginning capital stock did not match the ending of the previous year, we checked the data using *Moody's Transportation Manual* and the *Jet Airliner Production List*.

To match our various data sets, to have meaningful numbers of each type of aircraft, and to make the estimation manageable, we condensed the aircraft into 16 types—707s, 727-100s, 727-200s, 737-200s, 737-300s, 747s, 757s, 767s, DC-10s, DC8-5, DC8-6s, DC9-1s, DC9-3s, DC9-5s, MD-80s and L-1011s. More detailed models within each category were combined. For some airlines the data are reported combining two types planes—727-100s and -200s, for example—as if they were all 727-200s. In those cases we were forced to treat them all as 727-200s. We excluded propeller aircraft and aircraft which lacked either price or cost data for the entire sample as well as planes which were not owned by more than two airlines for at least three years. In terms of the major jet types, this only ruled out Airbus planes.

The data on operating costs for 1978-1985 come from the Civil Aviation Board's *Aircraft Operating Cost and Performance Report* and for 1985-1994 come from AVMARK's *Quarterly Aircraft Operating Costs and Statistics*. These data present the average costs for every type of aircraft for every carrier in each year. We use cash operating expenses which includes fuel, maintenance, labor and insurance but not accounting depreciation. It is measured in dollars per hour and we turn it into real dollars using the GDP deflator. To get it into the same units as the user cost we multiply the costs times the average hours flown per plane-year (3400) and get a cost per plane-year in millions. These operating cost data also provide the total revenue block hours flown (total hours flown in all planes of that type) and a measure of the total aircraft in operation. We corrected obvious outliers in the hours series which usually involved simply a misplaced decimal point on the part of the publication.

The cost data is kept at a finer level of detail than the fleet data so we create a weighted average of the operating expenses where the weights are the total revenue block hours for the plane type. This was necessary in some cases for airlines where costs are listed for domestic and international operations of the same plane (like the 747), in some cases slightly different models of the same basic plane are listed separately and in a few cases the small number of planes forced us to combine similar planes into one category (737-300s and 737-400s). In some years, there was cost data for a type of plane for a given airline but no record in the fleet data of them owning the plane and vice versa. We checked the annual reports and the *Jet Airliner Production List* to determine the truth. If the cost data were missing, we use the average for that plane type in that year for other airlines. If the fleet data were missing we located them from the other sources.

There were a few large mergers in our sample. In general we treated small mergers as standard investments, facing the same adjustment costs as other purchases. In four cases, however, the mergers/acquisitions basically doubled the size of the carriers in one year: Delta

purchasing Western, Northwest buying Republic, USAir buying Piedmont and Continental acquiring People's Express and merging with Tower Air and Texas Airways. For the first three mergers we treated the planes in question as if they were part of the major airline from the beginning of the sample. For Continental, since each of the mergers alone was not especially large and since they purchased new planes in addition to the planes they acquired in that year, we treated the acquisitions as normal investment. Our results were not sensitive to the merger choices. As the results in the paper show, we calculated the same elasticities of substitution even excluding all the merging airlines.

The prices of the planes come from the *AVMARK Newsletter* from 1978-1994 and from *The Airliner Price Guide*, summer edition, 1994. The *Guide* gives the average sale price for every type of airplane sold in our sample from 1978-1995. The *Newsletter* gives estimates of the used value of a standard model of each type of aircraft in each year. Since some aircraft cease being produced in the middle of our sample and we do not know whether airlines are buying new or used planes, we factored all prices up to a "simulated new" price when new prices were not available. To do so we used the following procedure. Since 5 of our 16 planes are sold new in all 17 years of our sample, we fit a regression of the real price of these five planes on plane dummies and a time trend to get a real price trend for new aircraft (neglecting the quality change in the same model over time). We then apply this trend starting in the last year of the real price data for models which cease being produced part way through our sample. For the DC-8s and the 707s in our sample, there are no new prices observed but we do observe, for several planes, the new and used price for the same aircraft in the same year and averaging those, the estimated used value of a plane is about 40% of its new equivalent. In 1978 we scale up the used price of the DC-8s and 707s to a simulated new price and then apply the same downward trend in real prices discussed above. We do not use the decline in used prices as that also takes into account depreciation caused from using the plane. Calculating the prices in alternate ways did not change our results.

The data on the capital stock using accounting data from COMPUSTAT sets the capital stock to the book value of property, plant, and equipment in the year before the sample starts and then applies a perpetual inventory method with a depreciation rate of 10% over the sample. In other words, we assume that 10% of the capital stock depreciates per year and that capital expenditures are added to the capital stock in each year.

2. Bandwidth Selection for Kernel Regressions

To standardize the bandwidth selection on the different samples we use a simple plug-in rule which Silverman (1986) shows to be valid for density estimation and modify it for an Epanechnikov kernel. For a uni-modal distribution, the optimal bandwidth is approximately $b = 2.347 * s * n^{-1/5}$ where s is the standard deviation of the independent variable and n is the number of observations. To prevent over-smoothing due to outliers, we cap any percent changes to investment or hours at 250% for this calculation (this is not relevant for the ranges considered in the paper). For consistency, we only display figures in the range where the density is greater than or equal to 0.3. This keeps us from concluding something about the shape based on a very small number of observations. For the regression of the change in hours on desired hours (Figure 8) we used a slightly larger bandwidth to smooth out high frequency variation.

This does not change any results.

TABLE 1 : CAPITAL STOCK BY TYPE AT START OF SELECTED YEARS

	<u>1979</u>	<u>1982</u>	<u>1985</u>	<u>1988</u>	<u>1991</u>	<u>1994</u>
<u>American</u>						
707	76	0	0	0	0	0
727-100	57	53	37	39	40	0
727-200	79	55	125	125	125	114
737-200	0	0	0	21	12	0
737-300	0	0	0	8	8	0
747	11	4	0	2	2	0
757	0	0	0	0	26	75
767	0	0	10	29	45	64
DC10	28	28	51	60	59	57
MD-80	0	0	31	118	212	260
<u>United</u>						
727-100	88	50	50	49	24	0
727-200	65	104	104	104	104	79
737-200	59	49	49	89	74	69
737-300	0	0	0	0	105	158
747	18	9	18	26	39	51
757	0	0	0	0	24	88
767	0	0	19	19	19	42
DC10	37	31	53	56	54	52
DC8-5	31	26	11	0	0	0
DC8-6	39	23	30	29	19	0
<u>TWA</u>						
707	85	81	0	0	0	0
727-100	35	34	26	22	11	8
727-200	39	56	56	56	55	51
747	11	18	19	18	17	11
767	0	0	10	11	11	10
DC9-10	14	7	7	7	0	0
DC9-30	0	0	0	38	41	39
DC9-50	0	0	0	0	0	6
MD-80	0	0	0	15	30	33
L1011	30	32	33	33	32	20

Source: I.C.A.O., *Fleet Statistics 1978-1995*

Table 2: Estimates of Frictionless Capital Stock and Elasticity of Substitution

	(1)	(2)	(3)	(4) No mergers	(5)
Capital -->	Ending	Assigned	Ending	Ending	Hours
ln(COSTS)	-1.502 (.3688)	-1.842 (.3797)		-1.252 (.5278)	
ln(COSTS) with taxes			-1.319 (.3736)		
ln(COSTS) only operating					-1.324 (.2822)
<u>Dummies:</u>					
Plane Type	Yes	Yes	Yes	Yes	Yes
Airline- Year	Yes	Yes	Yes	Yes	Yes
n	815	790	815	428	815
R ²	.97	.97	.97	.97	.97

Notes: The dependent variable in columns (1)-(4) is the log of the capital stock for an airline-plane-year as defined at the top of the column and in the text. The dependent variable in (5) is the log of total hours flown by airline-plane-year. Standard errors are in parentheses. The ln(COSTS) variable is the log of aircraft operating costs plus the rental price of capital for one year, defined as the user cost in the text. The cost variable adds taxes in column (3) and the removes the rental price in column (5). Column (4) excludes firms that participated in a merger.

References

- Abel, A. and J. Eberly (1994), "A Unified Model of Investment Under Uncertainty." *American Economic Review* 84(5): 1369-84.
- (1996), "Optimal Investment with Costly Irreversibility," *Review of Economic Studies*, vol.63, no.4, pp.581-593, August 1996.
- (1999), "The Effects of Irreversibility and Uncertainty on Capital Accumulation," *Journal of Monetary Economics*, Vol. 44, no. 3 (December 1999): 339-377.
- Abrevaya, J. and W Jiang (1999), "A Simple Statistic for Testing Curvature," Mimeo, University of Chicago, GSB.
- Antoniou, Andreas (1991), "Economies of Scale in the Airline Industry: The Evidence Revisited," *Logistics and Transportation Review*; 27(2) June, pp. 159-84.
- Attanasio, O. (1995), "Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules," NBER Working Paper # 5282 (September).
- Barnett, Steven A. and Sakellaris, Plutarchos (1998), "Nonlinear Response of Firm Investment to Q: Testing a Model of Convex and Non-convex Adjustment Costs," *Journal of Monetary Economics* 42(2), pp. 261-88
- Berry, S., M. Carnal, and P. Spiller (1997), "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity." Working Paper, Yale University.
- Bertola, G. and R. Caballero (1990), "Kinked Adjustment and Aggregate Dynamics" in *NBER Macro Annual*, 5 O. Blanchard and S. Fischer, eds. MIT Press: Cambridge.
- (1994), "Irreversibility and Aggregate Investment." *Review of Economic Studies* 61(2): 223-246.
- Blackorby, and Schworm (1983), "Aggregating Heterogeneous Capital Goods in Adjustment-Cost Technologies." *Scandinavian Journal of Economics*, 85(2), 202-222.
- Borenstein, S. (1992), "The Evolution of U.S. Airline Competition." *Journal of Economic Perspectives* 6(2) 45-73.
- (1991), "The Dominant-Firm Advantage in Multi-Product Industries: Evidence from the U.S. Airlines." *Quarterly Journal of Economics*, November 1991, 106, 1237-1266..
- (1989), "Hubs and High Fares: Airport Dominance and Market Power in the U.S. Airline Industry." *Rand Journal of Economics*, Autumn 1989, 20, 344-365.
- Bresnahan, T. and V. Ramey (1994), "Output Fluctuations at the Plant Level." *Quarterly Journal of Economics*, 109(3), 593-624.
- Breuckner, J., N. Dyer, and P. Spiller (1992), "Fare Determination in the Airline Hub-and-Spoke Networks." *Rand Journal of Economics*, 23, 309-333.

- Breuckner and P. Spiller (1994), "Economics of Traffic Density in the Deregulated Airline Industry." *Journal of Law and Economics*, 37, 379-415.
- Caballero, R. (1993), "Durable Goods: An Explanation for Their Slow Adjustment," *Journal of Political Economy* 101(2) April 1993, pp. 351-84.
- Caballero, R. and E. Engel (1994), "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach." Working Paper, M.I.T.
- Caballero, R. and J. Leahy (1996), "Fixed Costs: The Demise of Marginal q ." NBER Working Paper # 5508.
- Caballero, R., E. Engel, and J. Haltiwanger (1995), "Plant Level Adjustment and Aggregate Investment Dynamics." *Brookings Papers on Economic Activity*, 1995:2, 1-54.
- (1997), "Aggregate Employment Dynamics: Building from Microeconomic Evidence." *American Economic Review*, 87(1), 115-37.
- Carey, Christopher (1997), "Pilots' Pay Gap Widens." *St. Louis Post-Dispatch*, April 13.
- Chirinko, R. (1993), "Multiple Capital Inputs, Q , and Investment Spending." *Journal of Economic Dynamics and Control*, 17(5-6), 907-928.
- Cooper, R., J. Haltiwanger, and L. Power (1995), "Machine Replacement and the Business Cycle: Lumps and Bumps." Mimeo, Boston University.
- Dixit, A. & R. Pindyck (1994), *Investment Under Uncertainty*. Princeton University, Princeton.
- Doms, M. and T. Dunne (1997), "Capital Adjustment Patterns in Manufacturing Plants." forthcoming, *Review of Economic Dynamics*.
- Eberly, J. (1994), "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases," *Journal of Political Economy*; 102(3): 403-36.
- Goolsbee, A. (1999), "Taxes and the Quality of Capital," Mimeo, University of Chicago, GSB.
- (1998), "Investment Tax Subsidies, Prices, and the Supply of Capital Goods." *Quarterly Journal of Economics*, 113(1), February 1998, pp. 121 - 148.
- Gross, D. (1997), "The Investment and Financing Decisions of Liquidity Constrained Firms." Working Paper, University of Chicago Graduate School of Business.
- Hamermesh, D. (1992), "General Model of Dynamic Labor Demand." *Review of Economics and Statistics* 74(4), 733-737.
- Hamermesh, D. and Gerard Pfann (1996), "Adjustment Costs in Factor Demand." *Journal of Economic Literature*, 34(3), 1264-92.
- Hardle, W. (1990), *Applied Nonparametric Regression*. Cambridge University Press; Cambridge.

- Hayashi, F. and D. Inoue (1991), "The Relation between Firm Growth and Q with Multiple Capital Goods: Theory and Evidence from Panel Data on Japanese Firms." *Econometrica*, 59(3), 731-53.
- Hubbard, R. G. (1997), "Capital-Market Imperfections and Investment." forthcoming, *Journal of Economic Literature*.
- Kumbhakar, Subal C. (1990), "A Reexamination of Returns to Scale, Density and Technical Progress in U.S. Airlines," *Southern Economic Journal*, 57(2), October, pp. 428-42.
- James, G. (1982), *Airline Economics*, Lexington Books, Lexington.
- Morrison, Steven and Clifford Winston (1986), *The Economic Effects of Airline Deregulation*. Brookings Institution, Washington D.C.
- O' Connor, W. (1989), *An Introduction to Airline Economics*, 4th edition. Praeger. New York.
- Shapiro, M. (1986), "Capital Utilization and Capital Accumulation: Theory and Evidence." *Journal of Applied Econometrics*, 1(3), 211-234.
- Silverman (1986), *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London.
- Wildasin (1984), "The Q-Theory of Investment with Many Capital Goods." *American Economic Review*, 74(1), 203-210.




Figure 1: Total Stock of Aircraft for Major Airlines

STATA™




Figure 2: Plane Effects by Aircraft Size

STATA

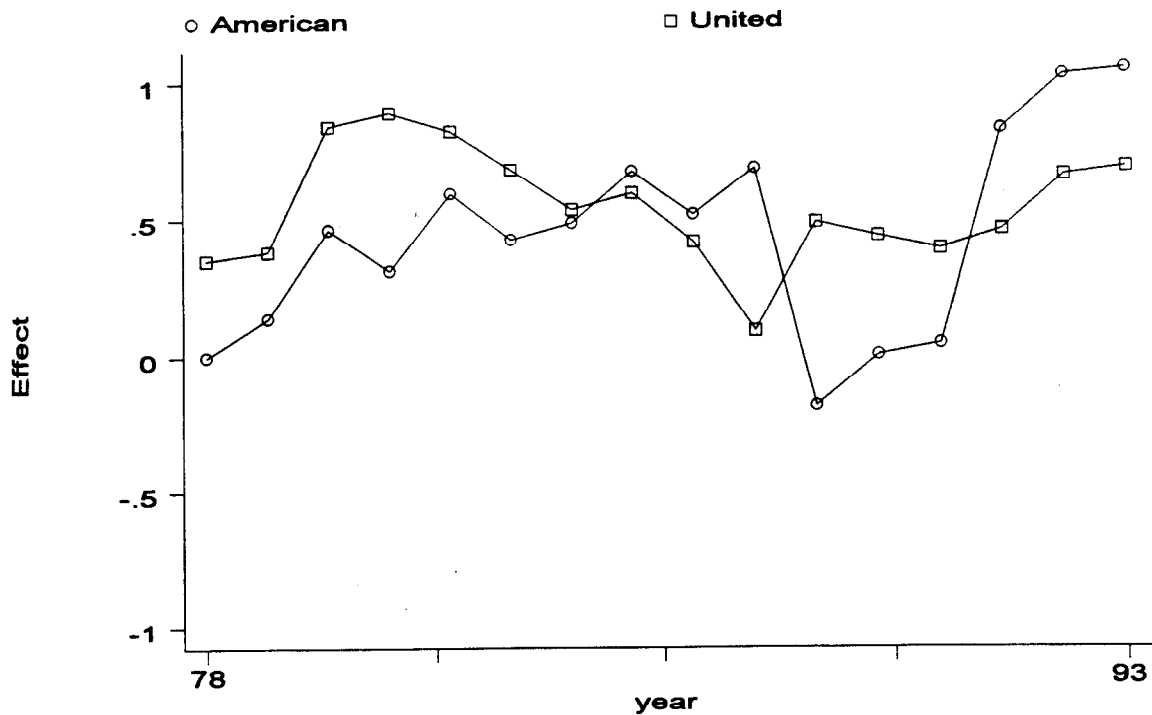


Figure 3A: Year Effects for Successful Airlines

STATA™

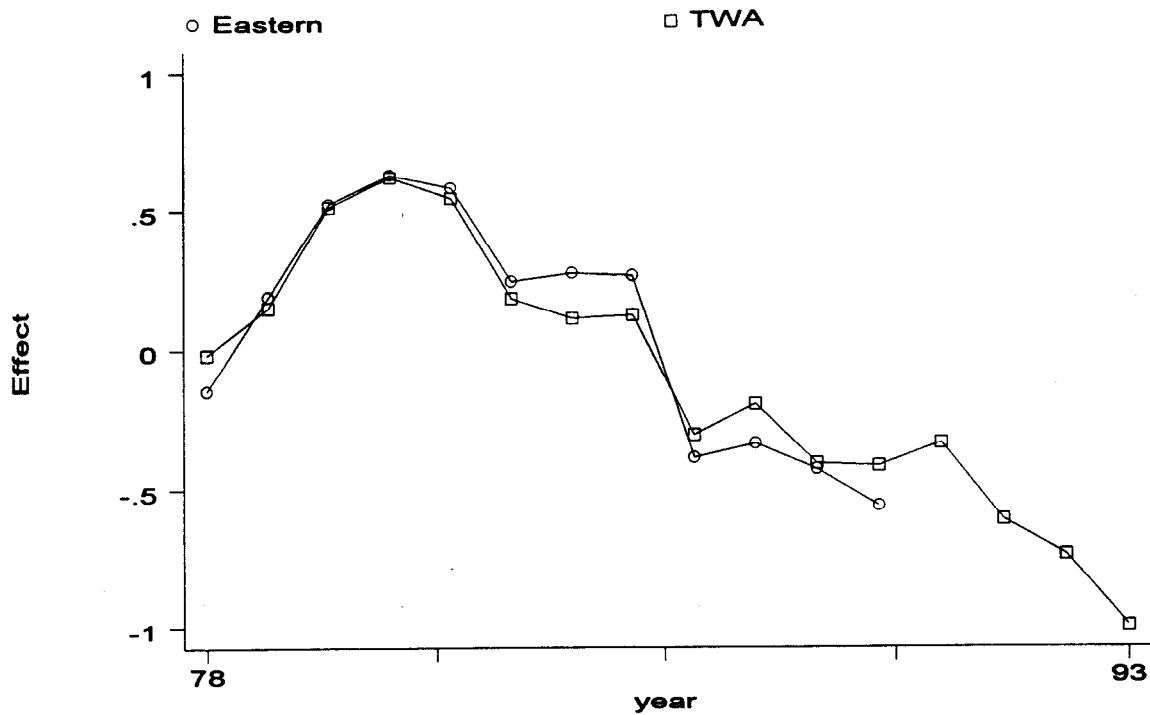



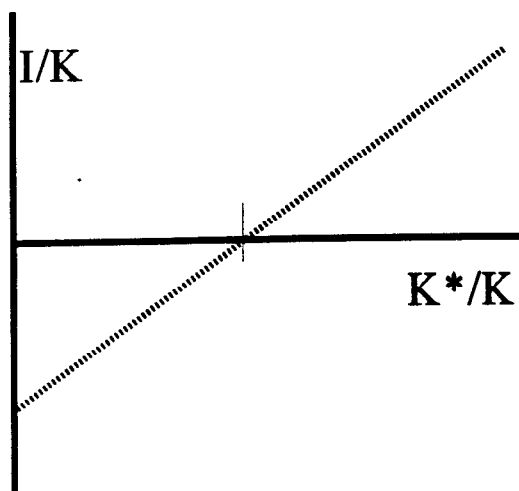
Figure 3B: Year Effects for Unsuccessful Airlines

STATA™

FIGURE 4:
Implied Shape for Various Adjustment Costs

Quadratic

Irreversibility



Non-Convexity

Price Wedge


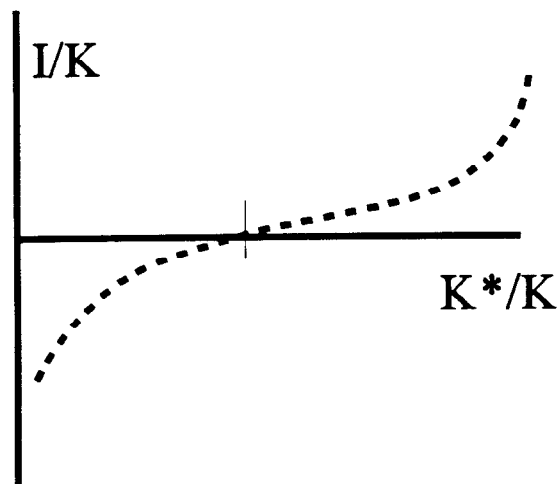


Figure 5: Net Investment




Figure 6:
Gross Investment and Retirements





Figure 7:
Net Investment with Plane-Year Effects






Figure 8:
Change in Hours

Figure 9:
Investment with Market Value of Capital




Figure 10:
Investment with Book Value of Capital




FIGURE II:
Net Investment




Figure 12:
Change in Hours and Net Investment

