International Journal of Industrial Organization 11 (1993) 407–435. North-Holland

Dynamic oligopoly behaviour in the airline industry

James A. Brander

Faculty of Commerce and Business Administration, University of British Columbia, Vancouver, B.C., V6T 1Z2, Canada

Anming Zhang*

Department of Economics, University of Victoria, Victoria, B.C., V8W 3P5, Canada

Final version received February 1993

This paper examines the dynamic interaction between United Airlines and American Airlines on a set of duopoly routes. We estimate quarterly 'conduct parameters' (or 'conjectural variations') and test constant behaviour Bertrand, Cournot, and collusive models. We also test two regime-switching models derived from the Green–Porter price war model, one based on reversion to Bertrand behaviour during price wars, the other on reversion to Cournot behaviour. The three constant behaviour models are all rejected. The regime-switching models can both describe the data but the Cournot-based version is preferred.

1. Introduction

In the aftermath of the U.S. Airline Deregulation Act of 1978, the U.S. airline industry has continued to be a focus of public attention. One concern is that the 'effective' number of firms has actually fallen [see Morrison and Winston (1990)], and most routes continue to be served by only one, two, or three airlines, suggesting that small numbers oligopoly is the dominant market structure in the industry. In addition, there appears to be considerable price instability. Taken together, these two observations suggest that dynamic strategic interaction could be an important aspect of airline behaviour. If so, the industry might offer an important opportunity for economists to apply empirical tests to the large theoretical literature on dynamic strategic environments. In addition, of course, competition in the

Correspondence to: J.A. Brander, Faculty of Commerce and Business Administration, University of British Columbia, Vancouver, B.C., V6T 1Z2, Canada.

*We thank two anonymous referees and the editor (Paul Geroski) for very helpful and perceptive comments. We also thank John Cragg, Murray Frank, Tae Oum, Serge Nadeau and Michael Tretheway. Financial support from the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

0167-7187/93/$06.00

airline industry is worth understanding for its own sake, especially given the debate going on in many countries over airline deregulation.

In this paper we examine the dynamic pattern of firm conduct using time-series and firm-specific data for a set of duopoly airline routes. Our main purpose is to infer from the data whether the dynamic pattern of market conduct can be well explained by leading models of oligopoly behaviour.

One possibility is that dynamic behaviour in the industry can be well approximated by the simple repetition of one-shot static equilibria of the Bertrand, Cournot, or cartel type. Alternatively, there may be a meaningful dynamic aspect to the interaction. One dynamic possibility described by Friedman (1971) is that firms might be able to use a grim trigger strategy to support the cartel outcome.

More flexible dynamic outcomes are provided by regime-switching models of the type examined by Green and Porter (1984). [See also Porter (1983a).] In regime-switching models, firms move between periods of 'punishment' (price wars) and periods of tacit collusion. Like Green and Porter (1984), we impose the idea that punishment phases are based on reversion to either Bertrand or Cournot one-shot outcomes. Our implementation is, however, more flexible than the strict Green–Porter model, as we do not require that price wars last a particular number of periods, but focus simply on whether regime switches (into and out of price wars) occur. We therefore refer to the models we test as RS (regime–switching) models.

Quite a few papers have investigated which static model best describes a 'snapshot' of a particular industry. Examples include Iwata (1974), Gollop and Roberts (1979), Appelbaum (1982), Geroski et al. (1985), Dixit (1988), and Brander and Zhang (1990), which examine the airline industry. Empirical analysis of dynamic oligopoly interaction is relatively rare. Porter (1983b) investigates cartel stability for the U.S. railroad industry during the 1880–1886 period. Some of the empirical work undertaken here follows the methodology of Porter (1983b). Hajivassiliou (1989) provides a valuable methodological commentary on Porter (1983b) and provides a reworking of Porter's data. Bresnahan (1987) analyzes changes of firm conduct in the mid-1950s for the U.S. automobile industry. These studies find evidence of switches of firm conduct between a Nash punishment regime and a collusive regime. Slade (1987, 1989) develops and tests an alternative to the Green–Porter model of price wars in which firms engage in price wars not to punish rivals but to learn about the economic environment in which they operate.

At the theoretical level, several alternative regime-switching models have been developed, including Abreu (1986), Rotemberg and Saloner (1986) and Harrington (1990a, b). The relationship between these papers and our work is described at various points in the text.

The airline industry has been the subject of extensive study and comment. Useful papers include Morrison and Winston (1987), Whinston and Collins

(1992), Berry (1992), and Hurdle et al. (1989), all of which examine empirically the 'contestability' of the airline markets. These studies show that deregulated airline markets are not perfectly contestable, and that there is a positive relationship between concentration and fares. Borenstein (1989, 1990) and Berry (1990) investigate relationships between hub–spoke route structures and airline pricing behaviour. They find strong evidence that an airline's domination of a hub results in higher average fares for routes from that hub. Thus dominant airlines are able to charge higher prices in their 'home' markets than their competitors. These pricing and contestability studies are, like ours, concerned with oligopoly conduct (at least implicitly). What distinguishes our paper is the explicit estimation of conduct parameters, explicit characterization and estimation of dynamic interaction, and the attempt to isolate the particular strategic setting of (roughly) symmetric duopoly. Other useful references on deregulated airline markets include Bailey et al. (1985), Levine (1987), and Keeler (1991).

The paper is organized as follows. Section 2 describes the theoretical structure underlying the paper. Section 3 discusses our empirical and statistical methodology; section 4 describes the data and data construction; and section 5 presents and interprets our main results. Model selection tests based on Bayes factors are described in section 6, and section 7 presents some sensitivity analysis. Section 8 contains qualifications and comments and section 9 contains concluding remarks.

2. Theoretical structure

Consider an industry consisting of two firms producing a homogeneous product. Let p_t be the market price of the product in time period t, and x_t^i be the corresponding output of firm i, $i=1,2$. Under certainty, the inverse demand function may be written as $p_t = p(X_t)$, where $X_t = x_t^1 + x_t^2$ is the aggregate output. Using $C_t^i = C^i(x_t^i)$ to denote costs for firm i in period t, the single-period profit function can be written as

$$\pi_t^i = x_t^i p(X_t) - C^i(x_t^i). \tag{1}$$

If we regard output, x_i, as the choice variable, and if we restrict attention to a single period, t, then the Nash equilibrium is represented by first-order conditions $\mathrm{d}\pi_t^i/\mathrm{d}x_t^i = p_t + x_t^i p' - c_t^i = 0$, where c_t^i represents marginal cost. This Nash quantity (or 'Cournot') solution might fail to hold for a variety of reasons, such as if price rather than quantity is the choice variable, or if the firms manage to collude and maximize combined profits. More generally, we can write

$$p_t + x_t^i p' - c_t^i = \delta_t^i, \tag{2}$$

where δ_t^i takes on the value 0 in the Cournot case. This equation is purely definitional and imposes no particular interpretation on δ. In general we can expect δ_t^i to vary with output levels, and possibly with other variables or parameters as well.

The conjectural variation parameter emerges by assuming that each firm views industry output, X, as a function of own output, x_i. The associated first-order condition is $\mathrm{d}\pi_t^i/\mathrm{d}x_t^i = p_t + x_t^i p'(1+v_t^i) - c_t^i = 0$, where v_t^i is interpreted as the covariation of the output of the other firm with own output, and is called the conjectural variation. We can rewrite this first-order condition as $p_t + x_t^i p' - c_t^i = -x_t^i p' v_t^i$, which corresponds to eq. (2) above with $\delta = -x_t^i p' v_t^i$. Rearrangement yields

$$v_t^i = \frac{(p(X_t) - c_t^i)}{p(X_t)} \frac{\eta_t}{s_t^i} - 1, \tag{3}$$

where $\eta_t = -(\mathrm{d}X_t/\mathrm{d}p_t)(p_t/X_t)$ is the (positive) elasticity of market demand, and s_t^i is the market share of firm i.

The conjectural variation terminology is somewhat controversial, but without entering into problematic interpretations, v is a convenient variable for parameterizing different theories of oligopoly. For symmetric duopoly with homogeneous products, values of -1, 0, and 1 for v represent Bertrand, Cournot, and cartel one-shot solutions, respectively. The higher the level of v, the greater the price–cost margin and hence the more apparently collusive is firm conduct. We refer to v simply as a 'conduct parameter'.

If firms undertake repeated interactions it seems natural for them to use contingent strategies, in which the price or output chosen in any period depends on the actions that have occurred in previous periods. Friedman's (1971) grim strategy described in the introduction implies, given a sufficiently low discount rate, that the (credible) threat of perpetual punishment would, under certainty, insure full collusion. Price wars would not actually occur. By introducing uncertainty about changes in market conditions, Green and Porter (1984) and Porter (1983a) develop a theory of oligopoly that can explain price wars. In their theory, firms cannot tell with certainty whether a fall in own profits has been caused by a rival's deviation from the implicit collusive output, or by worsened market conditions. Each firm starts by playing a collusive output level, but is prepared to punish the rival by playing the static non-cooperative level of output if price falls below some trigger level. Firms will play the punishment output level for some specified number of periods, then return to more (but not fully) collusive levels of output. The threat of punishment allows firms to maintain collusive outputs below the static non-cooperative output levels, but does not generally allow the full cartel outcome to be achieved.

Other models of price wars in dynamic games have also been developed.

Rotemberg and Saloner (1986) propose a model in which demand is uncertain and firms learn the demand realization before setting price. Given that demand shocks are independently and identically distributed (iid) it follows (somewhat counterintuitively) that price wars will tend to occur during booms. The basic idea is that if a firm is going to cheat, it is better off doing it during a high realization of demand or a 'boom', because the benefits of cheating are high relative to the costs of punishment anticipated in future periods.

We consider regime-switching (RS) models of the generalized Green–Porter type. Two different models are obtained by allowing the punishment phases to involve reversion to either Nash price (Bertrand) or Nash quantity (Cournot) one-shot solutions. In the quantity-based RS model we have $(v_n^1, v_n^2) = (0,0)$ in punishment periods, and $(0,0) < (v_c^1, v_c^2) < (1,1)$ in collusive periods. In the price-based version, we have $(v_n^1, v_n^2) = (-1,-1)$ in price war phases and $(-1,-1) < (v_c^1, v_c^2) < (1,1)$ in collusive phases.

In summary, the first three models of duopoly – Bertrand, Cournot, and full collusion – imply that v_t^i are constant over time and equal to -1, 0, and 1 respectively, whereas the Green–Porter and related RS models predict that there will be periodic switches of oligopoly conduct. Such models also predict that the degree of collusion in the collusive regime will be higher than in the punishment regime, but lower than the static cartel solution would imply. By estimating a time series of airline conduct parameters we can examine whether the data are consistent with the Bertrand, Cournot, cartel, or RS models.

3. Empirical and statistical methodology

We use a pooled cross-section and time-series data set of Chicago-based duopoly airline routes involving American Airlines and United Airlines. Dealing with the same duopolists on each route and using just Chicago-based routes corrects for many firm-specific and site-specific effects. It also alleviates the airline network effect associated with whether a route involves a hub city for a particular airline, as Chicago is a major hub for both American and United, while the connecting cities in our data are not significant hubs for either airline. Finally, working with the same set of routes over time increases comparability over different periods.[1]

In section 2 we considered the case of homogeneous products. In fact,

[1]Using Chicago-based routes dominated by American and United 'freezes' or controls for several potential confounding factors. American and United have a long stable competitive relationship on Chicago-based routes, suggesting some potential for observing long-run tacit collusion. One would expect different results if, for example, Fly-by-Night Airlines were competing with a large incumbent. The Chicago site itself has the advantage of being very large and very central. Also, staying with the same central hub equalizes the effects of climate, airport delays, and other site-specific factors.

airlines produce 'product lines' consisting of various fare classes, which are often categorized into first class, standard economy, and discount. Volume in the first-class category is very small and we exclude it from our analysis. Standard economy and discount categories are aggregated together and treated as a homogeneous output. It is possible to treat them separately, but the distinction between them is unreliable and not consistent across airlines, suggesting that aggregating the two categories is a better practice.[2] We use 'local traffic' as our measure of output. Local traffic for a city-pair consists of passengers who originate in one city and terminate in the other without changing planes (but allowing stops), and therefore excludes connecting passengers. We treat the outputs of American Airlines and United Airlines as homogeneous, noting that the two carriers are reasonably symmetric in factors (most notably, network) which might affect non-price product attributes. An airline's average fare is used as the product price for the airline.

The theoretical models of section 2 have implicitly assumed that firms make one choice of strategy variable per 'period'. In practice, it is not clear how long a decision period is. The shortest length permitted by the data is one quarter, as the principal data source for volume and price contains only quarterly information. We consider a quarter as our base case decision period, and use semi-annual data in sensitivity analysis.

Our basic 'experimental' assumption is that each airline route represents a similar strategic situation. Using k as the index for routes, eq. (3) becomes

$$v_{tk}^{i}=\frac{p(X_{tk})-c_{tk}^{i}}{p(X_{tk})}\,\frac{\eta_{tk}}{s_{tk}^{i}}-1. \tag{4}$$

We then calculate the route-specific conduct parameters, v_{tk}^{i}, for American and United. Note that v_{tk}^{i} as described above refers to the conduct parameter in pure duopoly. We have tried to restrict attention to airline routes where American and United did have something close to pure duopoly, but there are some marginal carriers on some of the routes. The presence of other carriers can create difficulties in the calculation and interpretation of results.

No problems are created if the output of additional firms is treated simply as an exogenous subtraction from demand that is independent of the actions of the two dominant firms. If, however, the output of other firms is understood by the two market leaders to depend on what they do, then the presence and size of marginal firms can change the reference values of conduct parameter v expressing the interaction betweeen the two market

[2]This is reinforced by the fact that 'price wars' consist largely of releasing a much higher fraction of seats into deep discount categories.

leaders.[3] Thus the presence of extra firms with endogenous output implies that the formula given by (4) is no longer exactly correct, but should be adjusted in some way for the presence of extra firms. There is no obvious way of making an appropriate adjustment. One partial solution, however, is to 'correct' for the presence of additional firms in the calculation of conduct parameters by including the market share of additional firms as an explanatory variable for the route-specific conduct parameters.

Section 2 set out the basic conjectural variation algebra for the case of certainty. However, the Green–Porter and other RS models require uncertainty. In addition, our econometric implementation also requires uncertainty. We impose uncertainty in the following way:

$$v_{tk}^i = v_t^i + \rho_i r_{tk} + \varepsilon_{tk}^i, \tag{5}$$

where r_{tk} denotes the market share of residual carriers and ρ^i is a parameter to be estimated. Thus the route-specific conduct parameter for airline i in period t is thought of as the sum of a general conduct parameter reflecting overall competitive interaction for airline i in period t, a route-specific effect reflecting the market share of marginal carriers, and an error.

Eq. (5) could be rewritten as $v_{tk}^i - \rho^i r_{tk} = v_t^i + \varepsilon_{tk}^i$, where v_t^i is the expected value of the adjusted conduct parameter for firm i in period t, and ε_{tk}^i is a random iid deviation of the 'observed' conduct parameter from its expected value (with mean zero).

We have chosen this point to insert the specification of random error into our estimation problem. A theoretically preferable alternative would be to derive the structure of the estimation error from an explicit specification of demand uncertainty. However, we emphasize that the specification given in (5) is consistent with some specification of demand (and cost) uncertainty, although not with demand uncertainty being additive in the inverse demand function.[4] It is, of course, somewhat arbitrary as to where simple additive iid shocks are imposed in the system. As a practical matter, imposing them at this stage is required to make our actual estimation problem tractable. In addition, tests of residuals indicate that this specification of the error appears reasonable. We take the view that despite our imposition of uncertainty at this stage in the model development, demand shocks should be viewed as one of the major sources of uncertainty underlying this error term. We can

[3]For example, in the competitive fringe model with homogeneous products, cartel behaviour among the two market leaders implies a conduct parameter that exceeds one, while Bertrand behaviour implies a conduct parameter that exceeds -1. The Cournot value of v remains at zero, irrespective of the presence of other firms.

[4]More specifically, if inverse demand can be written as $p = p^*(X) + xp'\varepsilon_d^i$ and if marginal cost uncertainty can be written as $c = c^* + xp'\varepsilon_c^i$, where p^* and c^* are the continuing or certain components of demand and cost, and ε_d and ε_c are iid shocks to demand and cost, then the error in eq. (4) becomes additive iid as assumed here.

also interpret the error as possibly arising from our inability to observe prices exactly. This helps accommodate our assumption of homogeneous products with the observation that prices for the two airlines, while close, are not identical.

Let v_t denote $(v_t^1, v_t^2)'$. Consider general situations where v_t shifts between two values, $v_c \equiv (v_c^1, v_c^2)'$ and $v_n \equiv (v_n^1, v_n^2)'$, representing collusive periods and non-collusive periods. Suppose D_t is an indicator variable which equals one when the market is in a collusive phase and zero when the market is non-collusive, yielding

$$v_t = v_\text{n} + (v_\text{c} - v_\text{n}) D_t. \tag{6}$$

Note that $v_t = v_\text{c}$ if $D_t = 1$, and $v_t = v_\text{n}$ if $D_t = 0$. Using eq. (5), we convert (6) to the stochastic specification:

$$v_{tk} = v_\text{n} + (v_\text{c} - v_\text{n}) D_t + \rho r_{tk} + \varepsilon_{tk}, \tag{7}$$

where $v_{tk} \equiv (v_{tk}^1, v_{tk}^2)'$, $\rho \equiv (\rho^1, \rho^2)'$, $\varepsilon_{tk} \equiv (\varepsilon_{tk}^1, \varepsilon_{tk}^2)'$, $k = 1, 2, \dots, K$, and $t = 1, 2, \dots, T$. The Cournot model arises when $D_t \equiv 0$ and $v_\text{n} = (0, 0)'$ [hence, $v_t = (0, 0)'$ for all t]. The Bertrand model and the stationary collusive model arise, under identical costs, when $D_t \equiv 0$ and $v_\text{n} = (-1, -1)'$ and when $D_t \equiv 1$ and $v_\text{c} = (1, 1)'$, respectively.

For a RS model of the Green–Porter type, we would have $v_\text{n} = (-1, -1)'$ and $v_\text{c} = (v_p^1, v_p^2)'$ with $-1 < v_p^i < 1$ if punishment occurs at the one-shot Nash price solution; and $v_\text{n} = (0, 0)'$ and $v_\text{c} = (v_q^1, v_q^2)'$ with $0 < v_q^i < 1$ if punishment occurs at the one-shot Nash quantity solution. In both cases, if the D_t sequence, $\{D_t\}$, is known for the sample period, then v_c can be estimated using a linear regression with D_t as a dummy variable. In fact, however, $\{D_t\}$ is unknown. We assume, as in Porter (1983b), that D_t is Bernoulli distributed: $D_t = 1$ with probability λ, and $D_t = 0$ with probability $1 - \lambda$. We also assume that ε_{tk} is identically and independently distributed $\text{N}(\boldsymbol{0}_2, \boldsymbol{\Sigma})$, where $\text{N}(\cdot, \cdot)$ stands for the normal distribution, $\boldsymbol{0}_2$ is a 2×1 vector of zeros and $\boldsymbol{\Sigma}$ is the (unknown) 2×2 covariance matrix.

While it is not obvious that ε_{tk} should be normal iid (niid), there are some plausible reasons for this specification. Since v_{tk}^i is a standardized variable that is not affected by market size, it seems more reasonable that its errors would be iid than is the case for, say, price or quantity, where we would expect errors to be heteroscedastic. Assuming normality is neither unreasonable nor particularly compelling, but it greatly simplifies an otherwise very complicated and possibly infeasible estimation procedure. We tested for normality, time dependence, and heteroscedasticity of the residuals and were unable to reject the niid null hypothesis. The probability density function of v_{tk} is

$$f(v_{tk}, r_{tk} \mid \Omega) = \lambda f^{c}_{tk} + (1-\lambda) f^{n}_{tk}, \tag{8}$$

where $\Omega \equiv (\lambda, v_c, \rho, \Sigma)$ denotes the parameters to be estimated;

$$f^{c}_{tk} \equiv f^{c}(v_{tk}, r_{tk} \mid D_t = 1) = (2\pi)^{-1} |\Sigma|^{-1/2} \times \exp\{-\tfrac{1}{2}(v_{tk} - v_c - \rho r_{tk})' \Sigma^{-1} (v_{tk} - v_c - \rho r_{tk})\}, \tag{9}$$

$$f^{n}_{tk} \equiv f^{n}(v_{tk}, r_{tk} \mid D_t = 0) = (2\pi)^{-1} |\Sigma|^{-1/2} \times \exp\{-\tfrac{1}{2}(v_{tk} - v_n - \rho r_{tk})' \Sigma^{-1} (v_{tk} - v_n - \rho r_{tk})\}. \tag{10}$$

The associated likelihood function is given by

$$l(\Omega \mid v_{tk}, r_{tk}) = \prod_{t,k} f(v_{tk}, r_{tk} \mid \Omega). \tag{11}$$

Given v_{tk} and r_{tk}, estimates of $(\lambda, v_c, \rho, \Sigma)$ can be obtained by maximizing $l(\Omega \mid v_{tk}, r_{tk})$ with respect to Ω.

The estimation problem is similar to the problem of mixture distributions studied by, for example, Day (1969) in which a sample is generated from two regimes for each time period. Define $w_{tk} \equiv \Pr(D_t = 1 \mid v_{tk}, r_{tk})$, i.e. w_{tk} is the probability of $D_t = 1$ given observations v_{tk} and r_{tk}. We use the following algorithm to estimate $\Omega = (\lambda, v_c, \rho, \Sigma)$ and w_{tk}. Given an initial estimate of (λ, v_c, ρ), namely, $(\lambda^{(0)}, v^{(0)}_c, \rho^{(0)})$, an estimate of w_{tk} can be obtained using

$$w^{(1)}_{tk} = (1 + \exp\{\beta'(v_{tk} - \rho^{(0)} r_{tk}) + \gamma\})^{-1}, \tag{12}$$

where

$$\beta = \frac{S_N^{-1}(v_n - v^{(0)}_c)}{1 - \lambda^{(0)}(1-\lambda^{(0)})(v_n - v^{(0)}_c)' S_N^{-1}(v_n - v^{(0)}_c)}, \tag{13}$$

$$\gamma = \ln\left(\frac{1-\lambda^{(0)}}{\lambda^{(0)}}\right) - \tfrac{1}{2}\beta'(v_n + v^{(0)}_c), \tag{14}$$

$S_N = (1/N)\sum_{t,k}(v_{tk} - \bar{m})(v_{tk} - \bar{m})'$ is the overall sample covariance matrix, $N = TK$, and $\bar{m} = (1/N)\sum_{t,k} v_{tk}$ is the overall sample mean vector. We then update $(\lambda, v_c, \rho, \Sigma)$ by using

$$\lambda^{(1)} = \frac{\sum_{t,k} w^{(1)}_{tk}}{N}, \tag{15}$$

$$\rho^{(1)} = \frac{N\lambda^{(1)}\sum_{t,k} v_{tk} r_{tk} - (\sum_{t,k} w^{(1)}_{tk} r_{tk})(\sum_{t,k} w^{(1)}_{tk} v_{tk}) - N\lambda^{(1)}(\sum_{t,k}(1 - w^{(1)}_{tk}) r_{tk}) v_n}{N\lambda^{(1)}\sum_{t,k} r^2_{tk} - (\sum_{t,k} w^{(1)}_{tk} r_{tk})^2}, \tag{16}$$

$$v_{\mathrm{c}}^{(1)}=\frac{\sum_{t,k}w_{tk}^{(1)}(v_{tk}-\rho^{(1)}r_{tk})}{\sum_{t,k}w_{tk}^{(1)}}, \tag{17}$$

$$\Sigma^{(1)}=\frac{1}{N}\sum_{t,k}(w_{tk}^{(1)}(v_{tk}-v_{\mathrm{c}}^{(1)}-\rho^{(1)}r_{tk})(v_{tk}-v_{\mathrm{c}}^{(1)}-\rho^{(1)}r_{tk})'$$

$$+(1-w_{tk}^{(1)})(v_{tk}-v_{\mathrm{n}}-\rho^{(1)}r_{tk})(v_{tk}-v_{\mathrm{n}}-\rho^{(1)}r_{tk})'). \tag{18}$$

Once $(\lambda^{(1)}, v_{\mathrm{c}}^{(1)}, \rho^{(1)})$ is obtained, new estimates of w_{tk} and Ω can be obtained using eqs. (12)–(18). Denoting the new estimates by $w_{tk}^{(2)}$ and by $\Omega^{(2)}=(\lambda^{(2)}, v_{\mathrm{c}}^{(2)}, \rho^{(2)}, \Sigma^{(2)})$, we continue this iterative procedure until convergence occurs. Suppose it converges at the nth iteration. Denote $\hat{w}_{tk}=w_{tk}^{(n)}$ and $\hat{\Omega}=\Omega^{(n)}$. It can be shown [see, for example, Day (1969)] that $\hat{\Omega}$ will be a solution of the maximum likelihood equations. There may, however, be more than one solution, so we repeat the iterative procedure from a number of different starting points. In the case of multiple local maxima, we choose the estmates at which the value of the likelihood function is highest.

After estimation is completed, we view each $\hat{w}_{tk}$ as an observation which, for a given period, consists of a common element and an error:

$$\hat{w}_{tk}=w_t+e_{tk}, \tag{19}$$

where w_t is the expected value of $\hat{w}_{tk}$ in period t, and e_{tk} is a random deviation of $\hat{w}_{tk}$ from its expected value (with mean zero). A period is classified as collusive if $w_t \geqq 0.5$, and as non-collusive otherwise.[5] We estimate the mean and standard deviation of the regime classification parameter w_t, and infer which periods might be considered as collusive or non-collusive.

In a finite set of time periods, we might of course fail to observe regime switches, even if the firms are in a Green–Porter environment. If, for example, firms were engaged in a price war thoughout the sample period, then we would be unable to distinguish between the Cournot model and the RS–quantity model or, correspondingly, between the Bertrand model and the RS–price model. If, however, we did observe static Cournot behaviour throughout the sample, this would certainly increase our confidence in the Cournot model.

[5]Each observation, v_{tk}, may be allocated to either the collusive regime or the competitive regime (but not both). A sensible classfication rule could be determined by minimizing the expected cost of misclassification (ECM). Using Fisher's linear discriminant function, one can show that the classification rule that minimizes the ECM is given by: allocate v_{tk} to the collusive regime if and only if $\hat{w}_{tk} \geqq 0.5$.

Table 1

Route, distance, fares and volume.

Route	Distance	AA fare	UA fare	Industry passengers	AA&UA share
Grand Rapids	134	112	116	1,448	0.92
Buffalo	467	155	146	2,130	0.99
Rochester	522	169	168	1,873	0.99
Wichita	591	173	172	964	0.99
Syracuse	601	159	159	1,435	0.99
Oklahoma	692	164	167	1,566	0.78
Albany	717	172	174	1,425	0.99
Hartford	778	181	185	3,930	0.99
Providence	842	162	169	1,576	0.99
Austin	972	154	153	1,130	0.83
Tucson	1,441	154	149	1,032	0.95
Las Vegas	1,521	146	122	5,907	0.67
Reno	1,680	164	146	753	0.95
Ontario, CA	1,707	190	186	1,526	0.95
Sacramento	1,790	200	195	875	0.97
San Jose	1,837	231	229	1,395	0.99

4. The data

To calculate the route-specfic conduct parameters using eq. (5), we need data on price, market share, marginal cost, and demand elasticity, all of which are route-, carrier-, and quarter-specific. In addition, we need data on industry output to compute the residual carriers' market share. We obtained quantity and price data from I.P. Sharp Associates. The I.P. Sharp data derive from Databank 1A of the U.S. Department of Transportation (DOT) *Origin and Destination Survey*. This data set, OD1A, is a 10% sample of all tickets that originate in the United States on significant domestic carriers. The basic unit of our price data is the directional fare, which is either a one-way fare or half of the round-trip (or excursion) fare. Our volume measure is the associated number of directional passengers in the 10% sample for each airline on each route. Estimated total passengers for each route would therefore be about 10 times the numbers reported in our tables.

We took all Chicago-based city pairs on which American and United Airlines had a combined market share of over 90% in 1985, and on which each carrier had a market share of at least 33% and at least 100 passengers per quarter in the sample. This left us with 16 routes for the period beginning in the fourth quarter of 1984 and ending in the fourth quarter of 1988. The data from the first and third quarters of 1988 were not in usable form. Our data set thus contains the same 16 Chicago-based city-pair routes for each quarter from the fourth quarter of 1984 to 1987, and for the second and fourth quarters of 1988.

Table 1 lists the name of each connecting city and distance (in miles) of

Table 2

Mean prices and volumes, quarterly.

Quarter	AA fare	UA fare	AA passenger	UA passenger	AA & UA share
4Q84	177	175	738	881	0.96
1Q85	175	160	666	891	0.96
2Q85	160	136	1,065	851	0.96
3Q85	158	158	944	962	0.97
4Q85	160	162	769	883	0.97
1Q86	158	162	1,208	832	0.97
2Q86	157	155	813	968	0.95
3Q86	165	157	673	972	0.92
4Q86	164	163	627	879	0.93
1Q87	157	147	620	956	0.92
2Q87	151	149	823	904	0.92
3Q87	179	185	651	749	0.88
4Q87	195	197	613	599	0.88
2Q88	182	190	897	566	0.90
4Q88	180	172	815	634	0.90
Overall average	168	165	792	835	0.93

each route, in ascending order of distance. It also displays the mean prices in U.S. dollars (over the 15 sample quarters) for American Airlines and United Airlines on each route. Industry volume (for the 10% sample) is given in column 5. The last column gives the combined mean market shares of the two airlines on each route.

Table 2 displays the mean values of price and volume for each quarter (the average now is taken over the routes). Average fares were roughly equal for the two carriers, and fares did fluctuate noticeably over the sample period. The prices were on average relatively low in the middle quarters. The combined average market shares of the two carriers were high and fairly stable, ranging from 88% to 97% over the sample quarters.[6]

Determining marginal costs raises serious conceptual issues. Consider the following linear specification of the total cost for an airline on a city-pair route market [as used by Schmalensee (1977) and Panzar (1979), among others]:

$$C^i(x^i_{tk}) = a^i_{tk} x^i_{tk} + b^i_{tk} g_i(x^i_{tk}) + F^i_{tk}, \tag{20}$$

where a^i_{tk} is the (direct) cost of serving a passenger, g_i is the number of flights, b^i_{tk} is the cost of a flight, and F^i_{tk} is the (firm-specific) fixed cost. The first two components of C^i form the variable cost.

[6]We note that the market might not be in equilibrium for 2Q85 as a result of a pilots strike against United Airlines. The strike, starting 17 May, prevented United from operating its full schedules for 29 days. Immediately after the strike ended, United began rebuilding its strike-reduced operations by offering a series of fare promotions. This may account for United's lowest average fare in 2Q85, accompanied by American's high passsenger volume.

Table 3

Costs per passenger-mile (cents).

	American	United
1984	13.14	11.89
1985	12.33	12.61
1986	11.27	11.00
1987	11.89	11.42
1988	12.07	11.57

The marginal cost of carrying a passenger is then

$$c_{tk}^{i}=c^{i}(x_{tk}^{i})=a_{tk}^{i}+b_{tk}^{i}g_{i}'(x_{tk}^{i}). \tag{21}$$

For carrying a standby passenger, one would expect $g_i'=0$ and hence $c_{tk}^{i}=a_{tk}^{i}$. On the other hand, if carrying an additional passenger leads the airline to consider adding one more flight, items such as flight operating cost (e.g. crew, fuel, landing fee), aircraft maintenance, and depreciation must be part of marginal cost. In general, the marginal cost of serving a representative passenger would be the direct cost of passenger service plus a share of cost incurred to make flights and seats available. We therefore use the average variable cost to approximate the marginal cost.

One proxy for variable cost is operating expenses, as reported by the U.S. DOT based on its Form 41 (see *Air Carrier Financial Statistics*). The DOT also reports revenue passenger miles (i.e. total miles flown by passengers) (see *Air Carrier Traffic Statistics*). We divide operating costs by passenger-miles to obtain a measure of variable cost per passenger-mile. Table 3 provides the yearly operating costs per passenger-mile (cpm_t^i) for the two airlines.

The cost per passenger-mile in table 3 must be converted into route-specific costs per passenger. Cost per passenger varies with distance, but not linearly: cost per passenger is increasing but concave in distance (e.g. cost per passenger on a 500 mile flight is less than twice the cost per passenger on a 250 mile flight). To obtain the route-specific cost per passenger, we use each carrier's average flight length for the U.S. market as as whole, denoted AFL_t^i, and the elasticity of cpm^i with respect to flight distance. We obtain AFL_t^i by dividing aircraft revenue miles by aircraft revenue departures performed for each year. (Average flight lengths are very similar for the two airlines, with overall means of 769 and 765 miles for American and United, respectively.) The route-specific marginal cost is then calculated as follows:

$$c_{tk}^{i}=cpm_{t}^{i}\left(\frac{DIS_{k}}{AFL_{t}^{i}}\right)^{-\theta}DIS_{k}, \tag{22}$$

where DIS_k is the distance of route k, and θ is the (positive) elasticity of cpm^i

with respect to distance.[7] Several studies in the airline literature suggest a value of about 0.5 for θ [see Brander and Zhang (1990)]. We take $\theta=0.5$ as our base case value for θ, and conduct sensitivity analyses using other elasticity values.

To complete our data collection, we need to know η_{tk}, the elasticity of market demand for route k in period t. Oum et al. (1986) estimated a two-stage consumer demand system for U.S. domestic air travel routes, using a cross-sectional sample of 200 routes in 1978. A translog demand system was employed, and fareclasses were sub-divided into first-class, standard economy, and discount. Their estimated elasticities for the discount fareclass range from 1.5 to 2.0, and the range for regular economy is 1.2 to 1.4 suggesting an overall average of 1.6. More recently, Oum et al. (1993) estimated elasticities for U.S. airlines over the 1981–1988 period and obtained an average of 1.6 for the combined economy and discount category. In our base case analysis, we use a demand elasticity of 1.6 to approximate η_{tk}, and conduct some sensitivity analyses using demand elasticities of 1.2 and 2.0.

5. Results

We first calculate[8] route-specific marginal costs for each airline and each quarter (c^i_{tk}). The two carriers had very similar costs overall, so we therefore continue to treat -1 and 1 as the Bertrand and cartel solutions. The calculated route-specific marginal costs are then used to construct $\{v^i_{tk}\}$. To see whether the data are consistent with the Cournot, Bertrand, or cartel models, we estimate the statistical model given by (5) for each airline. The estimated mean conduct parameters for each quarter (v^i_t) are shown in table 4.[9] We also estimate a model of the form

$$v^i_{tk}=v^i+\rho^i r_{tk}+\varepsilon^i_{tk} \tag{23}$$

for each airline, and the estimates of the overall conduct parameters (v^i) are given in the last row of the same table.[10]

The last row of table 4 shows that the overall mean conduct parameters of the two airlines are remarkably close to each other, and are between the

[7]It is argued in Bernardino and Blinder (1991) that airline cost data can be improved by adjusting for aircraft type. We would argue that trying to adjust costs for aircraft type is unreliable and may well add error.

[8]All calculations, including the mixture distribution estimation, were programmed in version 2.0 of GAUSS.

[9]The estimates for ρ^i are -2.16 for American and -1.31 for United, with standard errors of 0.33 and 0.46, respectively.

[10]The estimates for ρ^i in this case are -1.79 for American and -0.91 for United, with standard errors of 0.33 and 0.47, respectively.

Table 4

Estimated mean conduct parameters, Quarterly: base case.

	American		United	
Quarter	Mean	Standard error	Mean	Standard error
4Q84	0.35	0.15	0.47	0.21
1Q85	0.59	0.15	0.19	0.21
2Q85	0.04	0.15	0.06	0.21
3Q85	0.09	0.15	0.14	0.21
4Q85	0.18	0.15	0.20	0.21
1Q86	0.10	0.15	0.85	0.21
2Q86	0.47	0.15	0.37	0.21
3Q86	0.76	0.15	0.30	0.21
4Q86	0.65	0.15	0.38	0.21
1Q87	0.64	0.15	0.03	0.21
2Q87	0.28	0.15	0.23	0.21
3Q87	0.67	0.16	0.76	0.22
4Q87	0.77	0.16	0.87	0.22
2Q88	0.50	0.15	1.35	0.21
4Q88	0.55	0.15	0.87	0.21
Mean	0.42	0.05	0.44	0.06

Cournot value of 0 and the cartel value of 1, but somewhat closer to the Cournot value. Classical hypothesis tests would strongly reject the Bertrand value of -1, and would also reject both the cartel value and the Cournot value. Thus, the base case data appear to be inconsistent with the Bertrand, Cournot, and cartel models.

Table 4 also reveals that Bertrand behaviour is (statistically) inconsistent with the base case data in every sample quarter. This would imply, under the price version of the RS model, that every sample quarter would be classified as collusive. Regime classification under the RS–quantity model, however, is not clear by looking at table 4.

We now report the estimation results based on the RS model. Table 5 displays the maximum likelihood estimates of the two versions of the RS model, with estimated standard errors in parentheses.[11] The estimates of the RS–price model are displayed in columns 2 and 3. As expected, the estimated regime probability, λ, is one that would imply that the carriers behaved collusively in every sample quarter. The estimates of (v_p^1, v_p^2) are 0.42 for American Airlines and 0.44 for United Airlines, with standard errors of 0.06 and 0.08, respectively. The estimates are consistent with the prediction of the Green–Porter model that $(-1,-1)<(v_p^1, v_p^2)<(1,1)$, i.e. collusive period

[11]The standard errors in table 5 are obtained by computing the inverse of the Fisher information matrix evaluated at the corresponding estimates. The latter is a consistent estimate of the asymptotic covariance matrix of the parameter estimates.

Table 5

Estimation results of the regime-switching model: Base case.

	RS–price		RS–quantity	
Variable	American	United	American	United
v_n	−1	−1	0	0
v_c	0.42 (0.06)	0.44 (0.08)	0.68 (0.08)	0.75 (0.10)
ρ	−1.79 (0.74)	−0.91 (0.76)	−1.62 (0.30)	−0.84 (0.43)
Σ	0.40 (0.07)	0.18 (0.09)	0.29 (0.04)	0.04 (0.04)
	0.18 (0.09)	0.79 (0.13)	0.04 (0.04)	0.63 (0.07)
λ	1.00 (0.12)		−0.58 (0.07)	
Log-likelihood	−530.01		−528.48	

prices exceed those implied by competitive price-setting, but are less than those consistent with static cartel solutions.

The estimates of the RS–quantity model are displayed in columns 4 and 5 of table 5. The estimates of (v_q^1, v_q^2) are 0.68 for American and 0.75 for United, with standard errors slightly larger than those based on the RS–price model. The result is roughly consistent with the theoretical prediction that $(0,0)<(v_q^1, v_q^2)<(1,1)$, i.e. collusive period outputs are less than those implied by Cournot behaviour, but exceed those that maximize single-period expected joint profits.

The overall estimated degree of collusion implied by the RS model is given by $\hat{\lambda}\hat{v}_c+(1-\hat{\lambda})v_n$ [using (6)]. From table 5, $\hat{\lambda}\hat{v}_c+(1-\hat{\lambda})v_n=(0.42,0.44)'$ under the RS–price model, and $(0.40,0.43)'$ under the RS–quantity model. The estimates are virtually identical for the two models, and are almost identical to the sample means. Estimation of the two models, however, yields different log-likelihood function values: the RS–quantity value is greater than the RS–price value. The two models give quite different estimates of regime probability λ.

To see which quarters might be classified as collusive, we estimate, using eq. (19), the mean and standard deviation of the regime classification parameter, w_t. The estimation shows clearly that under the RS–price model, each of our sample quarters should be classified as collusive. Essentially, the firms were never sufficiently close to marginal cost pricing on these routes to conclude that we were observing Bertrand-style punishment. This is not the case for the quantity version of the RS model, however. The estimated mean regime classification parameters for the RS–quantity model, along with the standard error and 95% confidence interval for each mean, are reported in table 6.

The table shows that under the RS–quantity model, the estimated mean regime classification parameters (w_t) are less than 0.5 for 2Q85, 3Q85, 4Q85,

Table 6

Estimated mean regime classification parameters w_t: Base case.

	RS–quantity		
Quarter	Mean	Standard error	95% confidence interval
4Q84	0.530	0.061	(0.401, 0.660)
1Q85	0.589	0.059	(0.463, 0.715)
2Q85	0.348	0.060	(0.220, 0.476)
3Q85	0.421	0.072	(0.269, 0.574)
4Q85	0.456	0.072	(0.303, 0.609)
1Q86	0.553	0.084	(0.375, 0.732)
2Q86	0.576	0.074	(0.419, 0.733)
3Q86	0.666	0.068	(0.522, 0.811)
4Q86	0.643	0.076	(0.483, 0.804)
1Q87	0.595	0.082	(0.422, 0.768)
2Q87	0.494	0.081	(0.323, 0.665)
3Q87	0.698	0.081	(0.527, 0.870)
4Q87	0.730	0.076	(0.570, 0.891)
2Q88	0.708	0.079	(0.539, 0.876)
4Q88	0.672	0.078	(0.506, 0.838)

and 2Q87. Consequently, these quarters are classified as Cournot punishment periods.[12]

Finally, we test explicitly for behavioural switches by the airlines during the sample period. Two null hypotheses are tested. One hypothesis is that only non-collusive behaviour is observed in the data. This means, for the RS–price model, $v_c = v_n = (-1, -1)'$, i.e. the airlines behaved in Bertrand fashion, whereas for the RS–quantity model, $v_c = v_n = (0, 0)'$. Likelihood ratio tests are employed. Specifically, for a given model (RS–price or RS–quantity), let L_0 be the (maximized) value of the log-likelihood function under the null hypothesis, and L_1 be the value of the log-likelihood function under the mixture distribution estimation (given in the bottom row of table 5). Then, $2(L_1 - L_0) \sim \chi_3^2$, where χ_3^2 is a chi-squared distribution with three degrees of freedom. The likelihood ratios for testing this null hypothesis, $2(L_1 - L_0)$, are reported in the second row of table 7.

The table shows that the null hypothesis that only non-collusive behaviour is observed is rejected against both the RS–price and the RS–quantity models. (The critical value at level of significance $\alpha = 0.01$ is $\chi_3^2(0.01) = 11.34$.) The rejection of Bertand behaviour is overwhelming, whereas Cournot behaviour is rejected to a much lesser, but still highly significant, degree.

[12]Our construction of 95% confidence intervals has implicitly assumed that for each period, observations $\hat{w}_{tk}$ are normally distributed (with mean w_t). With 16 observations, this might result in the width of confidence intervals being underestimated. We calculated the distribution-free Chebychev 95% confidence interval for each w_t, which yielded similar results.

Table 7

Likelihood ratios for testing behaviour switch: Base case.

Hypothesis	RS–price	RS–quantity
Non-collusive only	414.56	86.06
Collusive only	0.00	3.06

The other null hypothesis to be tested is that conduct was stable throughout the sample period: either all periods were collusive or all were in a price war phase: $\lambda=0$ or $\lambda=1$. The alternative is that λ lies strictly between 0 and 1. As can be seen from table 5, the estimated value of λ is 1 for the RS–price model, while for the RS–quantity model it is significantly greater than 0 but less than 1. The likelihood ratio test can be used to test the null hypothesis. In this case, $2(L_1-L_0)\sim\chi_1^2$, and the likelihood ratios are reported in the last row of table 7. It can be seen from the table that the hypothesis that only cooperative behaviour is observed is not rejected for the RS–price model. For the RS–quantity model, the hypothesis is not rejected at $\alpha=0.05$ ($\chi_1^2(0.05)=3.84$) but is rejected at $\alpha=0.1$.

The above analysis suggests that we cannot reject the hypothesis that no behavioural switch occurred under the RS–price model, but we may do so under the RS–quantity model. In the latter case, we may further conclude that reversion to Cournot behaviour occurred in 2Q85–4Q85 and in 2Q87.

6. Bayesian model choice

We now illustrate a model choice criterion based on Bayes factors and use it to identify the most preferred model among the five oligopoly models: Bertrnad, Cournot, cartel, RS–price, and RS–quantity. Rewrite eq. (5) in the following matrix form:

$$y^i=Z\mu^i+\varepsilon^i, \tag{24}$$

where the superscript i denotes the airlines, $y^i\equiv(v_{11}^i,\ldots,v_{1K}^i,v_{21}^i,\ldots,v_{TK}^i)'$, $\mu^i\equiv(v_1^i,v_2^i,\ldots,v_T^i,\rho^i)'$, and $\varepsilon^i\equiv(\varepsilon_{11}^i,\ldots,\varepsilon_{1K}^i,\varepsilon_{21}^i,\ldots,\varepsilon_{TK}^i)'$. $Z\equiv(I_T\otimes \boldsymbol{1}_K,r)$ is an $N\times(T+1)$ matrix, where the first T columns are the design matrix $I_T\otimes \boldsymbol{1}_K$ ($\otimes$ denotes the Kronecker product, I_T is the $T\times T$ identity matrix, and $\boldsymbol{1}_K$ is a $K\times 1$ vector of ones), and the last column is $r\equiv(r_{11},\ldots,r_{1K},r_{21},\ldots,r_{TK})'$. The Bayes factor in favour of the jth model relative to the kth model, given sample y^i, is defined as

$$BF_{jk}^i=\frac{\Pr(y^i\,|\,M_j)}{\Pr(y^i\,|\,M_k)}. \tag{25}$$

Bayesian analysis of model selection based on loss functions would suggest that the jth model is 'preferred' to the kth model if $BF^i_{jk}>1$ [see, for example, Zellner (1971)]. In other words, if $BF^i_{jk}>1$, the expected loss associated with choosing the jth model (for the purpose of decision-making) would be less than the expected loss associated with choosing the kth model.

In our case, the Bayes factor may be calculated for each airline. Suppose that $\varepsilon^i \sim \mathrm{N}(\boldsymbol{0}_N, h_i^{-1} I_N)$, where $\boldsymbol{0}_N$ is an $N\times 1$ vector of zeros and h_i is the unknown precision (the reciprocal of the variance). Furthermore, the prior distribution for (μ^i, h_i) takes the normal–gamma conjugate form. More specifically, for the jth model, the prior conditional distribution of μ^i, given h_i, is normal with mean vector μ^i_{j0} and covariance matrix $(h_i N^i_{j0})^{-1}$, and the prior marginal distribution of h_i is gamma with parameters s^i_{j0} and τ^i_{j0}. Under these assumptions, the Bayes factor can be computed. [The expression for $\Pr(y^i \mid M_j)$ is given in Zellner (1971) and Leamer (1978).] We consider here a simple symmetric case in which $s^i_{j0}=s^i_0$, $\tau^i_{j0}=\tau^i_0$, and $N^i_{j0}=N^i_0$ for each prior model j. The Bayes factor in this case can be shown to equal

$$BF^i_{jk}=\left(\frac{\tau^i_0 s^i_0+(y^i-Zu^i)'(y^i-Zu^i)+q^i_j}{\tau^i_0 s^i_0+(y^i-Zu^i)'(y^i-Zu^i)+q^i_k}\right)^{-(\tau^i_0+N)/2}, \tag{26}$$

where $u^i \equiv (Z'Z)^{-1}Z'y^i$ is the estimated parameters using the sample information, and q^i_j (and q^i_t) are defined by

$$q^i_j \equiv (u^i-\mu^i_{j0})'((Z'Z)^{-1}+(N^i_0)^{-1})^{-1}(u^i-\mu^i_{j0}). \tag{27}$$

q^i_j may be interpreted as measuring the incompatibility of sample information and prior models. We have, from (25), that $BF^i_{jk}<1$ if $\delta^i_j>\delta^i_k$. That is, the higher the degree of incompatibility of sample information and the prior of the jth model relative to the prior of the kth model, the smaller will be the Bayes factor BF^i_{jk} in favour of the jth model.

To calculate q^i_j we need to specify the prior parameters μ^i_{j0} and N^i_0 for each model. For the Bertrand model the prior assessments of $(v^i_1,\ldots,v^i_T)$ are $(-1,\ldots,-1)$, and the associated prior, ρ^i_0, is estimated from (25) using constraints $v^i_t=-1$ for $t=1,\ldots,T$. The priors of the Cournot and cartel models are similarly defined ($v^i_t=0$ for Cournot and 1 for cartel). The priors associated with the RS model are obtained using the results given in table 5. In particular, periods other than 2Q85–4Q85 and 2Q87 are classified as collusive periods in the RS–quantity model, whereas in the RS–price model every period is regarded as cooperative. Finally, the prior parameter N^i_0 is specified by considering the 'equal weights' situation in which we assign equal weights between the prior beliefs and the current data, implying

Table 8

Incompatibility of prior and sample, q_j^i: Base case.

Prior	American	United
Bertrand	193.05	209.32
Cournot	23.11	34.50
Cartel	38.09	44.62
RS–price	6.88	16.17
RS–quantity	6.46	16.06

$N_0^i = Z'Z$.[13] Table 8 reports q_j^i calculated for each airline under each of the five oligopoly models.

Our model choice criterion based on Bayes factors shows that, given our base case evidence, the Bertrand model is least preferred among the five models. The Cournot model is most favoured among the first three models (Bertrand, Cournot, cartel), but is inferior to the two RS models. The comparison between the two RS models suggests that the quantity version is preferred. Thus, the Bayes factor criterion suggests that behavioural switches by the airlines did occur.

7. Sensitivity analysis

In the previous two sections, results were obtained using base case data. We now conduct some sensitivity analyses to see how the results would respond to alternative values of the demand elasticity, η, the distance elasticity, θ, and the length of a decision period. Consider first the effects of changing the elasticity estimates on airline conduct estimation. The estimated overall mean conduct parameters of American Airlines are displayed in table 9 for $\eta = 1.2$, 1.6, 2.0 and $\theta = 0.25$, 0.50, 0.75. There are nine different combinations of cases, including the base case $(\eta, \theta) = (1.6, 0.50)$. The standard errors of the means are in parentheses. Table 10 displays the results for United Airlines.

Tables 9 and 10 show that an increase in η of 0.40 would induce an increase of 0.38 or less in the estimated overall mean conduct parameters. (Thus, an increase in η of 0.10 would result in an increase in the conduct parameters of about 0.09.) In a similar way, the conduct parameters increase with θ. Given that our variations in η and θ are fairly substantial, neither effect appears highly significant. The tables also show that the standard errors of the means increase with η but fall as θ becomes large. The effects,

[13]It can be shown [see, for example, Leamer (1978)] that the posterior estimate of μ^i, given the normal–gamma conjugate prior, is a weighted average of prior μ_0^i and sample estimate u^i. The 'equal weights' situation arises when $N_0^i = Z'Z$. We also note that our basic results concerning the relative accuracy of the five models (see table 8) would remain for the unequal weights cases.

Table 9

Sensitivity of overall mean conduct parameters, American.

	$\theta=0.25$	$\theta=0.50$	$\theta=0.75$
$\eta=1.2$	−0.04 (0.05)	0.06 (0.03)	0.13 (0.03)
$\eta=1.6$	0.28 (0.06)	0.42 (0.05)	0.51 (0.04)
$\eta=2.0$	0.60 (0.08)	0.77 (0.06)	0.88 (0.05)

Table 10

Sensitivity of overall mean conduct parameters, United.

	$\theta=0.25$	$\theta=0.50$	$\theta=0.75$
$\eta=1.2$	−0.01 (0.06)	0.08 (0.05)	0.14 (0.04)
$\eta=1.6$	0.32 (0.08)	0.44 (0.06)	0.52 (0.06)
$\eta=2.0$	0.64 (0.10)	0.81 (0.08)	0.90 (0.07)

however, are small. Finally, we note that the Bertrand hypothesis is unlikely to hold even for the extreme case $(\eta,\theta)=(1.2, 0.25)$ in which the degree of competitiveness is the highest. The lowest degree of competitiveness is reached at the highest elasticity estimates of the ranges, and has conduct parameters that are less than the cartel conduct of 1 but are much closer to 1 than to the Cournot conduct.

As long as the same η and θ are used for each period, variations in these estimates would not affect the basic patterns of conduct parameters across time. It would still be the case, for example, that the airlines appeared to behave less collusively in the 2Q85–4Q85 period than in the 2Q86–4Q86 period. Changes in η and θ would nevertheless affect the level of firm competitiveness in each period, as seen above.

Our assumption that demand elasticities are constant abstracts from variation across routes and over time. Strictly speaking, this is an error in variables problem which leads to a possible understatement of the statistical significance of our results. One way of trying to address this problem is by doing some Monte Carlo based sensitivity analysis of variation in η and θ.

We assumed that the elasticity of demand, η, was distributed normally with mean 1.6 and standard deviation 0.4. We also allowed θ, the cost tapering factor, to be normally distributed with mean 0.5 and standard deviation 0.175. These imply fairly broad confidence ranges for η and θ. We then allowed η and θ to be selected randomly from these distributions for each route and time period in the data set. This process was repeated 1,000 times, then 2,000 times.

Table 11
Estimated mean conduct parameters: Semi-annual.

Season	American		United		RS–quantity classification parameter	
	Mean	Standard error	Mean	Standard error	Mean	Standard error
4Q84–1Q85	0.45	0.14	0.35	0.17	0.61	0.05
2Q85–3Q85	0.05	0.14	0.15	0.17	0.46	0.06
4Q85–1Q86	0.14	0.14	0.52	0.17	0.56	0.07
2Q86–3Q86	0.60	0.14	0.39	0.18	0.67	0.06
4Q86–1Q87	0.65	0.15	0.27	0.18	0.66	0.07
2Q87–3Q87	0.46	0.15	0.52	0.18	0.63	0.07

The pattern of results with 2,000 trials was very similar to the pattern with 1,000 trials, so we seem to have enough trials to draw reasonable conclusions. The time-series pattern of conduct parameters was relatively unaffected by random variations in η and θ. The overall level of estimated implicit collusion did fall slightly, but by less than a standard deviation: to 0.37 and 0.39 (from 0.42 and 0.44) for American Airlines and United Airlines, respectively. We conclude that our analysis is robust to random variations over time and across routes in both demand elasticity and tapering factor and we suspect that exact measurement of these variables would strengthen our results, given the central errors in variables problem induced by our approach.

It is reasonable, however, to argue that errors in estimates of demand elasticity are not likely to be independent, but may occur consistently in a given route over time. This remains a legitimate caveat about our analysis.

We now consider the effects of changing the length of a decision period. Some industry observers may suggest that carriers choose their capacities on a semi-annual basis with a 'high demand' season and a 'low demand' season in each model year. The high demand season is composed of the second and third quarters of a (calendar) year, whereas the low demand season is composed of the other two quarters. Table 11 displays the estimated mean conduct parameters for each airline in each season using the base case elasticity estimates, along with the standard errors of the estimates. It also includes the RS–quantity classification parameter and standard error. (The three single quarters, 4Q87, 2Q88, and 4Q88, are not included in the table.) These results show that using semi-annual data tends to mask price wars through an averaging effect. In particular, the one quarter price war that was estimated for 2Q87 is 'lost' when this quarter is averaged in with 3Q87.

We are forced by the data to use the quarter as the smallest decision period. It is possible, of course, that the actual decision period might be less than one quarter. This could tend to mask regime changes as, for example, a one month price war could be averaged in with two months of more accommodating tacit collusion. We therefore suggest that our results probably tend to understate the importance of RS behaviour.

8. Qualifications and comments

In order to implement our analysis we have made many abstractions and approximations that serve to qualify our analysis and that deserve some comment. One issue that we have abstracted from is multimarket contact. When firms come into contact in more than one market (e.g. on different routes) the possibilities for strategic interaction are substantially complicated and enlarged. In our analysis, we simply assume that each route is an independent observation of a common underlying strategic situation. This represents a particular but limited kind of multimarket interaction.

A standard source on the theory of multimarket contact is Bernheim and Whinston (1990). They find that if different markets and different firms are identical and technologies exhibit constant returns to scale, then multimarket contact does not enhance the ability of firms to sustain collusive prices. They also find that relaxation of these assumptions to include market or firm asymmetries or scale economies implies that multimarket contact can facilitate tacit collusion. Evans and Kessides (1992) obtain the interesting empirical finding that multimarket contact does seem to cause a reduction in competition in the airline industry. Specifically, they found that fares tend to be higher when competing carriers have extensive inter-route contacts. Thus it is possible that our results are not representative of individual markets, but are best taken as represenative of situations where multimarket contact is high.

A related empirical question is whether different routes are characterized by different behaviour. Our assumption that each route is a repetition of the same underlying strategic situation suggests that the systematic components of route-specific behaviour should move together over time, with some variation being induced by the independent and random components of route-specific fluctuations. One way of checking this is to look at the correlation matrix of proportional price–cost percentage margins ($(p-c)/p$) on the different routes as these margins change over time.

There are 16 routes, implying a 256-cell symmetric correlation matrix of which 16 cells contain 'own' correlations and are therefore 1.0, leaving 120 independent cross-correlations. For United Airlines, all 120 of these cross-correlations were non-negative (rounding off to one decimal point), and 25 were 0.8 or above. For American, 21 of 120 cross-correlations were negative, but 13 of these were for the Las Vegas route. Excluding Las Vegas, the average cross-correlation was about 0.44 for American and 0.62 for United. Thus, price–cost margins tended to move together over time on the different routes, supporting our assumption that different routes represent replications of a common strategic environment.

Our analysis implicitly relies on exogenous demand shocks to generate RS behaviour and takes account of changes in costs. Despite the fact that demand and cost shocks are accounted for, it might be interesting to ask whether any obvious external shocks could be associated with regime changes. We added oil prices (meant to capture cost shocks) and real GNP (mean to capture demand shocks) as regressors in eq. (23). We find that the oil price appears to be statistically and economically insignificant in these regressions for both carriers. This is reassuring given that oil prices are already accounted for through the role of costs in determining conduct parameters. Real GNP (on the other hand) has a significant positive effect (t-stat. = 3.5) on the United Airlines conduct parameter and an insignificant positive effect on the American conduct parameter.[14]

Another exogenous event that might have had an effect on competition was a strike by United Airline pilots during the first part of the second quarter of 1985. As the results show, an apparent price war began in this quarter. The possibility that this was triggered by the pilots strike certainly seems reasonable. One possible interpretation is that the pilots strike increased the level of uncertainty each firm had about its rival's costs and strategic intentions. This would seem to make defection more likely. Our analysis is not inconsistent with this hypothesis.

We also assumed that there are only two strategic 'players' in this set of markets. We corrected for the presence of fringe firms in the data by including the market share of other firms as an explanatory variable in our regressions, but we did not explicitly consider entry. Harrington (1989b) provides a useful theoretical investigation of the relationship between barriers to entry and the ability to maintain tacit collusion. Rather obviously, the possibility of entry can affect the competitive interaction of incumbent firms. It does not, however, change the possible reliance on RS as a strategy for maintaining some degree of tacit collusion. Thus our analysis would still be reasonable even if entry barriers were low. If entry barriers differed across routes, this could explain some of the (modest) variation in conduct across routes. However, we would point out that the large hubs maintained by American and United in Chicago give these two airlines a strong incumbency advantage in the markets we examine.

One interesting question is whether the two carriers we examine are strictly symmetric. As pointed out by Harrington (1989a) asymmetries in discount factors (and other asymmetries) can have a major impact on the ability of firms to maintain tacit collusion and on the distribution of gains

[14]This finding might be taken as modest evidence against the Rotemberg and Saloner (1986) formulation suggesting that aggressive conduct would tend to occur during periods of high demand.

from tacit collusion. There is nothing in our analysis that strictly rules out such asymmetries between firms, and our analysis could be thought of as a test of how symmetric the firms are. Our results do show some differences in behaviour between the two airlines, but these differences seem relatively modest.

A related assumption is that products can be treated as homogeneous. This is obviously a substantial simplification. However, differences across fare classes for a single airline are probably more significant than differences between airlines as a source of heterogeneity. As discussed earlier, we try to take the most reasonable aggregate of fare classes, excluding first class. As for heterogeneity between airlines, American and United are in a very similar situation and offer very similar ranges of service on these routes. One possible complication is that even if the airlines offer similar products, these markets might display 'switching costs' induced by consumer loyalty and, more importantly, by frequent flier programs. As shown by, for example, Klemperer (1989), the presence of switching costs can induce price wars and, in general, certainly affects the dynamic interaction between firms. We therefore emphasize that our interpretation of the results should be treated with caution as switching costs could certainly be part of the story.

Another possible specification problem is that average variable cost (and marginal cost) per passenger-mile are treated as constant with respect to quantity on a given route. (They are, however, allowed to vary by flight length.) This is a fairly common practice, but it obviously can induce errors. Specifically, as indicated by the valuable survey on empirical studies of market power by Bresnahan (1989), if costs are in fact decreasing or increasing, then one can mis-estimate market power. Specifically, one mis-measures the true price–marginal cost margin and ends up mis-attributing what is really the slope of the marginal cost curve to market power. More precisely, if there are local increasing returns to scale, then assuming constant returns could account for the strong rejection of Bertrand pricing. Prices could exceed average costs, distance adjusted, and yet still be close to marginal costs.

A related concern is that connecting traffic might vary significantly across routes. Thus if economies of scale are important, the fact that we ignore connecting traffic would add error to our cost estimates. There is, however, little reason to believe this effect would bias the results in a particular direction.

We assumed that the demand elasticity, η, and cost-tapering factor, θ, were constant over routes and time. We also did sensitivity analyses to check on the effect of random variations in η and θ. There is, of course, a potential problem if variations in η or θ are systematic in some way. Most importantly, the elasticity of demand might have a cyclical pattern over time. If so, then we might be mistakenly attributing variations in demand elasticity to variations in conduct. One encouraging note, however, is that the entire sample period was contained within a single macroeconomic recovery, so we do not mix observations from very different macroeconomic environments. Furthermore, an analysis of residuals based on runs tests does not indicate any serial correlation over time in the errors. We believe, therefore, that the time-series structure of our results is not an artifact of varying demand elasticity.

Some readers might question whether the version of uncertainty assumed in the Green–Porter model and related models is relevant for airlines. The crucial point is that firms cannot distingish between a rival's 'cheating' and a bad realization of demand. We would argue that this is reasonable in the airline industry, at least in the short run. Specifically, while airlines can learn each other's fare schedules readily enough, they cannot readily observe the number of seats allocated to deep discount categories, and therefore do not know the average price (or the capacity) of rivals on a given route. Such information can be inferred with a lag, using the data set used in this paper. Strictly speaking, it would be possible to base contingent strategies on information learned with a long lag, but observation (and common sense) suggests that it would be very difficult to base tacit collusion on such information. We would argue, therefore, that the Green–Porter type of model is reasonable for the airline industry, although there are many natural avenues for both refinement and generalization.

Finally, it should be emphasized that the Green–Porter and related RS models are examples of the 'rational equilibrium' approach to non-cooperative games: players are assumed to think through the game fully, make optimal choices, and implement equilibrium strategies from the beginning of the game. An alternative approach [described briefly in Tirole (1988, p. 261)], is called evolutionary or adaptive game theory. Players are viewed as adopting heuristic strategies which may subsequently be adapted in response to various factors such as past success and the past behaviour of other players. One influential example is Axelrod's (1984) experimental analysis of repeated Prisonser's Dilemma games, in which 'tit-for-tat' strategies performed well.[15]

A tit-for-tat strategy implies that each rival begins by colluding, then follows the strategy pursued by the rival in the previous period. Firms might make mistakes or be subject to random influences, so the full series would not be continuous collusion, and a tit-for-tat result should be observable in the data. Our analysis is in the rational equilibrium tradition and we do not undertake a careful study of heuristic evolutionary strategies here. Fifteen

[15]It is interesting to ask whether tit-for-tat or other successful evolutionary strategies can emerge as equilibrium maximizing solutions. As described by Fudenberg and Tirole (1991, ch. 5), one can specify such games, but most games would not yield a rational equilibrium tit-for-tat solution.

periods is, in any case, too short to look seriously at evolutionary strategies. However, one ad hoc way of getting some insight about the possibility of tit-for-tat strategies is to regress each firm's conduct parameter on its rival's one-period-lagged conduct parameter: $v^i_{tk} = a + bv^j_{t-1,k} + \varepsilon$ for $(i, j = 1, 2)$. The result of such a regression yields parameter estimates of 0.36 and 0.48 for American and United Airlines, respectively, with associated t-statistics of 7.75 and 5.60. Note that a pure tit-for-tat strategy would suggest coefficients of 1, and can apparently be rejected. The fact that the coefficients are significantly different from 0 suggests, however, that there may be an evolutionary element in the responses. These simple regression results just reported do not imply serial correlation in the errors of the estimated Green–Porter models and, more generally, cannot be taken as direct evidence against the results reported in the rest of the paper. Evolutionary strategies are, however, an interesting possible alternative to the equilibrium models we estimate.

9. Concluding remarks

The study uses a panel data set of Chicago-based duopoly airline routes to investigate the dynamic interaction between American Airlines and United Airlines during the 1984–1988 period. We are able to reject the hypothesis that the dynamic path of prices and quantities was characterized by mere repetition of the Bertrand, Cournot, or cartel one-shot solution. We also test two versions of a regime-switching (RS) model of the Green–Porter type, allowing for the possibility of shifts between 'punishment' and 'collusive' phases. The two versions are 'price-setting' and 'quantity-setting' versions depending, respectively, on whether firms revert to one-shot Bertrand or Cournot solutions in punishment phases.[16]

The results of estimating the RS models suggest that, if we take the price-setting version as correct, then the firms were apparently in a collusive phase throughout the period. If, on the other hand, the quantity-setting version is taken as a maintained hypothesis, then the industry did apparently go through regime changes, with punishment phases occurring in the last three quarters of 1985 and the first quarter of 1987. Model selection according to Bayes factors favours the quantity-setting version. The basic evidence underlying these findings is that there was substantial instability in price–cost margins, but the periods of low price–cost margins were much closer to the Cournot outcome than to the Bertrand outcome. In addition, the periods of higher price–cost margins were definitely more collusive than the Cournot model would imply, but were not at the cartel point. Our conclusion is that

[16]More generally, we might expect that firms could use other price–quantity levels in punishment phases, as considered by Abreu (1986). Characterizing general solutions of this type in an empirically useful way remains, however, an interesting and significant challenge.

the data seem more consistent with the quantity-setting RS model than with the other alternatives considered here.

Overall, there seems to be substantial implicit coordination between American and United Airlines. However, during periods that would be deemed collusive, the extent of collusion is not very great. Even in the quantity-based RS model, the mean conjectural variation for collusive periods is less than 0.6. Furthermore, under the RS–quantity model, the frequency of price wars was fairly high (4 out of 15 quarters in our sample). More generally, our analysis suggests that it is important to explicitly consider dynamic stochastic interaction in small-numbers oligopoly, as a static 'snapshot' of an industry gives a very limited picture of the interaction.

References

Abreu, D., 1986, External equilibria of oligopolist supergames, Journal of Economic Theory 39, 191–225.

American Petroleum Institute, Petroleum Industry Statistics, 1990, Washington, American Petroleum Institute 10, no. 3.

Appelbaum, E., 1982, The estimation of the degree of oligopoly power, Journal of Econometrics 19, 287–299.

Axelrod, R., 1984, The evolution of cooperation (Basic Books, New York).

Bailey, E.E., D.R. Graham and D.P. Kaplan, 1985, Deregulating the airlines (MIT Press, Cambridge, MA).

Bernheim, D.B. and M.D. Whinston, 1990, Multimarket contact and collusive behavior, The RAND Journal of Economics 21, 1–26.

Berry, S., 1992, Estimation of a model of entry in the airline industry, Econometrica 60, 889–917.

Berry, S., 1990, Airport presence as product differentiation, American Economic Review (Paper and Proceedings) 80, 394–399.

Bernardino, F. and K. Blinder, 1991, Market conduct in the airline industry: Comment, Discussion paper (Gellman Research Associates, Jenkintown, PA).

Borenstein, S., 1989, Hubs and high fares: Dominance and market power in the U.S. airline industry, Rand Journal of Economics 20, 344–365.

Borenstein, S., 1990, Airline mergers, airport dominance, and market power, American Economic Review (Papers and Proceedings) 80, 400–404.

Brander, J.A. and A. Zhang, 1990, Market conduct in the airline industry: An empirical investigation, Rand Journal of Economics 21, no. 4, 567–583.

Bresnahan, T.F., 1987, Competition and collusion in the American automobile industry: The 1955 price war, Journal of Industrial Economics 35, 457–482.

Breshnahan, T.F., 1989, Empirical studies of industries with market power, in: R. Schmalensee and R. Willig, eds, Handbook of industrial organization, vol. 2 (North-Holland, Amsterdam) 1011–1057.

Day, N.E., 1969, Estimating the components of a mixture of normal distributions, Biometrika 56, 463–474.

Dixit, A.K., 1988, Optimal trade and industrial policies for the U.S. automobile industry, in: R.C. Feenstra, ed., Empirical methods for international trade (MIT Press, Cambridge, MA).

Evans, W.N. and I.N. Kessides, 1992, Living by the 'golden rule': Multimarket contact in the U.S. airline industry, Discussion paper (Department of Economics, University of Maryland, MD).

Friedman, J.W., 1971, A noncooperative equilibrium for supergames, Review of Economic Studies 28, 1–12.

Fudenberg, D. and J. Tirole, 1991, Game theory (MIT Press, Cambridge, MA).

Geroski, P., L. Phlips and A. Ulph, 1985, Oligopoly, competition and welfare: Some recent developments, Journal of Industrial Economics 33, 369–386.

Gollop, F. and M. Roberts, 1979, Firm interdependence in oligopolistic markets, Journal of Econometrics 10, 313–331.

Green, E.J. and R.H. Porter, 1984, Noncooperative collusion under imperfect price information, Econometrica 52, 87–100.

Hajivassiliou, V.A., 1989, Testing game-theoretic models of price-fixing behaviour (Cowles Foundation for Research in Economics, Yale University).

Harrington, J.E., 1989a, Collusion among asymmetric firms: The case of different discount factors, International Journal of Industrial Organization 7, 289–307.

Harrington, J.E., 1989b, Collusion and predation under (almost) free entry, International Journal of Industrial Organization 7, no. 3, 381–401.

Hurdle, G.J., R.L. Johnson, A.S., Joskow, G.J. Werden and M.A. Williams, 1989, Concentration, potential entry, and performance in the airline industry, Journal of Industrial Economics 38, 119–139.

I.P. Sharp Associates Ltd, 1988–1989, Department of transportation origin and destination survey, Data Bank 1A.

Iwata, G., 1974, Measurement of conjectural variations in oligopoly, Econometrica 42, 949–966.

Keeler, T., 1991, Airline deregulation and market performance, in: D. Banister and K. Button, Transport in a free market economy (Macmillan, London) ch. 5, 121–170.

Klemperer, P., 1989, Price wars caused by switching costs, Review of Economic Studies 56, 405–420.

Leamer, E.E., 1978, Specification searches: Ad hoc inference with non-experimental data (Wiley, New York).

Levine, M.E., 1987, Airline competition in deregulated markets: Theory, firm strategy, and public policy, Yale Journal on Regulation 4, 393–494.

Morrison, S.A. and C. Winston, 1987, Empirical implications and tests of the contestability hypothesis, Journal of Law and Economics 30, 53–56.

Morrison, S.A. and C. Winston, 1990, The dynamics of airline pricing and competition, American Economic Review (Paper and Proceedings) 80, 389–393.

Oum, T.H. and D. Gillen and D. Noble, 1986, Demand for fareclasses and pricing in airline markets, Logistics and Transportation Review 23, 195–222.

Oum, T.H., A. Zhang and Y. Zhang, 1993, Inter-firm rivalry and firm-specific price elasticities in deregulated airline markets, Journal of Transportation Economics and Policy (forthcoming).

Panzar, J., 1979, Equilibrium and welfare in unregulated airline markets, American Economic Review (Papers and Proceedings) 69, 92–95.

Porter, R.H., 1983a, Optimal cartel trigger price strategies, Journal of Economic Theory 29, 313–338.

Porter, R.H., 1983b, A study of cartel stability: The Joint Executive Committee, 1880–1886, Bell Journal of Economics 14, 301–314.

Rotemberg, J. and G. Saloner, 1986, A supergame-theoretic model of price wars during booms, American Economic Review 76, 390–407.

Schmalensee, R., 1977, Comparative static properties of regulated airline oligopolies, Bell Journal of Economics 8, 565–576.

Slade, M.E., 1987, Interfirm rivalry in a repeated game: An empirical test of tacit collusion, Journal of Industrial Economics 35, 499–516.

Slade, M.E., 1989, Price wars in price-setting supergames, Econometrica 56, 295–310.

Tirole, J., 1988, The theory of industrial organization (MIT Press, Cambridge, MA).

U.S. Department of Transportation, 1984–1988, Air Carrier Financial Statistics, and Air Carrier Traffic Statistics, Various issues (Washington, DC).

U.S. Department of Commerce, 1989, Bureau of Economic Analysis, Business Statistics 1961–1988, 26th edn. (1989), A Supplement to the Survey of Current Business (U.S. Government Printing Office, Washington, DC).

Whinston, M.D. and S.C. Collins, 1992, Entry, contestability, and deregulated markets: An event study analysis of People Express, Rand Journal of Economics 23, 445–462.

Zellner, A., 1971, An introduction to Bayesian inference in econometrics (Wiley, New York).