What's On The Table: Revenue Management And The Welfare Gap In The US Airline Industry

Yiwei Chen Vivek F. Farias *†

October 17, 2012

Abstract

The past decade has been a difficult one for the US airline industry. On the one hand, airline profits have been highly variable with net losses over the last ten years standing in the tens of billions of dollars. On the other hand, consumers continue to complain of predatory pricing and other such tactics. Our goal here will simply be to get an estimate of what is possible moving forward. We approach this task from an econometric perspective: we produce a status-quo dollar estimate of total welfare in the US airline industry. We then compute a number of benchmarks that we posit are conservative estimates of what optimal welfare in the industry. The key feature that admits the plausibility of these benchmarks is that the mechanisms that achieve them resemble existing dynamic pricing practices from airline revenue management (RM). In fact, it is the structure of these benchmark mechanisms in particular, and RM practice in general, that inform our modeling choices.

Our benchmark estimates will leverage a unique, proprietary dataset set on ticket purchases via the 'micro' BLP approach Berry et al. [2004]. We will show that the welfare gap is surprisingly large, raising the possibility that a combination of innovative selling mechanisms, legislation and network capacity allocation can make a dramatic difference to airline profitability and consumer surplus alike.

1. Introduction

A primary quantity of interest for any industry is the allocative efficiency that results from the industry's structure and other details of its operations such as sales mechanisms. In simple terms, one asks: are limited 'resources' effectively allocated to those who value these resources the most. It is easy to see that allocative efficiency is important not just to consumers within the industry but to producers as well; an inefficient industry signals not just an opportunity to improve the economic utility customers derive from that industry's activity but also the potential opportunity for sellers to improve their revenues. It is thus not surprising that measuring the allocative efficiency of an industry is a task one may view as being of 'first order' importance.

The present paper undertakes a study of the allocative efficiency of the US airline industry. We measure allocative efficiency in terms of the 'welfare gap' in the airline industry where welfare

*YC is with the Sloan School of Management, MIT. VF is with the Operations Research Center and Sloan School of Management, MIT. Emails: {ywchen, vivekf@mit.edu}

†We gratefully acknowledge the help of Tom Magnanti, Richard Ratliff, Barry Smith, and Ben Vinod for their support in obtaining data relevant to this project. We further acknowledge useful conversations with Benjamin Van Roy and Gabriel Weintraub that led to the formulation of the questions here.

is defined as the sum of airline revenues and consumer surplus generated as a result of ticket purchases, measured in dollars. Computing the welfare gap entails estimating welfare in the airline industry in its current state on the one hand, and comparing this estimate with an 'optimal' welfare benchmark on the other; this study is the first to produce an optimal welfare benchmark that is 'operationally' consistent with airline pricing practices. As we will discuss momentarily, both of the above econometric quantities will be derived from a mix of publicly available industry level data and auxiliary 'micro'-level proprietary data made available to us by a large US Global Distribution System (GDS). We are motivated to study the allocative efficiency of the airline industry for two simple reasons:

1. First, the flourishing research area of Revenue Management (RM) has reached a point wherein practical improvements derived from an RM innovation are measured in fractions of a percentage point. Should our analysis reveal a large degree of allocative *in*efficiency, this would serve as a strong impetus for airlines to consider more substantiative changes such as altogether new mechanisms for the sale of air tickets or carefully revisiting their long term allocation of capacity across the US network.

2. Consumers frequently complain of predatory pricing by airlines (See Wirtz et al. [2003]) while airlines, in turn, point to such pricing as simply arising of the necessity of allocating a relatively expensive resource among a population with substantially heterogenous values for the resource. A measure of allocative efficiency would effectively provide a scientific view of this issue one way or the other.

Our study will make the following key contribution: *We establish that the allocative inefficiency of the US airline industry is approximately 12.5 %. In particular, we value the allocative inefficiency of this industry at approximately eight billion dollars a quarter in 2006. This is a large figure and we believe its magnitude provides a great deal of support for the serious consideration of, among other things, alternative ticket selling mechanisms, the rationalization of airline schedules, and the future role legislation might play in this industry, particularly in reducing concentration on routes.*

1.1. Challenges and Approach

Our task calls for the estimation of a suitable structural model describing consumer utility in conjunction with a benchmark model describing 'optimal' welfare. Doing so presents us with two principle challenges; we discuss our approach and contributions in the context of these challenges.

1. Data: Data available publicly (through the DoT) will typically not suffice. While the reason for this will be clear in subsequent sections, we present a brief bottom line explanation here: price discrimination in the airline industry is common practice. Moreover, it is natural to imagine that such discrimination will be crucial in any posted price mechanism that seeks to sell tickets effectively. As such, it is crucial for us to develop a structural understanding of this discrimination which is not possible with publicly available data. We have obtained, in addition to publicly available data on ticket coupon sales, a dataset from a large GDS that describes customer attributes, consideration sets and eventual purchase decisions for a small set of markets. While this dataset covers a substantially smaller number of ticket sales than the DoT data (and so does not suffice for a welfare estimate by itself), it suffices for us to build a reduced model that captures the key features of price discrimination under current RM practices. A number of qualitative features of the utility and price discrimination models we

estimate show *remarkable* agreement with expectations from RM practice, estimates produced in other studies where relevant, and conclusions drawn about the industry via entirely different means (such as, for instance, quality surveys).

2. A Pragmatic Benchmark: By far the biggest impediment to our study is a believable welfare benchmark. For instance, one benchmark one might consider is the following: the airline waits for all potential demand to realize and then conducts a (welfare maximizing) second price auction to assign seats to potential customers. This mechanism is not practical relative to the requirements of the airline industry where a commitment to sell must be made at the time of a customers arrival. As such, any welfare benchmark should be implementable via a mechanism seen as practical; an example of such a mechanism given current RM practice might be a (dynamic) posted price mechanism. Our dataset will permit us to incorporate concrete customer 'type' data in our structural model. This type data will allow us to calibrate a benchmark based on type specific posted prices. The types we consider will coincide with a customer's time of arrival and as such our benchmark will translate to what is, in essence, a dynamic pricing policy that is consistent with extant RM practice. We will also conjecture on the sources of inefficiency in the status quo with the support of some simple stylized analysis.

The overall methodology we will eventually employ in calibrating the structural models our study will call for is somewhat complex. The need for this complexity arises from the necessity to utilize both aggregate, market level data in conjunction with 'micro-level' customer data effectively. We will employ the 'micro' BLP methodology pioneered by Berry et al. [2004]. We hope that in this regard the present work can serve as a 'users guide' for future applications of this useful methodology in operations management where a number of situations call for the incorporation of a price discrimination model.

The remainder of the paper will be organized as follows. In the next Section, we will present our models for consumer utility and price discrimination. We will then present the structural equations that will eventually allow us to estimate these models. Finally, we will develop and discuss optimal welfare benchmarks. Section 3 will describe our estimation procedure. Whereas the general procedure is by now somewhat standard in the econometrics/ IO literature, our setup calls for a few modifications, and as such, we present a self contained description. The main contribution of this paper is presented in Section 4 where we present the estimated structural model and our welfare estimates. Before proceeding, we will next present a literature review that places our work in the context of the (large) body of empirical research on the airline industry. This work stems largely from the Industrial Organization community and to a lesser extent, the Operations Management community.

1.2. Literature Review

There are several streams of literature that are relevant to our paper. In this relatively concise literature review we attempt to touch on each of these streams; the papers we reference themselves point to important antecedent literature.

Characterization of Fares: Borenstein and Rose [1994] find evidence of the existence of price dispersion in the US airline industry of a magnitude as large as 36% of the airline's average ticket price. They also find evidence that prices dispersion increases as markets become more competitive. Dana [1999a] and Dana [1999b] propose a theory to explain that ticket prices increase and are more dispersed as load factor increases. Gale and Holmes [1992] propose a theory to explain that more discounted 'advance purchase' seats are sold in off-peak demand periods. Puller et al. [2008] find the evidence that an increase in the load factor is associated with a very modest increase in average fares and a modest decrease in price dispersion. Dana and Orlov [2009] report that internet penetration has positive impact of flight load factors. Borenstein [1989], Borenstein [1991], and Borenstein [2005] report the evidence of the hub premium for flights out and into the hub airport. Forbes [2008] find that the price fell by $1.42 on average for each additional minute of flight delay at LaGuardia Airport after the legislative change in takeoff and landing restrictions was made.

Industry Structure and Welfare: Morrison and Winston [1995] provide a comprehensive review of the airline industry's evolution, both before and after the Airline Deregulation Act in 1978. The book identifies the problems that the industry faces, analyzes their causes, and suggests fixes. Borenstein and Rose [2007] conduct a comprehensive review of the regulatory reform in the US airline industry. It discusses events that lead to deregulation of the industry and evaluate the impact of those reforms. Netessine and Shumsky [2005] study the impact of airline industry competition on seat allocation. They show that more seats are protected for higher-fare passengers under horizontal competition (two airlines compete for passengers on the same flight leg), and booking limits may be higher or lower under vertical competition (different airlines fly different legs on a multileg itinerary), which depends on the demand for connecting flights in each fare class. Li et al. [2011] find empirical evidence that the percentage of strategic customers ranges from 4.9% to 49.9%, measured by the 5th and 95th percentiles. Ciliberto and Tamer [2009] invent a novel methodology to investigate the empirical importance of firm heterogeneity as a determinant of market structure in the US airline industry. The paper finds evidence of heterogeneity across airlines in their profit functions. Berry and Jia [2009] explore the impact of tremendous turmoil in the US airline industry in the early 2000's, where there are four major bankruptcies and two major mergers. The paper finds that air-travel demand was 8% more price sensitive and passengers are more preferable to non-stop flight. Bamberger et al. [2004] find that the average fares fell $5-7\%$ after the creation of alliances on those city pairs affected by the alliances and total traffic increased by at least 6%. Park and Zhang [1998] and Park and Zhang [2000] study the impact of international airline alliances between US and foreign carriers on the variety of flight options and markets. Goolsbee and Syverson [2008] investigate how incumbent responds to threat of entry by competitors. The paper uses the evolution of Southwest Airlines' route network to identify particular routes where the probability of future entry rises abruptly. It finds that incumbents cut fare significantly when threatened by Southwest's entry. The evidence on whether incumbents are seeking to deter or accommodate entry is mixed. Peter [2006] simulates post-merger prices for six major airline mergers and finds that the effect of ownership transfer on price incentives plays key role on post-merger price changes.

As for welfare, Armantier and Richard [2008] explore the impact of code-sharing alliance between Continental Airlines and Northwest Airlines on consumer welfare. The paper finds significant welfare gains for passengers who take connecting flights, whereas the welfare of passengers who take nonstop flight were hurt sharply. Park [1997] also study the impact of airline alliances on economic welfare. As one might expect, they find that complementary alliances increase economic welfare, and parallel alliances decrease it.

Revenue Management: Talluri and van Ryzin [2004] provide a comprehensive review on both general dynamic capacity allocation heuristics and bid-price controls in the network revenue management system. Vulcano et al. [2002] present dynamic mechanisms for airline revenue management that show the potential to outperform classical dynamic pricing policies. Akan et al. [2011] use a mechanism design approach to construct innovative advance purchase contracts for revenue management with customers with heterogenous valuations; these are optimal mechanisms.

Estimation Methodology: The seminal work of Berry et al. [1995] (BLP hereafter), establish a procedure for estimation of random-coefficient discrete choice demand models, which incorporate unobserved product attributes and endogenous prices using *aggragate* market level data. Nevo [2000] provide a 'users guide' to this approach. Dube et al. [2009] develop a novel computational approach to the same problem. There is a huge body of applied research using BLP or its variants. To name a few in the airline industry, Berry et al. [2006] develop a two-type customer BLP model and use this framework to explore the impact of airline hubs have on fares. Berry and Jia [2009] use a similar model to study the impact of tremendous turmoil in U.S. airline industry in the early 2000's. In recent years, there is an emerging literature that incorporates micro (individual) level data into the BLP model which otherwise only uses macro (aggregate) level data ('micro' BLP hereafter). Berry et al. [2004] jointly use macro level data of automobile aggregate demand and a micro-level data set from General Motors which reveals the second-choice preference of people who bought a GM car with an aim to better characterize customer's purchasing behavior. Armantier and Richard [2008] use individual customer choice data to similar effect to study the impact of a merger on welfare.

We note that even though BLP and 'micro' BLP methodologies allow customers to be heterogeneous with different types, these methodologies typically do not associate types with customer traits identifiable from data. Rather these types are induced by some parametric distribution over coefficients in the utility model. Our research departs from this trend by defining types a-priori based on revenue management segmentation practices. This plays a key role in allowing us to build believable welfare benchmarks.

2. Structural Models and Optimal Welfare

Our study of airline efficiency can be conducted at various 'time scales'. Depending on the time scale it might be relevant, for instance, to consider an airlines capacity investment decisions. Our study will be conducted over the course of a quarter and as such we will make the assumption that all investment decisions and decisions on flight schedules have been made and may not be adjusted over this time span. In particular, airline costs are effectively sunk. This level of granularity is consistent with that in many modern empirical studies of the airline industry (see, for example, Armantier and Richard [2008]). As such, allocative inefficiency, if any, arises from (a) pricing decisions airlines make in selling tickets over the course of the quarter and (b) capacity allocation decisions across routes. Studies at longer time scales will potentially reveal further inefficiencies (for instance, in network formation) but are beyond the scope of the present study. We begin with defining a number of concepts relevant to our setup:

Markets: A market is defined as the collection of all 'itineraries' from a particular origin ('O') to a particular destination ('D') and back within a quarter. For instance, we understand by the Boston-San Francisco market the collection of all itineraries from Boston to San Francisco and back within a quarter. These itineraries may include intermediate stops. Notice, further that we distinguish the Boston-San Francisco market from the San Francisco- Boston market. Further notice that we ignore 'one way' markets. We do this primarily for tractability and since the fraction of such itineraries sold is small (less than 20%), We will index markets by m; $m = 1, 2, \ldots, M$. Market m will be associated with a market 'size', D_m. Colloquially one may think of this as the size of the pool of potential customers across all products sold in that market and following the example of other researchers such as Berry and Jia [2009], we will take this quantity to be the geometric mean of the population in the areas of the origin and destination airports.

Products: A product is what one might colloquially think of as a return ticket. In particular, a product is associated with a market. In addition a product specifies a number of attributes, including: (i) the *itinerary*, which is simple a sequence of airports including the origin and destination, and potentially intermediate stops, (ii) The carrying airline, (iii) A variety of other features including, for instance, whether either/both of the origin and destination are 'hubs', whether the flight is non-stop and the distance covered by the itinerary. A number of product features that are presumably relevant to consumer decision making are *not* observed (such as, for instance, an advance purchase or weekend stay requirement, or potentially , even product specific advertising effects).

We will index products in market m by j, $j = 1, 2, \ldots, J_m$. We denote by $N = \sum_m J_m$ the total number of products in the industry. Observed product features will be encoded by the variable X_{jm} and unobservable attributes by ξ_{jm}.

Consumers: Every market is associated with a set of consumers. These consumers are associated with a number of distinguishing features observable to the airline, and partially, to the econometrician. In particular, a customer is associated with a time of purchase, and a class and date of travel desired. Of course, the customer's desired origin and destination correspond to the market she is associated with.

We will index consumers in a given market, m by i. We will associate every consumer with a *type* which in general could correspond to some set of observable consumer features and take on one of finitely many values. Here we will take a consumers type to simply be the time of her purchase relative to the first departure date of the itinerary. We will let $R(i)$ denote the type of the ith consumer. As we will note in the sequel, it is important to be able to identify a customers type with customer data that one might expect is available to the airline; this will have important implications for our ability to construct a viable social welfare benchmark and simultaneously estimate a structural model of consumer utility that incorporates heterogeneity in a meaningful way. It is also worth mentioning here that our choice of how we define 'type' will be informed by revenue management insights – presumably the price discrimination factors that influence airline RM should make for good factors on which to distinguish customers. This is in contrast with extant econometric studies of the airline industry which treat customer type abstractly by allowing for, and fitting, general parametric distributions over coefficient vectors.

Finally, every consumer i in market m is associated with a consumer specific price quote for each product j in that market. We denote this price by p_{ijm}. The dependency of price on consumer attributes complicates our problem both in terms of data required as well as estimation but it is necessary: type specific price discrimination is a crucial feature of existing RM practice and is obviously an important lever in designing an alternative benchmark as well.

2.1. Consumer Utility and Price Dispersion

With the above setup, we are now in a position to state and understand a structural model for consumer utility. In particular, consumer i garners utility u_{ijm} from the jth product in market m, as given by:

$$u_{ijm} = -\alpha_{R(i)} p_{ijm} + \beta_{R(i)}^{\top} X_{jm} + \xi_{jm} + \epsilon_{ijm},\ \forall i, j \in m, m.$$

As discussed X_{jm} is a (say, $Q \times 1$) vector of observed product properties and ξ_{jm} represents the effect of unobserved product characteristics. The price and feature co-efficients $\alpha_{R(i)}$ and $\beta_{R(i)}$ must be estimated for each possible type $R(i) \in \{1, 2, \ldots, R\}$ in addition to a distribution, γ over these types. ϵ_{ijm} is idiosyncratic noise and assumed to be a standard Gumbel random variable. We will denote the value of the outside option in market m as u_{i0m} which we also assume to be a standard Gumbel random variable. Several salient features of this model are worth noting:

1. Recall that under our definition, a given 'product' can be offered at different prices, and as such in the sequel we will make the assumption that all products are available to a customer [1]. This is, in fact, a very reasonable assumption: load factors on most airlines are typically well below 90%; see Dana and Orlov [2009]. We further assume that upon arrival, a customer chooses the product that maximizes her utility; in particular, she picks a product in $\mathrm{argmax}_j\, u_{ijm}$. Customers do not strategize about timing their purchase.

2. While the above model is a random co-efficient model, our specification is such that we will be able to ascribe a specific set of co-efficients to a specific customer based on her (observable) type, i.e., $R(i)$ is observable to the econometrician as well as the airline. This will permit our welfare benchmark model to use type specific prices. Since the types we eventually define will be based on time of arrival, this will translate to a pricing scheme whose format resembles current practice. In the absence of such 'concrete' types it is difficult to construct a believable model to serve as a welfare benchmark.

3. We will allow for prices to be endogenous in that they are correlated with unobserved product features. More precisely, p_{ijm} is potentially correlated with ξ_{jm}.

4. Consider writing $p_{ijm} = \overline{p}_{jm} + e_{ijm}$ where e_{ijm} is zero mean, so that we interpret $\overline{p}_{jm}$ as quoted price for product j averaged over the population of consumers in market m. This yields a utility model of the form

$$u_{ijm} = -\alpha_{R(i)} \overline{p}_{jm} + \beta_{R(i)}^{\top} X_{jm} + \xi_{jm} + \epsilon_{ijm} + \hat{e}_{ijm}.$$

   As observed by Armantier and Richard [2008], $\hat{e}_{ijm}$ while zero mean is likely correlated with X_{jm} and depends on α so that one needs to impose further structure on the error term e_{ijm} to make progress here. We will refer to such a model as a 'price dispersion' model which we describe next.

Price Dispersion Model: Recall from our discussion above that we seek to write $p_{ijm} = \overline{p}_{jm} + e_{ijm}$ where e_{ijm} is zero mean. This price error term encapsulates the details of price discrimination, and we posit the following reduced form model to describe it. We assume:

$$e_{ijm} = c^{\top} D_{R(i)jm} + \eta_{ijm}$$

where η_{ijm} is an independent normal random variable with mean zero and variance $\sigma^2(\overline{p}_{jm})^{\zeta}$ and $D_{R(i)jm}$ is a vector capturing features of the customer and product. We will require that this feature vector have mean zero averaged over the population[2].

[1] Colloquially, products are also associated with a price, so that a product being unavailable in industry jargon corresponds in essence to a particular itinerary not being available at a particular price.

[2] which we may accomplish, for instance, by de-averaging.

Again, it is worth discussing a few salient features of the price dispersion model: First, notice that the features $D_{R(i)jm}$ are allowed to depend jointly on customer attributes and product features. Our data set[3] will allow us to use features that go beyond a separable specification. We will see this to be an important distinction since the impact of a customers type is apparently more relevant to certain types of products than others. Since mis-specifying this structure can have substantial implications on the estimation of price coefficients which in turn will strongly impact our welfare benchmark, we see this to be an important distinction. As a second point, our specification of η_{ijm} allows for heteroscedasticity which has been found to be an important feature of price dispersion in the airline industry in numerous pieces of research.

Incorporating this price dispersion model into our structural model for consumer utility yields:

$$\boxed{u_{ijm} = -\alpha_{R(i)} \left(\overline{p}_{jm} + c^\top D_{R(i)jm} + \eta_{ijm}\right) + \beta_{R(i)}^\top X_{jm} + \xi_{jm} + \epsilon_{ijm},\ \forall i, j \in m, m.}$$

We end with a brief comparison to structural models utilized in two recent empirical work on the airline industry, namely Berry and Jia [2009] and Armantier and Richard [2008]. Of these, Berry and Jia [2009] consider a model where the dependence of price on the customer is ignored altogether. However, since that piece of work seeks to characterize the evolution of the airline industry, a structural model relating prices to costs is needed there. Armantier and Richard [2008] consider a reduced form model that attempts to incorporate price discrimination, but with two important distinctions: first, the random co-efficient model there does not attribute co-efficients to identifiable customer types but instead is abstract. In particular, it would not be possible for us to construct our welfare benchmark using such a model. Second, the price dispersion model considered there does not consider the impact of customer attributes on price [4]; we know from RM practice that this is a *crucial* feature of pricing practice in the US airline industry.

2.2. Market Share and Observed Price Equations

Here we summarize the structural equations for market share and observed price that follow from the specification we have just presented for a consumer's utility. In particular, we denote by s^r_{jm} the expected fraction of type r consumers who purchase product j in market m. Further we denote by s_{jm} the overall fraction of customers in market m that purchase product j in the data. Recall here that we assume that γ_r denotes the expected fraction of type r consumers in a given market so that we must have

$$s_{jm} = \sum_{r=1}^{R} \gamma_r s^r_{jm},$$

where we plug in s_{jm} as an estimate for the expected market share of product j in market m. Recalling the definition of u_{ijm} then, we must have

$$s^r_{jm} = \mathsf{E}_{jm}\left[\frac{\exp(u_{ijm})}{1+\sum_{j'}\exp(u_{ij'm})}\middle| R(i) = r\right],$$

[3] in particular, our auxiliary data which will consist of a sample of consideration sets along with the purchase decision made and the time of purchase.

[4] the authors there claim that their model captures price shocks that depend on customer attributes in an additive manner, but this claim appears to be incorrect.

where the subscript on the expectation denotes that the expectation is over the random variable η_{ijm} with j and m understood as being fixed. This yields the following market share equations:

$$s_{jm} = \sum_r \gamma_r \mathsf{E}_{jm} \left[\frac{\exp(u_{ijm})}{1 + \sum_{j'} \exp(u_{ij'm})} \middle| R(i) = r \right] \quad \forall j, m \tag{1}$$

Our primary dataset does *not* indicate offered prices but rather prices at which a product was purchased. In particular letting $\overline{\mathbf{p}}_{\mathbf{jm}}$ denote the expected price paid for product j in market m conditioned on j being purchased from among the products available in market m, we have the following structural equations [5] relating the this quantity to the average price at which product j is offered in market m, $\overline{p}_{jm}$:

$$\overline{\mathbf{p}}_{\mathbf{jm}} = \overline{p}_{jm} + \frac{1}{s_{jm}} \sum_r \gamma_r \mathsf{E}_{jm} \left[\frac{\exp(u_{ijm})(c^\top D_{rjm} + \eta_{ijm})}{1 + \sum_{j'} \exp(u_{ij'm})} \middle| R(i) = r \right] \quad \forall j, m \tag{2}$$

In our estimation procedures, we will first estimate the price dispersion model from an auxiliary data-set. Following this we will use the relationships derived above (with s_{jm} and $\overline{\mathbf{p}}_{\mathbf{jm}}$ being estimated from our primary data-set) to estimate the remaining unknown parameters via a natural extension of the BLP methodology.

2.3. Criteria for Identification

Recall that we assume that the unobserved shock for product j in market m, ξ_{jm} is uncorrelated with product attributes but potentially correlated with the average offer price $\overline{p}_{jm}$. We will therefore seek L instruments Z^l_{jm} that are uncorrelated with the shock ξ_{jm} but explain the variability in p across products and markets. More precisely, we will make the following identification assumption: Let us denote by θ the vector of model parameters to be estimated (excluding average offer price $\overline{p}$, and ξ) and let θ^* denote its true value. For a given value of θ, let us denote by $\overline{p}(\theta), \Xi(\theta)$ values of $\overline{p}$ and ξ respectively such that $\theta, \overline{p}(\theta), \Xi(\theta)$ satisfy the market share and observed price equations. Armed with this notation, we make the following identification assumption:

Assumption 1. *For all l,*

$$\mathsf{E}\left[Z^l_{jm} \Xi(\theta)_{jm} \right] = 0$$

if and only if $\theta = \theta^*$.

Now denote by $\underline{X}_{jm}$ the observable product attributes X_{jm} excluding the constant term 1; we assume that $\underline{X}_{jm}$ is a Q dimensional vector whose components we index by q. Further, we denote by $f(j,m)$ the index of the firm that sells product j in market m, and by $\mathcal{F}_f$, the set of airline-itineraries (i.e. products) that are produced by firm f.

We then consider the following L instruments for ξ. The first Q of these instruments simply correspond to the Q observable product attributes, i.e.

$$Z^l_{jm} = \underline{X}_{jm,l} \quad \text{if } l \leq Q.$$

The next Q instruments correspond to the average value of each of the observable product attributes for all products produced by the same firm, excluding product j, so that

$$Z^l_{jm} = \frac{\sum_{j' \neq j, j' \in \mathcal{F}_{f(j,m)}} \underline{X}_{j'm,l}}{|\{\mathcal{F}_{f(j,m)}\} \setminus \{j\}|} \quad \text{if } Q+1 \leq l \leq 2Q.$$

[5] (2) relies on the market size being large

The third Q instruments correspond to the average value of each of the observable product attributes, averaged over all products produced by all other firms, so that

$$Z^l_{jm} = \frac{\sum_{j' \notin \mathcal{F}_{f(j,m)}} X_{j'm,l}}{|\cup_{f \neq f(j,m)} \mathcal{F}_f|} \quad \text{if } 2Q+1 \leq l \leq 3Q.$$

The preceding instruments are 'standard' choices; see Berry et al. [2004]. In addition, since a variable that impacts carrier costs but not demand is likely to be a good instrument, we also use HUB_D_{jm}, the indicator function for the destination of product j in market m being a hub as an additional instrument. More precisely,

$$Z^l_{jm} = \text{HUB\_D}_{jm} \quad \text{if } l = 3Q+1.$$

See Berry and Jia [2009] for a further discussion on this instrument.

2.4. Benchmarks

A number of econometric quantities will be of interest to us. We will primarily be interested in social welfare (measured in dollars), which in turn is a sum of airline revenues as a result of ticket sales and the consumer surplus (measured in dollars) generated by the same sales. Under our model, the expected surplus of consumer i in market m, $\mathsf{E}[\text{CS}^m_i]$ measured in dollars, is simply given by:

$$\mathsf{E}[\text{CS}^m_i] = \mathsf{E}\left[\frac{\max_j u_{ijm}}{\alpha_{R(i)}}\right] = \mathsf{E}\left[\frac{1}{\alpha_{R(i)}} \log\left(1 + \sum_j \exp(\hat{u}_{ijm})\right)\right],$$

where $\hat{u}_{ijm} = u_{ijm} - \epsilon_{ijm}$. The second equality follows from McFadden [1978]. The expected revenues earned by all firms in market m is given by

$$\mathsf{E}[\text{Rev}^m] = \sum_j D_m \overline{\mathbf{p}}_{jm} s_{jm}.$$

Having estimated our structural model, the above expressions make our estimate of social welfare under current pricing policies transparent; in particular, the expected social welfare is given by the expression $\sum_m D_m \mathsf{E}[\text{CS}^m_i] + \mathsf{E}[\text{Rev}^m]$. We next discuss establishing a benchmark for optimal welfare.

2.4.1. A Lower Bound on Optimal Welfare

Consider restricting attention to a scenario wherein all customers of type r in market m are offered a price $\hat{p}_{rjm}$ for product j. The expected consumer surplus of a customer in market m in such a scenario is then simply

$$\mathsf{E}[\text{CS}^m_i(\hat{p})] = \mathsf{E}\left[\frac{1}{\alpha_{R(i)}} \log\left(1 + \sum_j \exp(-\alpha_{R(i)}\hat{p}_{R(i)jm} + \beta^\top_{R(i)} X_{jm} + \xi_{jm})\right)\right];$$

relative to our earlier expression for consumer surplus, here the random price variable p_{ijm} is taken to be $\hat{p}_{R(i)jm}$. In a similar vain, the expression for total expected revenues for all firms in market m becomes

$$\mathsf{E}[\text{Rev}^m(\hat{p})] = D_m \sum_j \sum_r \gamma_r \hat{p}_{rjm} s^r_{jm}(\hat{p}),$$

where

$$s^r_{jm}(\hat{p}) = \frac{\exp(-\alpha_r \hat{p}_{rjm} + \beta_r^\top X_{jm} + \xi_{jm})}{1 + \sum_{j'} \exp(-\alpha_r \hat{p}_{rj'm} + \beta_r^\top X_{j'm} + \xi_{j'm})}$$

is the expected market share for product j in market m among type r customers.

We then posit two potential benchmark estimates for what one might consider 'optimal' social welfare. The first, more conservative benchmark allows us to reallocate, relative to the status quo, the assignment of tickets for any given product between the two types of customers but does *not* permit any reallocation of resources (i.e. seats on an itinerary leg) across products. This benchmark is given by the optimal solution to the following problem:

$$\begin{aligned} \max_{\hat{p}} \quad & \sum_m D_m \left(\mathsf{E}[\mathrm{CS}_i^m(\hat{p})] + \sum_j \sum_r \gamma_r \hat{p}_{rjm} s^r_{jm}(\hat{p}) \right) \\ \text{s.t.} \quad & \sum_r \gamma_r s^r_{jm}(\hat{p}) \leq s^0_{jm}, \forall j, m \end{aligned} \tag{3}$$

where s^0_{jm} represents the market share of product j in market m under *current* pricing practices. We will refer to the optimal value of this optimization problem as $OPT(SW1)$. This benchmark allows us to obtain *a lower bound on the welfare increase that obtains from simply allowing the social planner to re-allocate among the two types of customers, seats that are currently allocated in the status quo. In particular, no net increase in the number of seats allocated on any network leg is allowed.*

In addition, we might want to permit a reallocation of resources across products – for instance, we might want to increase the availability of a particular itinerary and doing so might require reducing the availability of some other itinerary that shares legs. In this sense, we permit a reallocation of the ability to travel among a broader group of customers. This particular benchmark is given by the optimal solution to the following problem:

$$\begin{aligned} \max_{\hat{p}} \quad & \sum_m D_m \left(\mathsf{E}[\mathrm{CS}_i^m(\hat{p})] + \sum_j \sum_r \gamma_r \hat{p}_{rjm} s^r_{jm}(\hat{p}) \right) \\ \text{s.t.} \quad & \sum_{m,j: l \in L(m,j)} D_m \sum_r \gamma_r s^r_{jm}(\hat{p}) \leq \sum_{m,j: l \in L(m,j)} D_m s^0_{jm}, \forall l \end{aligned} \tag{4}$$

where l indexes legs of the network and $L(m, j)$ is the set of legs that are part of the itinerary of product j in market m. We will refer to the optimal value of this optimization problem as $OPT(SW2)$; of course, by construction $OPT(SW2) \geq OPT(SW1)$ representing the gains from allowing a reallocation of resources across products. In addition to the previous benchmark, this benchmark also permits welfare gains due to *the re-allocation of seats on a given network leg across itineraries that use that leg. Again, no net increase in the number of seats allocated on any network leg is allowed.*

There are several points worth discussing with respect to these bounds:

1. A Meaningful Policy: As discussed earlier, an important concern with constructing a welfare benchmark is the form of the pricing policy implicit in this benchmark. The pricing policy implicit in the benchmarks above are of a practical nature. In particular, notice that since the customer types we will use will correspond with the time of a customers arrival, the pricing policies implicit in the benchmarks will correspond to dynamic, posted price policies which are of the form that the airline industry already uses.

2. Conservative Lower Bounds: In reality, the load factor (i.e. the fraction of capacity sold) for a given carrier is well below 100% (in recent years, the number has been approximately 80%). Both bounds above *do not allow seats that are unsold in the status quo allocation to be allocated in any way.* Since these unsold seats are likely unsold due to pricing policies (as opposed to a lack of demand), this represents a substantial restriction. Conversely, any conclusions we draw will in essence not count on increasing aggregate demand which we view as a robust feature.

3. Strategizing Consumers: Under fairly mild conditions, one may show that the prices that emerge from the optimal solutions to either of the benchmark optimization problems are immune to consumers that strategize on their time of purchase since the prices are constant over time. To see why this is the case, observe that if one allocated a common resource to two distinct groups of customers (say, for instance leisure travelers and business travelers) at different prices, a net welfare increase obtains by transferring a unit sold to the group that receives the lower price to the group that receives the higher price provided the sizes of both groups are sufficiently large to permit such a transfer.

4. Other Bounds: Further constraints can be added to the optimization problems (3) or (4) to compute other bounds of interest. For instance, one might consider bounds computed with the side constraint that the total revenues earned by producers are no worse than the status quo. In particular, this constraint would read

$$\sum_m D_m \sum_j \sum_r \gamma_r \hat{p}_{rjm} s^r_{jm}(\hat{p}) \geq \sum_m D_m \sum_j \overline{\mathbf{p}}^0_{jm} s^0_{jm}$$

where quantities with a 0 superscript represent the status quo.

3. Methodology

We devote this section to describing the methodology used in estimating the structural equations that underlie our analysis. The description of our methodology will also clearly state the nature of the data required; concrete details about this data will be given in the following section. The techniques employed extend the 'micro' BLP methodology Berry et al. [2004]. Since the use of this methodology is apparently new in the Operations Management literature we provide a concise, self contained outline in this section.

3.1. The Price Error Model

Recall from Section 2, that we need a reduced-form model to describe price discrimination. In particular, recall that we posited that the price offered to customer i for product j in market m, p_{ijm} depends on characteristics of the product j and the customer i according to

$$p_{ijm} = \overline{p}_{jm} + c^\top D_{R(i)jm} + \eta_{ijm}$$

where $\overline{p}_{jm}$ is the quoted price for product j averaged over the population of consumers in market m; $D_{R(i)jm}$ is a vector capturing features of the customer and product whose mean, averaged over the population, is zero; and η_{ijm} is an independent normal random variable with mean zero and variance $\sigma^2(\overline{p}_{jm})^\zeta$.

Our goal will be to estimate the co-efficient vector c and the exponent ζ describing the extent of heteroscedasticity in prices. We assume access to data of the following form: (a) A small subset of product-market pairs, (j, m), $\mathcal{A}$; (b) For each $(j, m) \in \mathcal{A}$, a representative sample set of customers $\mathcal{C}_{j,m}$; (c) The ***quoted*** price p_{ijm} for each customer $i \in \mathcal{C}_{jm}$.

Given the above data, one may hope to estimate the above model using maximum likelihood estimation. In particular, denote by θ_2 the set of parameters σ, ζ, c and by $\overline{p}$ the vector of prices $\overline{p}_{jm}$, for all (j, m). The likelihood of observing the price p_{ijm} as a function of θ_2 is then simply:

$$L(p_{ijm}; \theta_2, \overline{p}) \triangleq \frac{1}{\sqrt{2\pi \overline{p}_{jm}^{\zeta} \sigma^2}} \exp\left(\frac{\left(p_{ijm} - \overline{p}_{jm} - c^\top D_{R(i)jm}\right)^2}{2\overline{p}_{jm}^{\zeta} \sigma^2} \right).$$

The log-likelihood function for the observed data is then:

$$F(\theta_2, \overline{p}) \triangleq \sum_{j,m \in \mathcal{A}} \sum_{i \in \mathcal{C}_{jm}} \log L(p_{ijm}; \theta_2, \overline{p}).$$

Assuming that $F(\cdot)$ admits a unique maximum, we estimate

$$(\theta_2^*, \overline{p}^*) = \operatorname*{argmax}_{\theta_2, \overline{p}} F(\theta_2, \overline{p}).$$

The implicit optimization problem above is non-convex and high dimensional (owing to the fact that $|\mathcal{A}|$ will likely be a large set). We consider the following heuristic simplification. Observe that we must have (by the law of large numbers) that

$$\frac{1}{|\mathcal{C}_{jm}|} \sum_{i \in \mathcal{C}_{jm}} p_{ijm} \to \overline{p}_{jm}^*$$

as $|\mathcal{C}_{jm}|$ grows large. Consequently, we make the approximation

$$\overline{p}_{jm}^* \sim \frac{1}{|\mathcal{C}_{jm}|} \sum_{i \in \mathcal{C}_{jm}} p_{ijm}.$$

Since the standard error in this estimate is *substantially* smaller than any other parameters we estimate (since $|\mathcal{C}_{jm}|$ is large), we will subsequently treat $\overline{p}_{jm}^*$ as a known quantity. It will remain to estimate θ_2 as the presumed unique maximizer of $F(\theta_2, \overline{p}^*)$; this is a low (here, three) dimensional optimization problem. Solving this problem yield an estimate of our price error model. In the sequel, we will actually solve this problem jointly with an optimization problem that arises in estimating the remaining model parameters that we describe next.

3.2. Solving the Market Share and Observed Price Equations

Recall that we have the following market share equations that relate the share of product j in all sales in market m to parameters specifying our structural model for consumer utility:

$$s_{jm} = \sum_r \gamma_r \mathsf{E}_{jm} \left[\frac{\exp(u_{ijm})}{1 + \sum_{j'} \exp(u_{ij'm})} \middle| R(i) = r \right] \quad \forall j, m$$

where we recall that $u_{ijm} = -\alpha_{R(i)} \left(\overline{p}_{jm} + c^\top D_{R(i)jm} + \eta_{ijm} \right) + \beta_{R(i)}^\top X_{jm} + \xi_{jm} + \epsilon_{ijm}$. Further recall that η_{ijm} was assumed to be a zero mean normal random variable with standard deviation $\sigma \overline{p}_{jm}^{\zeta/2}$. We denote by θ_1 the collection of parameters α_r and β_r for $r = 1, \ldots, R$ and by ξ and $\overline{p}$ vectors that stack the components ξ_{jm} and $\overline{p}_{jm}$ respectively. As before, we denote by θ_2 the set of parameters σ, ζ, c. We then write the equations for market share compactly as

$$s_{jm} = \mathcal{S}_{jm}(\overline{p}, \xi, \theta_1, \theta_2) \qquad \forall j, m \tag{5}$$

Next recall the structural equations relating the expected price paid for product j in market m conditioned on j being purchased from among the products available in market m to the average price at which product j is offered in market m, $\overline{p}_{jm}$:

$$\overline{\mathbf{p}}_{\mathbf{jm}} = \overline{p}_{jm} + \frac{1}{s_{jm}} \sum_r \gamma_r \mathsf{E}_{jm} \left[\frac{\exp(u_{ijm})(c^\top D_{rjm} + \eta_{ijm})}{1 + \sum_{j'} \exp(u_{ij'm})} \middle| R(i) = r \right] \quad \forall j, m$$

Given the notation we have established, we write the above equation compactly as

$$\overline{\mathbf{p}}_{jm} = \overline{p}_{jm} + \overline{\mathcal{P}}_{jm}(\overline{p}, \xi, \theta_1, \theta_2) \qquad \forall j, m \tag{6}$$

Let $\mathcal{D}$ denote the set of values of $(\overline{p}, \xi, \theta_1, \theta_2)$ that simultaneously satisfy the market share and observed price equations, (5), (6). We make the following assumption:

Assumption 2. *If $(\overline{p}^1, \xi^1, \theta_1^1, \theta_2^1)$ and $(\overline{p}^2, \xi^2, \theta_1^2, \theta_2^2)$ are two points in $\mathcal{D}$, then $(\theta_1^1, \theta_2^1) \neq (\theta_1^2, \theta_2^2)$ if and only if $\xi^1 \neq \xi^2$.*

The above assumption can be verified in special cases; for instance if $\overline{p}$ is given, the equation (5) can be shown to have a unique solution. We will not verify the assumption here. Denote by Ξ the operator mapping values of (θ_1, θ_2) to values of ξ so that $(\overline{p}, \Xi(\theta_1, \theta_2), \theta_1, \theta_2) \in \mathcal{D}$ for some non-negative $\overline{p}$; for the remainder of this section we focus on computing this operator. We will adopt the following heuristic procedure that appeared to perform adequately for the estimation problem we faced in the present work:

1. Recall that our goal is to compute $\Xi(\theta_1, \theta_2)$ given θ_1, θ_2 and, of course, the data $s, \overline{\mathbf{p}}$. We set $\overline{p}^0 = \overline{\mathbf{p}}$.
2. In the $k+1$st iteration, set ξ^{k+1} as the unique solution to $s = \mathcal{S}(\xi^{k+1}, \overline{p}^k; \theta_1, \theta_2)$ [6].
3. Set $\overline{p}^{k+1} = \overline{\mathbf{p}} - \overline{\mathcal{P}}(\xi^{k+1}, \overline{p}^k; \theta_1, \theta_2)$.
4. If $\|p^k - p^{k+1}\|$ is sufficiently small set $\Xi(\theta_1, \theta_2) = \xi^{k+1}$; else go to step 2.

3.3. Estimating the Model

At this juncture, we recall our two identification conditions, namely:

$$\mathsf{E}\left[Z_{jm}^l \Xi(\theta_1, \theta_2)_{jm} \right] = 0$$

[6] This solution can be found via the iteration $\xi_{jm}^{i+1} = \xi_{jm}^i + \log s_{jm} - \log \mathcal{S}_{jm}(\xi^i, \overline{p}, \theta_1, \theta_2)$. which is easily shown to be a contraction mapping.

for all l, and

$$\mathsf{E}\left[\sum_{j,m\in\mathcal{A}} \nabla_{\theta_2} \log L(p_{i,j,m}, \theta_2, \overline{p}^*_{\text{aux}})\right] = 0.$$

We will proceed to estimate θ_1^* and θ_2^* via a standard GMM procedure using the empirical counterparts of the two moment conditions above. In particular, define the matrix Z with generic element $Z_{l;(jm)}$ and let us define

$$g(\theta_1, \theta_2) \triangleq \begin{bmatrix} Z\Xi(\theta_1, \theta_2) \\ \nabla_{\theta_2} F(\theta_2, \overline{p}^*_{\text{aux}}) \end{bmatrix}$$

Now the GMM procedure calls for the estimation of (θ_1^*, θ_2^*) as the optimal solution to the optimization problem

$$\min_{\theta_1,\theta_2} g'(\theta_1, \theta_2)\Phi^{-1} g(\theta_1, \theta_2) \tag{7}$$

where Φ is positive definite. So as to produce an estimator with minimal variance, the optimal choice of Φ is given by $\mathsf{E}\left[g(\theta_1^*, \theta_2^*)g'(\theta_1^*, \theta_2^*)\right]$. Of course, since this is not available to us, we employ the following standard two phase procedure:

1. Solve (7) taking Φ to be the identity matrix. Call the optimal solution (θ_1^I, θ_2^I).
2. Update the weight matrix Φ according to $\Phi = g(\theta_1^I, \theta_2^I)g'(\theta_1^I, \theta_2^I)$.
3. Solve (7) with the updated value of Φ. The optimal solution $(\tilde{\theta}_1, \tilde{\theta}_2)$ will be our estimate of (θ_1^*, θ_2^*). The (estimated) covariance matrix of this estimator is taken as $(G'\Phi^{-1}G)^{-1}$ where G if the Jacobian of g evaluated at $(\tilde{\theta}_1, \tilde{\theta}_2)$; see Newey and McFadden [1994].

3.3.1. Caveats

Having concluded our overall estimation procedure, it is worth raising a number of caveats that serve to question the validity of the procedure. For the most part, these caveats aren't specific to this particular exercise, but arise more broadly:

1. Auxiliary Model: We have conveniently assumed the first order conditions as identification conditions for the auxiliary model. Since the likelihood function there was *not* convex, it is unclear that these conditions are sufficient (and thus valid identification conditions). In order to assuage this concern, we conducted the following two stage estimation procedure: we used global optimization to solve the maximum likelihood problem (being a low dimensional problem, this approach become feasible). Using this estimate of θ_2, we estimated θ_1 via a GMM procedure much like the above but treating θ_2 as given. This resulted in essentially identical estimates with a somewhat larger variance.

2. Primary Model: We have not verified identification: we do not know whether the moment equations presented in the previous section uniquely identify θ_1. Unfortunately, this is a fairly common problem in models of this complexity. This is the most serious caveat here.

3. GMM Issues: It is unclear that the function g satisfies the conditions required for the consistency and asymptotic normality of a GMM estimator. For instance, we have essentially no understanding of the smoothness properties of Ξ. Again, this is fairly routine in the context of BLP procedures.

4. Estimation Results: Model and Welfare Gap

We present our results here. We begin with discussing our dataset. We then present our estimated model. We will then use our estimated model to estimate the various benchmarks of interest identified in Section 2 and consequently estimate welfare gaps under a variety of assumptions.

4.1. The Data

We draw chiefly on two sources of data. The first source of data is the Airline Origin and Destination Survey (DB1B), published by the US Department of Transportation (DOT). The DB1B data is a uniform 10% sample of airline 'coupons' (i.e. purchased tickets) from US domestic carriers. It provides detailed information on ticket purchase price, itineraries, ticketing carriers, flight mileage, and the number of passengers who travel on the itinerary at a given purchase price in each quarter. This data, by itself, does not suffice to build a model reflecting customer specific price dispersion; it does not give us information on products in a customer's consideration set or any information about the customer. To that end, we have also obtained a (proprietary) auxiliary data set from a major US ticketing global distribution service (GDS). This data set consists of the offer sets considered by $21,117$ distinct customers and what they eventually purchased making for a total of $172,234$ quotes. A great deal of information is available on each quote: in addition to price, we see a number of features specific to the itinerary and know the timing of the quote relative to the departure date. We next describe both data sources in greater detail.

Primary Data: Following Borenstein and Rose [1994] and Berry et al. [2006], we extract from the DB1B dataset round-trip itineraries within the continental US with at most five stops on both the outbound and return trip including origin and destination. We restrict our attention to economy class customers only. We use the data corresponding to the fourth quarter in 2006. Following Berry and Jia [2009], we eliminate purchased tickets with fare lower than $25, or with more than one ticketing carrier, or those which contain ground traffic as part of the itinerary. We focus on medium to large markets whose origin and destination airports are both located in metropolitan areas with populations (as per US Census Bureau information) exceeding $800,000$ in 2006. There are six metropolitan areas wherein each is served by more than one airport which are close to each other. We treat economy class demand at these airports as perfectly substitutable and group such airports together; we define markets based on grouped airports. This is similar to Berry and Jia [2009].

In addition to the constant, the observable product attributes X_{jm} will include the following 5 features:

- NON_STOP_{jm}: This is a dummy variable indicating whether or not product j in market m is a non-stop itinerary. Customers are likely to value a non-stop flight for a number of reasons including shorter travel time, the absence of the risk of missing a connecting flight, a perceived lower risk of lost baggage etc.

- HUB_{jm}: This is a dummy variable indicating whether or not the origin airport is a hub. Again, it is reasonable to posit that customers might value departing from a hub given that hubs offer a broader variety of services and conveniences, and will typically offer a number of alternatives in the event of a flight cancelation or if the customer misses her flight

- DISTANCE_{jm} and DISTANCE^2_{jm}: These quantities are defined as the round trip distance (and the square of this distance) for product j in market m. We hope that a combination of these two features will capture the utility customers will derive from using air travel (as opposed to slower modes of transportation) as well as any potential disutility for the long travel times associated with traveling very long distances.

- TICKETING CARRIER DUMMIES$_{jm}$: The ticketing carriers identity is a good proxy for a number of issues a customer is likely to consider important, including for instance the airlines reputation on a particular route.

Table 1 summarizes the primary dataset.

Auxiliary Data: Our auxiliary dataset consists of 172,234 choice sets considered by purchasing customers in the fourth quarter of 2006 along with information on what was eventually purchased. The data includes, for each quote in the consideration set, information on the fare (such as price, itinerary, potential restrictions etc.), the market, and the date of travel. In addition we know the date at which the offer set was considered. Based on this auxiliary data, we consider defining two customer types:

1. Type 1 Customers: In revenue-management speak, these would be termed 'leisure' customers. This set includes customer who make their purchases at least 8 days prior to departure.

2. Type 2 Customers: These are customers who make their purchase within 7 days of departure. In RM speak, these would be considered 'business' customers.

We believe this division of customers is perhaps most meaningful given airline revenue management practice. Moreover, a type specific pricing policy would then simply translate to a dynamic pricing policy which is standard practice. $D_{R(i)jm}$ are then dummy variables for the combination of customer type and whether or not the product j purchased consists of a non-stop itinerary. In particular,

$$D_{R(i)jm} = \begin{pmatrix} \mathbb{I}[\text{purchase time}_i \leq 7 \text{ days and } NON\_STOP_{jm} = 1] \\ \mathbb{I}[\text{purchase time}_i \leq 7 \text{ days and } NON\_STOP_{jm} = 0] \end{pmatrix}.$$

(Note that dummy's for the remaining combinations are excluded to prevent collinearity).

Table 2 summarizes key features of the auxiliary dataset. Notice that the average fares (corresponding to offered prices) in the auxiliary data set are higher than those in the primary data set (corresponding to purchase prices) as one might expect. It is also worth recalling at this point that the only assumption we must make in using the auxiliary data set in our estimation is that the structural model describing price discrimination is consistent across markets. We need not make any assumptions on the similarity of the populations in the main and auxiliary data set.

Table 1: Summary Statistics for the Primary Dataset.

	Mean	Std.
Fare (\$100)	4.39	3.51
NON_STOP	0.59	0.49
Distance (10^3 miles)	2.86	1.40
No. Markets	3,125	
No. Products	17,737	
No. Observations	261,151	
Per Market		
Mean Fare (\$100)	3.84	1.12
POP (10^7)	3.35	2.30
Mean NON_STOP	0.26	0.40
Mean Distance (10^3 miles)	2.86	1.46
No. Passengers (10^3)	24.77	55.54
No. Products	5.67	7.29
Per Product		
Mean Fare (\$100)	4.85	3.11
NON_STOP	0.09	0.29
Distance (10^3 miles)	3.45	1.57
No. Passengers (10^3)	4.37	25.67

Table 2: Summary Statistics for the Auxiliary Dataset.

	Mean	Std.
Fare (\$100)	6.17	3.35
NON_STOP	0.46	0.51
Fraction Type One	0.88	0.32
Fraction Type Two	0.12	0.32
No. Products/Market	8.68	10.31
No.Tickets/Market and Product	521.92	810.76
No. Markets	38	
No. Observations	172,234	

4.2. Estimated Model

We first discuss the price dispersion model estimated via our auxiliary dataset; recall briefly that the model took the following form: we posited that the price offered to customer i for product j in market m, p_{ijm} depends on characteristics of the product j and the customer i according to

$$p_{ijm} = \overline{p}_{jm} + c^\top D_{R(i)jm} + \eta_{ijm}$$

where $\overline{p}_{jm}$ is the quoted price for product j averaged over the population of consumers in market m; $D_{R(i)jm}$ is a vector capturing features of the customer and product whose mean, averaged over the population, is zero; and η_{ijm} is an independent normal random variable with mean zero and variance $\sigma^2(\overline{p}_{jm})^\zeta$. The features we used here were dummies for a combination of whether or not the customer was a leisure or business traveler (as determined by time of purchase) and whether or not the itinerary was non-stop.

Table 3: Estimated Price Dispersion Model.

$\mathbb{I}$[purchase time ≤ 7 days and $NON\_STOP = 1$]	1.78	(0.26)
$\mathbb{I}$[purchase time ≤ 7 days and $NON\_STOP = 0$]	1.29	(0.25)
σ	0.31	(0.04)
ζ	2.09	(0.13)

To summarize, we establish the following facts about the nature of price dispersion as a result of RM practices, all of which are in line with what one might expect:

1. Late purchasers pay a substantial premium. This is precisely what we would expect from a revenue management standpoint. For an econometric perspective, see Puller et al. [2008].

2. The premium is higher on non-stop routes, even in relative terms. This observation agrees with similar qualitative conclusions drawn by Berry and Jia [2009].

3. Price dispersion is heteroskedastic in nature; in fact since $\zeta \sim 2$ and $\sigma \sim 0.33$, we conclude that price shocks are of a magnitude roughly proportional to price. This is in excellent agreement with the landmark work of Borenstein and Rose [1994].

Having estimated our price dispersion model, we use the modified BLP procedure in Section 3 to estimate the remaining parameters of our random utility model. Recall that this model took the form

$$u_{ijm} = -\alpha_{R(i)} \left(\overline{p}_{jm} + c^\top D_{R(i)jm} + \eta_{ijm} \right) + \beta_{R(i)}^\top X_{jm} + \xi_{jm} + \epsilon_{ijm}, \ \forall i, j \in m, m.$$

where the features $D_{R(i)jm}$ and X_{jm} and the two types of customers we assume were described in Section 4.1. Of the features constituting X_{jm} the estimates of the coefficients of β for the HUB and carrier identity dummy, as well as for Distance and squared distance were not significantly different between the two customer types, and as such we report one set of coefficients for each of these features. Table 4 describes the learned model.

Table 4: Estimates in Customer Demand Model.

α_1 (Leisure)	0.80	(0.13)
NON_STOP1	1.69	(0.34)
α_2 (Business)	0.09	(0.03)
NON_STOP2	2.23	(0.46)
HUB	0.21	(0.07)
DISTANCE (10^3 miles)	0.77	(0.27)
DISTANCE^2 (10^6 miles^2)	-0.10	(0.04)
γ_1	0.78	(0.05)
γ_2	0.22	(0.05)
AMERICAN	0.41	(0.13)
CONTINENTAL	0.18	(0.05)
DELTA	0.11	(0.03)
JETBLUE	0.25	(0.09)
NORTHWEST	-0.07	(0.17)
SOUTHWEST	0.08	(0.04)
UNITED	0.25	(0.08)
USAIR	-0.16	(0.10)

The estimated model above captures a number of key features we might anticipate for air travelers *remarkably* well:

1. **Price Sensitivity:** Early purchasers are more sensitive to prices as witnessed by the price co-efficient estimated for type 1 (early) vs. type 2 (late) customers. The computed price elasticities (-1.27 for type 1 and -0.53 for type 2) tell a similar story [7]. This is in line with expectations; modern revenue management practices operate on the premise that later customers typically represent relatively inelastic demand. Moreover, the aggregate elasticity estimate is in excellent agreement with other studies compiled in Borenstein and Rose [2007].

2. **Value placed on Non-stop flights:** Type 1 customers appear to place a smaller premium on non-stop flights. Again, this is in line with the revenue management view of type 1 and type 2 customers as being 'leisure' and 'business' travelers predominantly. The difference in the coefficients for other features were not significant.

3. **Carrier Specific Effects:** We see that customers place a premium on certain carriers (such as American, United and JetBlue) and incur a disutility from other carriers (US Airways). Interestingly, this is roughly in line with the 2006 airline quality survey results Bowen and Headley [2007], wherein Continental, United and American were the preferred large carriers, JetBlue was the top rated low cost carrier and US Airways was ranked last.

4. **Distance:** The estimated model shows that consumer utility is a concave function of distance. Specifically, the coefficients estimated for DISTANCE and its square are such that the utility placed on distance traveled is increasing and concave in distance traveled up to approximately 3800 miles which covers the vast majority of domestic routes. This dependence on distance

[7] The aggregate price elasticity that measures change in total demand when the prices per unit percentage increase in all prices is -1.07.

is in excellent agreement with the Standard Industry Fare Level (SIFL) formula that is used to place a 'fair value' on air travel as a function of distance (see SIFL [2012]).

5. **Types Matter:** We can conduct all of our estimation under the assumption of a single customer type. In doing so, we will estimate an aggregate price elasticity of -5 which is substantially larger than any available estimates of this quantity. We see this as a strong sign that accounting for customer types is, in fact, quite important.

4.3. Sensitivity To Assumptions In Reduced Form Pricing Model

Recall that we used the following reduced form model to describe price dispersion across customers:

$$p_{ijm} = \overline{p}_{jm} + c^{\top} D_{R(i)jm} + \eta_{ijm}$$

where η_{ijm} is an independent normal random variable with mean zero and variance $\sigma^2(\overline{p}_{jm})^{\zeta}$. This section considers three variants to this model to understand the impact of the various modeling assumptions made in the above model. We also consider a fourth model variant that ignores customer heterogeneity altogether. In particular, we consider the following models:

1. Model I: The features used in the price dispersion model were dummies of whether the customer was a leisure or business traveler (as determined by time of purchase). In particular, we ignore the potential for the extent of price dispersion on non-stop flights being higher than on flights with multiple stops. This model is closest to out original model.

2. Model II: We do not allow for heteroskedasticity in price dispersion. This is a crucial feature with a longstanding historical precedent (Borenstein and Rose [1994]).

3. Model III: Non-existence of price dispersion (i.e. $c = 0$ and $\sigma = 0$). In effect this ignores airline revenue management altogether; we know is inter-temporal price discrimination is an integral part of airline pricing.

4. Model IV: No customer heterogeneity. This is, in effect closely related to Model IV. In addition to ignoring price discrimination, we ignore heterogeneity in customer tastes itself.

We anticipate models II, III and IV will lead to implausible results, reinforcing the need for the elements of the original model that are done away with in those respective models.

NOTE: III-I, I-II,IV-III,V-IV

Estimation results are reported in the Table 5 in the Appendix. [8].

[8] Estimated parameters c_1 and c_2 in Table 5 are coefficients for variables $\mathbb{I}[\text{purchase time} \leq 7 \text{ days and } NON\_STOP = 1]$ and $\mathbb{I}[\text{purchase time} \leq 7 \text{ days and } NON\_STOP = 0]$ respectively.

Table 5: Estimates of Different Specifications of Customer Demand Model and Price Dispersion Model.

	Benchmark	I	II	III	IV
Fare 1 ($100)	0.80	0.90	1.08	1.13	0.98
	(0.13)	(0.15)	(0.17)	(0.14)	(0.23)
Constant 1	-7.86	-7.94	-7.54	-6.98	-9.13
	(1.27)	(1.63)	(1.88)	(2.54)	(1.08)
NON_STOP 1	1.69	1.74	1.85	2.19	2.55
	(0.34)	(0.28)	(0.51)	(0.29)	(0.31)
Fare 2 ($100)	0.09	0.07	0.06	0.09	-
	(0.03)	(0.02)	(0.03)	(0.04)	-
Constant 2	-9.87	-9.84	-10.62	-9.92	-
	(2.37)	(2.11)	(3.15)	(2.53)	-
NON_STOP 2	2.23	2.25	2.50	2.57	-
	(0.46)	(0.37)	(0.69)	(0.43)	-
HUB	0.21	0.22	0.25	0.17	0.23
	(0.07)	(0.04)	(0.09)	(0.08)	(0.05)
Distance (10^3 miles)	0.77	0.75	0.85	0.90	1.90
	(0.27)	(0.25)	(0.32)	(0.28)	(0.31)
Distance^2 (10^6 miles^2)	-0.10	-0.10	-0.13	-0.11	-0.17
	(0.04)	(0.04)	(0.03)	(0.03)	(0.02)
γ_1	0.78	0.79	0.79	0.73	1.00
	(0.05)	(0.03)	(0.05)	(0.04)	-
γ_2	0.22	0.21	0.21	0.27	0.00
	(0.05)	(0.03)	(0.05)	(0.04)	-
c_1	1.78	1.56	1.84	0.00	0.00
	(0.26)	(0.41)	(0.36)	-	-
c_2	1.29	1.56	1.44	0.00	0.00
	(0.25)	(0.41)	(0.44)	-	-
σ	0.31	0.34	2.57	0.00	0.38
	(0.04)	(0.06)	(0.93)	-	(0.03)
ζ	2.09	2.16	0.00	0.00	2.11
	(0.13)	(0.19)	-	-	(0.17)

To see the implications of each of the four modeling alternatives estimated, we turn to a single key summary statistic, namely price elasticity. We measure price elasticity both for each customer type as also in the aggregate. The reason we choose to focus on elasticity is the plethora of econometric studies that have made elasticity estimates available over the years.

To examine the validity of above specifications, we compute the price elasticity of above models. As reported by Berry and Jia [2009], Gillen conducted a survey that collected 85 demand elasticity estimates from cross-sectional studies with a median of 1.33; the preponderance of these studies, allows us to gauge the plausibility of an elasticity estimate obtained via any of the alternate modeling choices discussed above. Table 6, computes elasticity estimates obtained under the four alternate models discussed in this section:

Table 6: Price Elasticity of Different Specifications of Customer Demand Model and Price Dispersion Model.

	Benchmark	I	II	III	IV
Type One	-1.27	-1.41	-5.17	-5.82	-2.98
Type Two	-0.53	-0.53	-0.49	-0.46	-
Aggregate	-1.09	-1.25	-1.92	-1.86	-2.98

In summary, we see that, as expected Model I yields results similar to those obtained via our benchmark model. Models II and III which do not capture key feature of priced dispersion yield results that while not egregious reflect a substantially large price elasticity than those computed in other studies for leisure travelers, as well as on the aggregate. Model IV seems altogether implausible with an elasticity estimate that is twice that reported in the summary of estimates provided in Berry and Jia [2009]. We see this as some support for the modeling assumptions we have chosen to make: doing away with them appears to yield implausible results.

4.4. Results: Optimal Welfare Benchmarks

Having estimated our model, we may now proceed to estimate the benchmarks we posited in Section 2. These results are summarized in Table 7:

Table 7: Optimal Welfare Benchmarks: Third Quarter, 2006

		Status Quo	$OPT(SW1)$	$OPT(SW2)$
CS($billion)	Type One	6.12	4.71	4.59
	Type Two	32.56	45.58	46.66
	Aggregate	38.67	50.29	51.25
	Gain	-	11.62	12.58
	Relative Gain (%)	-	30.04	32.53
Revenue ($billion)	Type One	12.51	11.67	11.93
	Type Two	14.68	11.07	10.89
	Aggregate	27.19	22.74	22.82
	Gain	-	-4.45	-4.37
	Relative Gain (%)	-	-16.37	-16.08
SW ($billion)	Aggregate	65.87	73.03	74.07
	Gain	-	7.16	8.21
	Relative Gain (%)	-	10.88	12.46
Demand(10^3/market)	Type One	15.75	12.13	11.83
	Type Two	9.04	12.66	12.96
	Aggregate	24.79	24.79	24.79
Avg Purchasing Price ($100)	Type One	3.21	4.11	3.42
	Type Two	5.89	4.11	3.42
	Aggregate	4.85	4.11	3.42

The table above provides estimates that answer the questions posed by this paper. Let us begin with characterizing the *status quo*. We see that:

1. By far, the largest share of consumer surplus obtains from business travelers. These travelers

account for approximately 84% of the surplus obtained by *all* air travelers. This is not surprising given that 'business' travel likely facilitates a substantially broader variety of economic output than does leisure travel.

2. Airlines obtain a larger (greater that 50%) share of their revenues from business travelers. Business travelers pay, on average, nearly twice as much as their leisure counterparts. Again, this is perhaps to be expected and highlight the importance (or, at the least, impact), of airline revenue management practices.

3. In spite of the above, the fraction of airline resources, i.e. seats, occupied by leisure travelers is more than twice that of business travelers.

Our welfare benchmarks give us a sense of how inefficient the status quo is, and where these inefficiencies arise from. In particular, we consider the two benchmarks in turn.

Inefficiencies due to Monopoly Power: This is captured best by the gains shown by the $OPT(SW1)$ benchmark. Recall that benchmark re-allocated seats between leisure and business travelers in a manner that maximized allocative efficiency. The ability to do this leads to a net increase of 11 billion dollars in consumer surplus – a 30% increase. This is a *dramatically* large number. While revenues are clearly hurt under this benchmark, the aggregate increase in welfare net of changes in revenue is approximately 7 billion dollars a quarter. This is again a very large sum. We attribute these efficiencies to market power – specific routes tend to be highly concentrated (monopolies and duopolies). The appendix describes a stylized model where qualitatively similar impacts arises due to monopoly power.

Inefficiencies due to Resource Allocation: An important reason for airline deregulation was the desire to rationalize the allocation of airline seats across routes in the airline network. $OPT(SW2)$ which to an extent allows transfers of seats across routes, captures some of the inefficiencies that remain due to an inefficient allocation of airline capacity across routes. In particular, this inefficiency was valued at an additional billion dollars of welfare loss per quarter in 2006 above the welfare losses described above. In reality, these losses are likely to be substantially larger: recall that our benchmark did not allow us to change the net resource availability on a given leg on the airline network.

5. Conclusion

To answer the question this paper posed, the airline industry remains *very* inefficient from a welfare standpoint. The welfare loss was at least eight billion dollars a quarter in 2006. The size of this loss is remarkable when compared with the revenue gains airlines have been focused on eking out through tactical operational changes in recent years.

This paper provides strong support to at least two broad directions in which economic policy and airline RM policies can eventually have impact. First, monopoly pricing power continues to remain a substantial friction on the industry. Since hub-spoke network structures naturally lend themselves to such concentration, incentive design that reduces monopoly pricing power and concentration on specific routes is a problem very worthy of our attention. In a different direction, a substantial welfare gain can probably be obtained from re-allocation of airline capacity across US routes. These capacity decisions evolve relatively slowly and potentially represent a quicker (albeit smaller) opportunity for improvement.

Turning toward our approach, the use of a reduced form model for price discrimination allowed us to incorporate an important feature of airline revenue management practices into this study using the sparse data we had available for the task. In future studies, it would be interesting to develop a more careful structural model to describe price discrimination across customer classes. Of course, given the complexity of airline RM practices, and the relative unavailability of relevant data, this is unlikely to be an easy problem, either from a data collection standpoint or from a modeling and estimation standpoint.

References

M. Akan, B. Ata, and J. Dana. Revenue management by sequential screening. *Working Paper*, 2011.

O. Armantier and O. Richard. Domestic airline alliances and consumer welfare. *RAND Journal of Economics*, 39(3):875–904, 2008.

G. Bamberger, D. Carlton, and L. Neuman. An empirical investigation of the competitive effects of domestic airline alliances. *Journal of Law and Economics*, 47(1):195–222, 2004.

S. Berry and P. Jia. Tracing the woes: An empirical analysis of the airline industry. *Working Paper*, 2009.

S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

S. Berry, J. Levinsohn, and A. Pakes. Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68–105, 2004.

S. Berry, M. Carnall, and P. Spiller. Airline hubbing: Costs, and demand. *in Advances in Airline Economics Vol. 1: Competition Policy and Antitrust*, D. Lee, ed. Elsevier Press:183–214, 2006.

S. Borenstein. Hubs and high fares: Airport dominance and market power in the u.s. airline industry. *RAND Journal of Economics*, 20:44–65, 1989.

S. Borenstein. The dominant firm advantage in multiproduct industries: Evidence fomr u.s. airlines. *Quarterly Journal of Economics*, 106:1237–1266, 1991.

S. Borenstein. U.s. domestic airline pricing, 1995-2004. *UC Berkeley Competition Policy Center working paper*, No. CPC05-48, 2005.

S. Borenstein and N. Rose. Competition and price dispersion in the u.s. airline industry. *Journal of Political Economy*, 102:653–683, 1994.

S. Borenstein and N. Rose. How airline markets work... or do they? regulatory reform in the airline industry. *NBER working paper*, No. 13452, 2007.

B. Bowen and D. Headley. Airline quality rating. *Website: http://www.airlinequalityrating.com/reports/2007aqr.pdf*, 2007.

F. Ciliberto and E. Tamer. Market structure and multiple equilibria in airline markets. *Econometrica*, 77:1791–1828, 2009.

J. Dana. Using yield management to shift demand when the peak time is unknown. *RAND Journal of Economics*, 30(Autumn):456–474, 1999a.

J. Dana. Equilibrium price dispersion under demand uncertainty: The roles of costly capacity and market structure. *RAND Journal of Economics*, 30(Winter):632–660, 1999b.

J. Dana and E. Orlov. Internet penetration and capacity utilization in the us airline industry. *Working Paper*, 2009.

J-P. Dube, J. Fox, and C-L. Su. Improving the numerical performance of blp static and dynamic discrete choice random coefficients demand estimation. *Working Paper*, 2009.

S. Forbes. The effect of air traffic delays on airline prices. *International Journal of Industrial Organization*, 26(5):1218–1232, 2008.

I. Gale and T. Holmes. The efficiency of advance-purchase discounts in the presence of aggregate demand uncertainty. *International Journal of Industrial Organization*, 10(3):413–437, 1992.

A. Goolsbee and C. Syverson. How do incumbents respond to the threat of entry? evidence from the major airlines. *The Quarterly Journal of Economics*, 123(4):1611–1633, 2008.

J. Li, N. Granados, and S. Netessine. Are consumers strategic? structural estimation from the air-travel industry. *Working Paper*, 2011.

D. McFadden. Modelling the choice of residential location. *in Spatial Interaction Theory and Planning Models*, A. Karlgvist, et al., eds. Amsterdam: North-Holland, 1978.

S. Morrison and C. Winston. *The Evolution of the Airline Industry*. Brooings Institution, Washington, DC, 1995.

S. Netessine and R. Shumsky. Revenue management games: Horizontal and vertical competition. *Management Sci.*, 51 (5):813–831, 2005.

A. Nevo. A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics and Management Strategy*, 9(4):513–548, 2000.

W. Newey and D. McFadden. Large sample estimation and hypothesis testing. *Handbook of Econometrics*, IV, edited by R.F. Engle and D.L. McFadden. Elsevier Science, 1994.

J. Park. The effects of airline alliances on markets and economic welfare. *Transportation Research*, 33(3):181–195, 1997.

J. Park and A. Zhang. Airline alliances and partner firms' output. *Transportation Research*, 34: 245–255, 1998.

J. Park and A. Zhang. An empirical analysis of global airline alliances: Cases in north atlantic markets. *Review of Industrial Organization*, 16:367–384, 2000.

C. Peter. Evaluating the performance of merger simulation: Evidence from the u.s. airline industry. *The Journal of Law and Economcis*, 49(2):627–649, 2006.

S. Puller, A. Sengupta, and S. Wiggins. Testing theories of scarcity pricing and price dispersion inthe airline industry. *Working Paper*, 2008.

SIFL. Standard industry fare level. *Website: http://ostpxweb.dot.gov/aviation/domfares/siflb.pdf*, 2012.

K. Talluri and G. van Ryzin. *The Theory and Practice of Revenue Management.* Springer Science+Business Media, 2004.

G. Vulcano, G. van Ryzin, and C. Maglaras. Optimal dynamic auctions for revenue management. *Management Sci.*, 48 (11):1388–1407, 2002.

J. Wirtz, S. Kimes, J. Theng, and P. Patterson. Revenue management: Resolving potential customer conflicts. *Journal of Revenue and Pricing Management*, 2(3):216–226, 2003.

A. Social Welfare Model

We consider a market with two type customers and one product. We denote $D_r(p_r)$ as the demand function of type r customer. We assume $D_r(\cdot)$ is a non-increasing and non-negative function. We assume D_r is invertible, i.e., for $d_r = D_r(p_r)$, there exists a function $P_r(\cdot)$ such that $p_r = P_r(d_r)$. We use C to denote total capacity. Since we are only interested in the case with limited resource, we assume that $C \leq \sum_{r=1}^{2} D_r(0)$.

The social planner solves the following problem (PSW):

$$\begin{aligned} \max_{D_1, D_2} \quad & \sum_{r=1}^{2} \int_0^{D_r} P_r(D) dD \\ \text{s.t.} \quad & \sum_{r=1}^{2} D_r \leq C. \end{aligned}$$

We use script SW to denote the optimal solution of social planner's problem. From KKT condition, we have $p_1^{SW} = p_2^{SW} \triangleq p^{SW}$. It is determined by solving the following problem:

$$\sum_{r=1}^{2} D_r(p_r^{SW}) = C.$$

We restrict our attention to the deterministic systems where price is fixed given a customer type, market, and product. We compare some key variables in the optimal and generic deterministic systems. We use script c to denote the variables in the current deterministic system. The results are shown in the following lemma.

Lemma 1. *For the current deterministic system with two-type customers. Suppose the pricing policy satisfies $p_1^c \leq p_2^c$, and we define $C = \sum_{r=1}^{2} D_r(p_r^c)$. Then we have following results:*

1. $p_1^c \leq p_1^{SW} = p_2^{SW} \leq p_2^c$.
2. $D_1^c \geq D_1^{SW}, D_2^c \leq D_2^{SW}$.
3. $CS_1^c \geq CS_1^{SW}, CS_2^c \leq CS_2^{SW}$.

4. $SW_1^c \geq SW_1^{SW}, SW_2^c \leq SW_2^{SW}$.

Proof. 1. As we discussed above, we note that in the social planner's problem, $p_1^{SW} = p_2^{SW}$. Since $\sum_{r=1}^{2} D_r(p_r^c) = \sum_{r=1}^{2} D_r(p_r^{SW})$, and the demand function $D_r(p)$ is non-increasing in p, then we have $p_1^c \leq p_1^{SW} = p_2^{SW} \leq p_2^c$.

2. Since $D_r(p)$ is non-increasing in p, and $p_1^c \leq p_1^{SW} = p_2^{SW} \leq p_2^c$, then $D_1^c \geq D_1^{SW}, D_2^c \leq D_2^{SW}$.

3. Since $CS_r = \int_0^{D_r} (P_r(D) - P_r(D_r))dD$ is increasing in D_r, and $D_1^c \geq D_1^{SW}, D_2^c \leq D_2^{SW}$, then $CS_1^c \geq CS_1^{SW}, CS_2^c \leq CS_2^{SW}$.

4. Since $SW_r = \int_0^{D_r} P_r(D)dD$ is increasing in D_r, and $D_1^c \geq D_1^{SW}, D_2^c \leq D_2^{SW}$, then $SW_1^c \geq SW_1^{SW}, SW_2^c \leq SW_2^{SW}$.

■