The Welfare Effects of Intertemporal Price Discrimination:

An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets

John Lazarev*

Graduate School of Business

Stanford University

Job Market Paper

This version: November 23, 2011†

Abstract

This paper studies how a firm's ability to price discriminate over time affects production, product quality, and product allocation among consumers. The theoretical model has forward-looking heterogeneous consumers who face a monopoly firm. The firm can affect the quality and quantity of the goods sold each period. I show that the welfare effects of intertemporal price discrimination are ambiguous. I use this model to study the time paths of prices for airline tickets offered on monopoly routes in the U.S. Using estimates of the model's demand and cost parameters, I compare the welfare travelers receive under the current system to several alternative systems, including one in which free resale of airline tickets is allowed. I find that free resale of airline tickets would increase the average price of tickets bought by leisure travelers by 54% and decrease the number of tickets they buy by 10%. Their consumer surplus would decrease by only 16% due to a more efficient allocation of seats and the opportunity to sell a ticket on a secondary market.

*I thank Lanier Benkard and Peter Reiss for their invaluable guidance and advice. I am grateful to Tim Armstrong, Jeremy Bulow, Liran Einav, Alex Frankel, Ben Golub, Michael Harrison, Jakub Kastl, Jon Levin, Trevor Martin, Michael Ostrovsky, Mar Reguant, Andrzej Skrzypacz, Alan Sorensen, Bob Wilson, Ali Yurukoglu and participants of the Stanford Structural IO lunch seminar for helpful comments and discussions. All remaining errors are my own.

Correspondence: jlazarev@gsb.stanford.edu

†The latest version is available at `http://www.stanford.edu/~jlazarev/Lazarev_JMP.pdf`

1 Introduction

This paper estimates the welfare effects of intertemporal price discrimination using new data on the time paths of prices from the U.S. airline industry. Who wins and who loses as a result of this intertemporal price discrimination is an important policy question because ticket resale among consumers is explicitly prohibited in the U.S., ostensibly for security reasons. Some airlines do allow consumers to "sell" their tickets back to them, but they also impose fees that can make the original ticket worthless. Just what motivates these practices is a matter of public debate[1]. Economic theory suggests that secondary markets are desirable because they facilitate more efficient reallocations of goods. Yet the existence of resale markets would also frustrate airlines' ability to price discriminate over time, which could potentially decrease overall social welfare.

Theoretically, the welfare effects of price discrimination are ambiguous (Robinson, 1933). I focus on three channels through which price discrimination can affect social welfare. First, price discrimination changes the quantity of output sold as some buyers face higher prices and buy less, while other buyers face lower prices and buy more.[2] Second, price discrimination can affect the quality of the product (Mussa and Rosen, 1978). For instance, a firm may deliberately degrade the quality of a lower-priced product to keep people willing to pay a higher price from switching to the lower-priced product (Deneckere and McAfee, 1996). Finally, price discrimination can result in a misallocation of products among buyers. Since consumers potentially face different prices, it is not necessarily true that the customers willing to pay the most for the product will end up buying it.

Empirically, we know little about the costs and benefits of intertemporal price discrimination.[3] There are several reasons why there has been little work on this problem. First, there is a lack of available data. In the airline industry, price and quantity data that are necessary to estimate demand have been available to researchers only at the quarterly level. Such data do not allow one to separate intertemporal discrimination for a given departure date from variation due to different days of departure. McAfee and te Velde (2007) used a sample of price paths but they did not have

[1] Consumer advocates speak out against these inflexible policies and question the legality of such practices. If you buy a ticket, they argue, it's your property and you should be able to use it any way you want, including giving it to a friend or selling it to a third party. For examples, see Bly (2001), Curtis (2007), and Elliot(2011).

[2] An increase in total output is a necessary condition for welfare improvement with third-degree price discrimination by a monopolist. Schmalensee (1981), Varian (1985), Aguirre et al (2010), and others have analyzed these welfare effects in varying degrees of generality.

[3] Exceptions include Hendel and Nevo (2011), Nair (2007).

access to the corresponding quantities. I solve this problem by merging daily price data collected from the web with quarterly quantity data using a structural model.

A second impediment to studying intertemporal price discrimination is that a structural model of dynamic oligopoly with intertemporal price discrimination would necessarily be very complicated. Among other difficulties, one would have to deal with the multiplicity of equilibrium predictions. I avoid these problems by focusing solely on monopoly routes. Finally, I use institutional details of the way that prices are set in practice in the industry to simplify the problem even further.

While I do observe the lowest available price on each day prior to departure, I only observe the quantity of tickets purchased at each price on a quarterly basis. As a result, it would be difficult to estimate demand and cost parameters directly. Instead, I estimate the parameters of consumers' preferences indirectly, based on a model of optimal fares. In the model, a firm sells a product to several groups of forward-looking consumers during a finite number of periods. Consumer groups differ in three ways: what time they arrive in the market, how much they are willing to pay for a flight, and how certain they are about their travel plans. The firm cannot charge different prices to different consumer groups but is able to charge different prices in different periods of sale. There is no aggregate demand uncertainty.[4] Under these assumptions, I show that a set of fares with positive cancellation fees and advance purchase requirements maximizes the firm's profit. By contrast, the market-clearing fare without advance purchase requirements or cancellation fees maximizes the social welfare defined as the sum of the airline's profit and consumers' surplus.

For each value of the unknown parameters, my model predicts a unique profit-maximizing path of fares as well as the corresponding quantities of tickets sold. I match these predictions with data collected from 76 U.S. monopoly routes. For every departure date in three quarters, I recorded all public fares published by airlines for six weeks prior to departure. Since quantity data are not publicly available, I use the model of optimal fares to predict quantities sold at each price level in each period. I then aggregate these predictions to the quarterly level and match them to data from the well-known quarterly sample of airline tickets. To estimate demand and cost parameters, I use a two-step generalized method of moments based on restrictions for daily prices, monthly quantities

[4] Aggregate demand uncertainty is another reason why an airline facing capacity constraints may benefit from varying its prices over time (Gale and Holmes, 1993, Dana, 1999). Puller et al (2009) found only modest support for the scarcity pricing theories in the ticket transaction data, while price discrimination explained much of the variation in ticket pricing.

and the quarterly distribution of tickets derived from the model of optimal fares.

For markets in my data sample, the estimates suggest that, on average, 76% of passengers travel for leisure purposes. A significant share of leisure travelers start searching for a ticket at least six weeks prior to departure. By contrast, 83% of business travelers begin their search in the last week. Business travelers are willing to pay up to six times more for a seat and they are less price-elastic. Business travelers tend to avoid tickets with a cancellation fee as the probability that they have to cancel a ticket is higher.

These estimates allow me to assess the welfare effects of intertemporal price discrimination. Compared to an ideal allocation that maximizes social welfare, the profit-maximizing allocation results in a 21% loss of the total gains from trade. To understand to what extent intertemporal price discrimination contributes to this loss, I use the estimates to calculate the equilibrium sets of fares for three alternative designs of the market.

The first scenario assesses the potential benefits and costs of allowing unrestricted airline ticket resale.[5] I model resale by assuming that there are an unlimited number of price-taking arbitrageurs who can buy tickets in any period in order to resell them later. Under this assumption, the profit-maximizing price path is flat. The welfare effects of a secondary market, however, are ambiguous. On the one hand, the secondary market increases the quality of tickets and eliminates misallocations among consumers. On the other hand, the secondary market can – and, for the markets I consider, does – reduce the total quantity of tickets sold in the primary market. I find that the average price of tickets bought by leisure travelers would increase from \$77 to \$118, and the number of tickets they buy would decrease by 10%. However, business travelers would face an average price decrease from \$382 to \$118, with quantity increasing by 49%. The consumer surplus of leisure travelers would decline by 16%, the consumer surplus of business travelers would increase by almost 100%, and the airline's profit would decrease by 28%. Overall, social welfare on the average route would increase by 12%, even though the total quantity of tickets sold would go down.

In a second scenario, I return to a market without resale and assume that the monopolist is not allowed to alter the quality of tickets by imposing a cancellation fee but can still charge different prices in different periods. I find that the monopolist would still discriminate over time but the

[5] Recent empirical literature on resale and the welfare effects of actual secondary markets includes Leslie and Sorensen (2009), Sweeting (2010), Chen et al (2011), Esteban and Shum (2007), Gavazza et al (2011). Ticket resale is explicitly prohibited in the U.S. airline industry.

equilibrium price path would become flatter, which would reduce misallocations of tickets among consumers. The average ticket price would go up from \$137 to \$157. Leisure travelers would benefit due to the increase in the quality of tickets but would lose from the increase in prices. The net effect on their consumer surplus would be still positive. Overall, social welfare would slightly increase.

Finally, to illustrate the welfare properties of the cancellation fee, I calculate profit-maximizing fares for different values of the fee. Most U.S. airlines charge the same cancellation fee in all domestic markets they serve. For the markets I study, this fee often exceeds the price of a ticket, significantly decreasing the tickets' quality. The estimation results show that there exists a lower cancellation fee that would increase the consumer surplus of each traveler group while leaving the airline's profit unchanged.

The paper informs three important empirical literatures. First, it contributes to the empirical price discrimination literature. Shepard (1991) considered prices of full and self service options at gas stations. Verboven (1996) studied differences in automobile prices across European countries. Leslie (2004) quantified the welfare effects of price discrimination in the Broadway theater industry. Villas-Boas (2009) analyzed wholesale price discrimination in the German coffee market. Second, it connects to empirical studies of durable goods monopoly. Nair (2007) estimated a model of intertemporal price discrimination for the market of console video games. Hendel and Nevo (2011) estimated that intertemporal price discrimination in storable goods markets increases total welfare. This paper arrives at a different conclusion for airline tickets. Finally, there are several related papers that analyze price dispersion in the U.S. airline industry (Borenstein and Rose, 1994, Stavins, 2001, Gerardi and Shapiro, 2009). To the best of my knowledge, this is the first paper to empirically estimate the welfare effects of intertemporal price discrimination in the airline industry.

The rest of the paper proceeds as follows. Section 2 gives background information on airline pricing. Section 3 presents a model of optimal fares. Section 4 describes the data used in the analysis. Section 5 shows how to use the model of optimal fares to infer demand and supply parameters from the collected data. Section 6 presents the results of estimation. Section 7 formally describes the alternative market designs and present the results of counterfactual simulations. Section 8 concludes.

2 Institutional Background

An airline can start selling tickets on a scheduled flight as early as 330 days before departure. At any given moment, the price of a ticket is determined by the decisions of two airline departments, the pricing department and the revenue management department. The pricing department moves first and develops a discrete set of fares that can be used between two airports served by the airline. The revenue management department moves second and chooses which of the fares from this set to offer on a given day.

The pricing department offers fares with different "qualities" to discriminate between leisure and business travelers. High-quality fares are unrestricted. Low-quality fares come with a set restrictions such as advance purchase requirements (APR) and cancellation fees. To secure cheaper fares, a traveler typically has to buy a ticket early, usually a few weeks before her departure date. If her travel plans later change, she may have to pay a cancellation fee, which often could make the purchased ticket worthless. These restrictions exploit the fact that business travelers are usually more uncertain about their travel plans than leisure travelers.

Figure 1 gives a snapshot of all coach-class fares that were published by American Airlines' pricing department for Dallas, TX – Roswell, NM flights departing on March 1st, 2011, six weeks prior the departure. Fares with advance purchase requirements include a cancellation fee of $150. Fares without advance purchase requirements are fully refundable.

The fact that the pricing department has published a fare does not imply that a traveler will be able to get that fare on the specific flight. The flight needs to have available seats in the booking class that corresponds to that fare. How many seats to assign to each booking class in each flight is the primary decision of the revenue management department.

Figure 2 shows the paths of coach-class prices for flights from Dallas, TX to Roswell, NM on March 1st, 2011. American Airlines is the only carrier that serves this route; there are three flights available during that day.

The behavior of ticket prices depicted is representative of monopoly markets. There are three main stylized facts in the data. First, prices increase in discrete jumps. Second, there are several distinct times when the lowest price for all flights jumps up simultaneously. As in the figure, these times typically occur 6, 13 and 20 days before departure. Third, between these jumps, prices are

Fares • 15 items

Fare Basis	Airline	Booking Class	Trip Type	Fare	Cabin	Effective Date	Expiration Date	Min / Max Stay	Adv Purchase Req
QA21ERD1	AA	Q	One-Way	138.00(USD)	E				21
SA21ERD1	AA	S	One-Way	146.00(USD)	E				21
SA14ERD1	AA	S	One-Way	154.00(USD)	E				14
GA07ERD1	AA	G	One-Way	211.00(USD)	E				07
VA07ERD1	AA	V	One-Way	225.00(USD)	E				07
WA07ERD1	AA	W	One-Way	242.00(USD)	E				07
LA00ERD5	AA	L	One-Way	249.00(USD)	E				
MA00ERD5	AA	M	One-Way	363.00(USD)	E				
KA00ERD5	AA	K	One-Way	463.00(USD)	E				
HA00ERDR	AA	H	One-Way	577.00(USD)	E				30
YA2AAD	AA	Y	One-Way	756.00(USD)	E				30
Y	AA	Y	One-Way	770.00(USD)	E				
YA0UPAMR	AA	A	One-Way	823.00(USD)	B				30
FA2AA	AA	F	One-Way	973.00(USD)	F				30
F	AA	F	One-Way	1237.00(USD)	F				

Figure 1: List of available fares from Dallas, TX to Roswell, NM for 03/01/2011, six weeks before departure

relatively stable.

This behavior results largely because of the institutional details surrounding the way airlines set ticket prices. The lowest price of a ticket for a given flight is determined by the lowest fare with available seats in the corresponding booking class. There are three reasons that the lowest price of an airline ticket for a given flight may change over time. First, if the number of days before departure is less than the APR, travelers cannot use that fare to buy a ticket. Less restrictive fares are usually more expensive, which results in a price increase. If we look at Figure 2 again, we can see that the first major price increase occurred 20 days before departure: the price went up from $138 to $154. This was the day when the advance purchase requirement for the two lowest fares became binding.

Second, the decision of the revenue management department to open or close availability in a certain booking class may change the lowest price. Eighteen days before departure, the revenue management department of American Airlines closed booking class S for flight AA 2705 but kept booking class G open. As a result, the lowest price for this flight went up from $154 to $211.

Finally, the pricing department can add a new fare, as well as update or remove an existing one. On very competitive routes, airline pricing analysts monitor their competitors very closely,

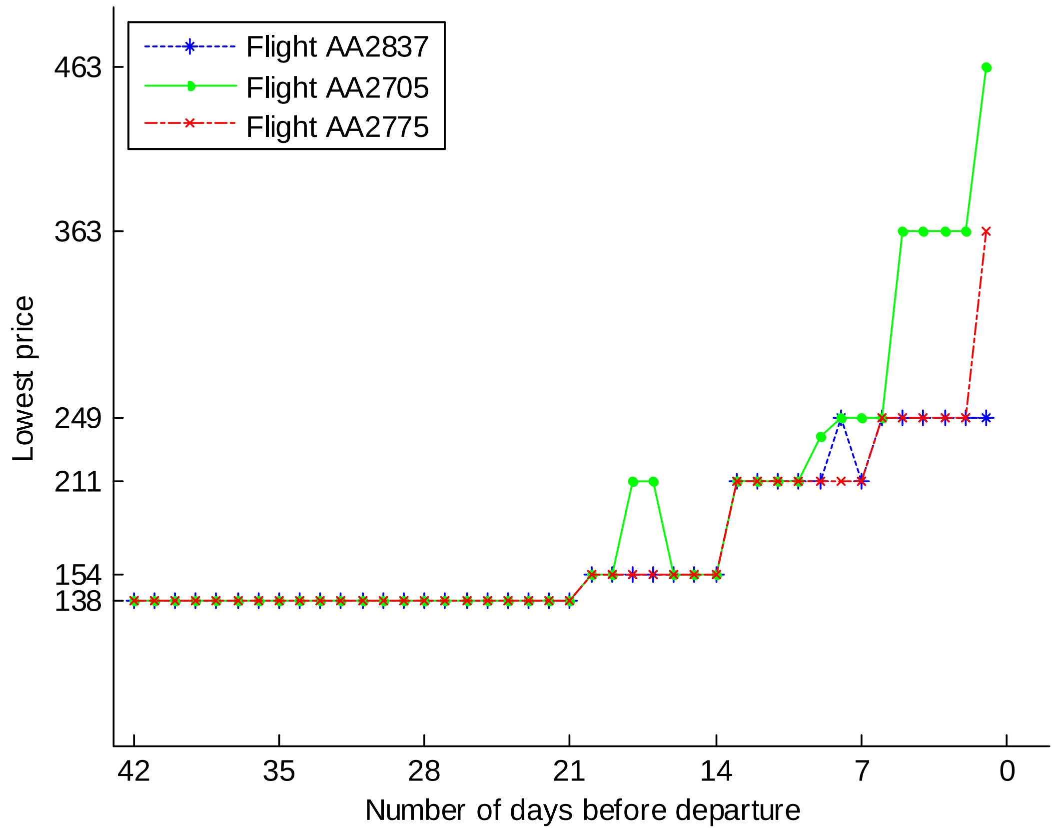

Figure 2: Example Price Path. Route: Dallas, TX - Roswell, NM. Departure Date: 03/01/11

often responding to their price moves on the same day (Talluri and van Ryzin, 2005). On routes with few operating carriers, the set of fares is usually stable. For example, during the time period depicted on Figure 2, the pricing department of American Airlines did not update fares for flights from Dallas to Roswell departing on March 1st, 2011. Changes in prices were caused primarily by APR restrictions or the decisions of the revenue-management department.

3 The Model of Optimal Fares

A theoretical model that is able to generate the stylized facts listed in Section 2 would have to include the decision problems of both the pricing and revenue-management departments. The solution of the pricing department's problem observed in the data is a finite set of fares that include advance purchase requirements. To construct an optimal set of such fares, the pricing department has to be able to calculate the value of the airline's expected profit for each possible set of fares. This value, in turn, depends on the strategy of the revenue management department that takes the set of fares as given and updates availability of each booking class in real time. Solving a model like that would be a very difficult task. Furthermore, the model would have to explain why it is optimal to impose advance purchase requirements with the particular lead times observed in the data. Any explanation would have to make somewhat artificial assumptions on changes in fundamental parameters that happen exactly on these days.

To reduce complexity, I focus on the decision problem of the pricing department only, and look not at the actual realizations of price paths, but at the set of fares that generate these paths. I collapse the window for selling tickets to five periods: 21 days and more, from 14 to 20 days, from 7 to 13 days, from 3 to 6 days, and less than 3 days before departure. These thresholds correspond to the advance purchase requirements observed in the data. For each period of sale, I determine the lowest fare satisfying the corresponding advance purchase requirement. These are the fares that a simplified model will have to explain.

To recover consumers' preferences (or, to be precise, the airline's expectations about consumers' preferences), I develop a model that shows how a set of parameters reflecting travelers' preferences transform into a path of profit-maximizing fares.[6] The model is initially formulated for a representative origin and destination and a representative departure date.

3.1 Market

Suppose there is a population of potential travelers of size $\tilde{M}$ flying from Origin to Destination on a particular date. Each traveler buys at most one ticket. Tickets are sold by a monopoly airline over the course of T periods. Each period the airline announces a price p_t at which it is willing to sell any number of tickets to any traveler. The vector $\mathbf{p} = (p_1, ..., p_T)'$ denotes a price path chosen by the airline.

A sold ticket can be cancelled by the traveler, with cancellation fee of f. The fee $f \geq 0$ is taken to be exogenous because in practice U.S. airlines have only one cancellation fee that applies to all domestic routes.

[6] I do not consider a more general problem of finding a profit-maximizing mechanism since the mechanism observed in the data is implemented through publicly posted prices. This problem has been studied by Gershkov and Moldovanu (2009), Board and Skrzypacz (2011), and Hoerner and Samuelson (2011), among others.

3.2 Consumers

Types, Arrival and Exit The population of potential travelers consists of I discrete types; types are indexed by $i = 1, ..., I$. (In the estimation, I assume that $I = 2$: leisure and business travelers.) The sizes of different types of potential buyers change over time for three reasons. First, each period new travelers arrive to the market.[7] The mass of new buyers of type i who arrive at time t is equal to $\tilde{M}_{it} = \tilde{\lambda}_{it} \cdot \gamma_i \cdot \tilde{M}$, where γ_i is the weight of each type in the population and $\tilde{\lambda}_{it}$ is the type-specific arrival rate. Second, those travelers who bought tickets in previous periods are not interested in purchasing additional ones. Third, each period a fraction of travelers who arrived in the previous periods learn that they will not be able to fly due to some contingency, so they cancel the ticket (if purchased) and exit the market. The probability that a traveler of type i learns that she will not be able to fly is equal to $(1 - \delta_i)$ in every period.

Preferences Travelers know their utilities conditional on flying but are uncertain if they are able to fly. If a traveler ι of type i buys a ticket in period t, she pays the price p_t and, conditional on flying, receives:

$$u_{\iota it} \equiv \mu_i + \sigma_i \left(\varepsilon_{\iota it} - \varepsilon_{\iota i0}\right), \tag{1}$$

where μ_i is type-i's mean utility from flying on this route measured in dollar terms, $\varepsilon_{\iota it}$ are i.i.d. Type-I extreme value terms that shift traveler ι's utility in each period, and σ_i is a normalizing coefficient that controls the variance of $\varepsilon_{\iota it}$. This coefficient captures the slope of the demand curve and hence the price sensitivity across the population of type-i travelers: the lower is the coefficient, the more sensitive are type-i travelers. The traveler learns all components of their utilities defined in equation (1) at $t = 0$. I use the extreme value errors to represent the consumer tastes that the airline and the researcher do not observe.

After purchase, the traveler can cancel a ticket. If she cancels a ticket at time t', she loses the price she paid, p_t, but may receive a monetary refund if the cancellation fee does not exceed the price. The refund is equal to $\max\left(p_t - f\,, 0\right)$. Since the refund does not exceed the price of the ticket, the traveler will cancel her ticket only if she learns that she is not able to fly. If the traveler

[7] Without this assumption, the profit-maximizing monopolist would forgo the opportunity to discriminate over time (Stokey, 1979). Board (2008) analyzes the profit-maximizing behavior of a durable goods monopolist when incoming demand varies over time.

doesn't fly, her utility is normalized to zero.

Travelers are forward-looking and make purchase decisions to maximize their expected utility. They face the following tradeoff: if they wait, they will receive more information about their travel plans but may have to pay a higher prices as the airline could increase prices over time.

Individual demand Consider the utility-maximization problem of a type-i traveler who is in the market at time τ. She has $T-\tau$ periods to buy a ticket. She buys a ticket at time τ only if it gives a higher utility than buying a ticket in subsequent periods or not buying a ticket at all. If she buys a ticket in period τ, then her net expected utility is given by:

$$\left[\delta_i^{T-\tau} u_{\iota i\tau} + R_{i\tau}\right] - p_\tau,$$

where $R_{i\tau}$ denotes the expected value of the refund:

$$R_{i\tau} = \left(1 - \delta_i^{T-\tau}\right) \max\left(p_\tau - f\,, 0\right).$$

Suppose the traveler considers waiting until period τ'. Then with probability $\left(1 - \delta_i^{\tau'-\tau}\right)$ she learns about a travel emergency and exits the market. With the remaining probability $\delta_i^{\tau'-\tau}$ she stays in the market. If she buys a ticket, she receives $\delta_i^{T-\tau'} u_{\iota i\tau'} + R_{i\tau'} - p_{\tau'}$. In this case, her expected utility at time τ is equal to

$$\delta_i^{T-\tau}\left[\mu_i + \sigma_i\left(\varepsilon_{\iota i\tau'} - \varepsilon_{\iota i0}\right)\right] + \delta_i^{\tau'-\tau}\left(R_{i\tau'} - p_{\tau'}\right).$$

Thus, the traveler buys a ticket in period τ if the following set of inequalities holds:

$$\begin{aligned}
\delta_i^{T-\tau}\left[\mu_i + \sigma_i\left(\varepsilon_{\iota i\tau} - \varepsilon_{\iota i0}\right)\right] + R_{i\tau} - p_\tau &> \delta_i^{T-\tau}\left[\mu_i + \sigma_i\left(\varepsilon_{\iota i\tau'} - \varepsilon_{\iota i0}\right)\right] + \delta_i^{\tau'-\tau}\left(R_{i\tau'} - p_{\tau'}\right) \\
&\quad \text{for all } \tau < \tau' \leq T \text{ and} \\
\delta_i^{T-\tau}\left[\mu_i + \sigma_i\left(\varepsilon_{\iota i\tau} - \varepsilon_{\iota i0}\right)\right] + R_{i\tau} - p_\tau &> 0.
\end{aligned}$$

These inequalities can be rewritten in a more convenient way:

$$\begin{aligned}
\frac{\delta_i^{T-\tau}\mu_i + R_{i\tau} - p_\tau}{\sigma_i\delta_i^{T-\tau}} + \varepsilon_{\iota i\tau} &> \frac{\delta_i^{T-\tau'}\mu_i + R_{i\tau'} - p_{\tau'}}{\sigma_i\delta_i^{T-\tau'}} + \varepsilon_{\iota i\tau'} \text{ for all } \tau < \tau' \leq T \text{ and} \\
\frac{\delta_i^{T-\tau}\mu_i + R_{i\tau} - p_\tau}{\sigma_i\delta_i^{T-\tau}} + \varepsilon_{\iota i\tau} &> \varepsilon_{\iota i0}.
\end{aligned} \tag{2}$$

Market demand To calculate the firm's expected demand for tickets, we need to know the demand of each traveler type as well as the size of each type in a given period. Denote by $s_{it\tau}$ the share of type-i buyers who arrived in period τ and purchase a ticket in period t conditional on not exiting the market. This share corresponds to the probability that traveler ι has a realization of $\varepsilon_{\iota it}$, $t = \tau, ..., T$ that satisfies inequalities defined in (2). Under the assumption that $\varepsilon_{\iota i\tau}$ is extreme value, this share is equal to

$$s_{it\tau} = \frac{\exp\left(\frac{\delta_i^{T-t}\mu_i + R_{it} - p_t}{\sigma_i \delta_i^{T-t}}\right)}{1 + \sum_{k=\tau}^{T} \exp\left(\frac{\delta_i^{T-k}\mu_i + R_{ik} - p_k}{\delta_i^{T-k}\sigma_i}\right)}.$$

Consider now the size of type-i buyers who arrived in period τ. By time t, only $\delta_i^{t-\tau}$ of the initial size has not exited the market due to a realized emergency. Thus, the total demand of type-i travelers is equal to:

$$D_{it} = \sum_{\tau=1}^{t} s_{it\tau} \delta_i^{t-\tau} \tilde{M}_{i\tau};$$

the market demand for tickets in period t is given by:

$$D_t = \sum_{i=1}^{I} D_{it}.$$

This is the demand faced by the monopoly airline.

3.3 Producers

The airline is the only producer in the market. Its revenue comes from selling tickets and collecting cancellation fees. Denote by $\tilde{c}$ the opportunity cost of flying an additional passenger. This cost can be interpreted in two ways. First, in a more general model, it may reflect the expected marginal revenue of adding an additional unit of capacity to the market. Second, if there are connecting passengers on the flight, this cost reflects the marginal revenue from flying those passengers.

Denote the payoff relevant parameters by $\tilde{\theta} = \left(\gamma, \mu, \sigma, \delta, \tilde{\lambda}, \tilde{c}, \tilde{M}\right)$. The set of all admissible parameters, $\tilde{\Theta}$, is defined as follows: $\gamma_i > 0$, $\mu_i > 0$, $\sigma_i > 0$, $0 < \delta_i < 1$, $\tilde{\lambda}_{it} > 0$, $\tilde{c} > 0$, $\tilde{M} > 0$, $\sum_i \gamma_i = 1$. The airline learns a realization of these parameters at $t = 0$. The airline is forward-looking and maximizes its total profit. Since there is no aggregate demand uncertainty, total profit is a deterministic function and is equal to:

$$\pi\left(\mathbf{p};\tilde{\theta}\right)=\sum_{i=1}^{I}\sum_{t=1}^{T}\left(p_t-R_{it}-\tilde{c}\delta_i^{T-t}\right)D_{it}.$$

A price path $\mathbf{p}$ chosen by the airline uniquely determines the supply of seats and their allocation among travelers.

3.4 Welfare

For each price path $\mathbf{p}$, we can calculate the sum of utilities for each type of travelers. Consider the group of type-i travelers who arrived at time τ and define the average aggregate utility of this group by $v_{i\tau}(\mathbf{p})$. Then,

$$v_{i\tau}(\mathbf{p})=\int_{\iota}\max_{\tau\leq\tau'\leq T}\left\{\delta_i^{T-\tau}\left[\mu_i+\sigma_i\left(\varepsilon_{\iota i\tau'}-\varepsilon_{\iota i0}\right)\right]+\delta_i^{\tau'-\tau}\left(R_{i\tau'}-p_{\tau'}\right),0\right\}d\iota.$$

Integrating with respect to the extreme value distribution, we get:

$$v_{i\tau}(\mathbf{p})=\delta_i^{T-\tau}\sigma_i\log\left(1+\sum_{t=\tau}^{T}\exp\left(\frac{\delta_i^{T-t}\mu_i+R_{i\tau}-p_\tau}{\delta_i^{T-t}\sigma_i}\right)\right).$$

Then, the total sum of traveler's utilities equals:

$$V(\mathbf{p})=\sum_{i=1}^{I}\sum_{\tau=1}^{T}v_{i\tau}(\mathbf{p})\tilde{M}_{i\tau}.$$

Define social welfare as the sum of travelers' ex-post utilities and the airline's profit. The supply and allocation of seats among travelers are efficient if they maximize social welfare. A price path $\mathbf{p}$ is called efficient if it induces efficient supply and allocation of seats.

Proposition 1 *An efficient price path exists only if the cancellation fee is equal to zero.*

Proofs of all Propositions are in Appendix C.

A higher cancellation fee makes a ticket less attractive to travelers. For this reason, I refer to it as a measure of ticket quality. Proposition 1 shows that the socially optimal quality of tickets is achieved only when the cancellation fee is equal to zero. A positive cancellation fee thus implies *inefficiency in the quality of production.*

Proposition 2 *If the cancellation fee is equal to zero, the unique efficient price path is flat and equal to $p_t\equiv\tilde{c}$.*

Proposition 2 reveals two impediments to efficient supply and allocation of seats: market power and dynamic pricing. First, if the price exceeds marginal cost, then the number of seats sold by the airline is lower than the socially efficient level. As a result, social welfare is lower than its maximum level due to *inefficiency in the quantity of production.* Second, if the price path is not flat, then the airline charges different prices in different time periods, which results in a misallocation of seats among travelers. In this case, social welfare does not achieve its maximum level due to *inefficiency in allocation.* Inefficiency in quality of production, inefficiency in quantity of production, and inefficiency in allocation are the three reasons why a price path may not induce an efficient outcome.

3.5 Optimal price path and its welfare properties

A price path $\mathbf{p}$ is called *optimal* if it maximizes the airline's profit $\pi\left(\mathbf{p};\tilde{\theta}\right)$. Denote by $\mathbf{p}^*\left(\tilde{\theta}\right)$ the optimal price path as a function of $\tilde{\theta}$. The next proposition shows that this price path exists and is unique.

Proposition 3 *For any $\tilde{\theta} \in \tilde{\Theta}$ and cancellation fee $f \geq 0$, the optimal price path $\mathbf{p}^*\left(\tilde{\theta}\right)$ is uniquely defined and satisfies the following system of equations:*

$$G\left(\mathbf{p}^*\left(\tilde{\theta}\right),\tilde{\theta}\right) = 0,$$

where $G : R^T \times R^{\dim\left(\tilde{\theta}\right)} \to R^T$ is the set of first order conditions:

$$G\left(\mathbf{p},\tilde{\theta}\right) = \left(\frac{\partial\pi\left(\mathbf{p};\tilde{\theta}\right)}{\partial p_1},...,\frac{\partial\pi\left(\mathbf{p};\tilde{\theta}\right)}{\partial p_T}\right)'.$$

Using Proposition 3, we can directly verify that the optimal price path cannot be efficient even when the cancellation fee is zero. In practice, airlines often impose a positive cancellation fee for lower fares, which can further decrease social welfare. The reason why airlines do this is simple: it increases their profit. Even though a positive cancellation fee diminishes the quality for all traveler groups, travelers with a higher probability of cancellation suffer from it more. If the probability of cancellation is positively correlated with the utility from flying, the fee screens travelers by their type. Fares observed in practice often come with advance purchase requirements. These requirements implement a price path that increases over time.

Thus, our theoretical analysis suggests that price paths observed in practice lead to all three types of inefficiency identified in the previous subsection: inefficiency in quality of production, inefficiency in quantity of production, and inefficiency in allocation. To assess welfare losses associated with each type of inefficiency, we need to know the estimates of the demand and cost parameters $\tilde{\theta}$.

4 Data

4.1 Monopoly Markets

A market is defined by three elements: origin airport, destination airport and departure date. A product is an airline ticket that gives a passenger the right to occupy a seat on a flight from the origin to the destination departing on a particular date.

To be included in my dataset, a domestic route has to satisfy five criteria. First, the operating carrier on the route was the only scheduled carrier in the time period I consider. Second, the carrier had to have been the dominant firm for at least a year before the period I consider. Specifically, its share in total market traffic had to be at least 95% in each month prior to the period of study. Third, at least 90% of the passengers flying from the origin to the destination must fly nonstop. Fourth, total market traffic on the route must be at least 1000 passengers per quarter. Fifth, there should be no alternative airports that a traveler willing to fly this route can choose. I do not include routes to/from Alaska or Hawaii. These criteria were chosen to limit ambiguities in markets and to ensure the markets were nontrivial.

In all, I have 76 directional routes that satisfy these criteria. A typical route has a major airline hub as either its origin or destination. There are six monopoly airlines in the dataset: American Airlines (26 routes to or from Dallas/Fort Worth, TX), Alaska Airlines (26 routes mainly to or from Seattle, WA), United/Continental Airlines (8 routes to or from Houston, TX), AirTran Airways (4 routes to or from Atlanta, GA), Spirit Airlines (6 routes to or from Fort Lauderdale, FL), and US Airways (6 routes to or from Phoenix, AZ). Figure A1 in the appendix shows a map of these routes. Table 1 gives summary statistics of route characteristics.

Table 1: Monopoly routes: summary statistics

	mean	st.d.
distance	401	213
median family income	$71,942	$8,432
average ticket price	$151	$136
quarterly traffic, passengers	16,663	11,854
share of major airline, traffic	0.9953	0.0188
share of nonstop passengers	0.9772	0.0255
share of connecting passengers	0.6511	0.2616
load factor	0.7104	0.0896

4.2 Data Sources

In the industry, fares are distributed by the Airline Tariff Publishing Company[8] (ATPCO), an organization that receives fares from all airlines' pricing departments. It publishes North American fares three times a day on weekdays, and once a day on weekends and holidays.[9] Until recently, the general public did not have access to information stored in global distribution systems. Yet a few websites have provided travelers with recommendations on when is the best time to book a ticket based on this information. In 2004, travelers received direct access to public fares and booking class availabilities through several new websites and applications. I recorded fares manually from a website that has access to global distribution systems subscribed to ATPCO data. This website is widely known among industry experts and regarded as a reliable and accurate source of public fares.[10] I recorded fares that were published six weeks before departure. The period of six weeks is motivated by three facts. First, few tickets are sold earlier than that period. Second, most travel websites recommend searching for cheap tickets six to eight weeks before departure. Third, when a pricing department updates fares it takes into account flights that depart in the next several weeks rather than flights that depart in the next several days. Thus, I believe that it is reasonable to assume that fares posted six weeks before departure reflect the optimal decision of pricing departments.

[8] Until recently, ATPCO was the only agency distributing fares in North America. In March 2011, SITA, the only international competitor of ATPCO, received an approval from the US Department of Transportation and the Canadian Transportation Agency to distribute data for airlines operating in the region.

[9] On weekdays, the fares are published at 10 am, 1 pm and 8 pm ET. On weekends, the fares are published at 5 pm. In October 2011, ATPCO added a fourth filing feed on weekdays – at 4 pm ET.

[10] In addition to public fares that are available to any traveler, airlines can offer private fares. Private fares are discounts or special rates given to important travel agencies, wholesalers, or corporations. Private fares can be sold via a GDS that requires a special code to access them or as an offline paper agreement. In the United States, the majority of sold fares are public.

I consider three quarters of departure dates between October 1, 2010 and June 30, 2011. Besides the data on daily fares described above, I use monthly traffic data from the T-100 Domestic Market database and the Airline Origin and Destination Survey Databank 1B that contains a 10% random sample of airline tickets within a given quarter. Both datasets are reported to the U.S. Department of Transportation by air carriers and are freely available to the public.

In the estimation, I control for several route characteristics, which allows me to compare different markets with each other. These characteristics include route distance, median household income in the Metropolitan Statistical Areas to which origin and destination airports belong, and population in the areas. A detailed description of this part of the data is in Appendix B

5 Estimation.

5.1 Econometric Specification[11]

My empirical model allows for two types of travelers. I refer to the first type as leisure travelers, and to the second type as business travelers. Leisure travelers are highly price sensitive customers who are ready to book earlier and are more willing to accept ticket restrictions. Business travelers, on the other hand, are less price sensitive, book their trips later and less likely to accept restrictions.[12] The demand parameters of the model of optimal fares are able to capture these distinctions.

For a given departure date $d = 1, ..., D$ and a given route $n = 1, ..., N$, the demand and cost parameters $\tilde{\theta}_{nd}$ determine the optimal price path $\mathbf{p}^* \left(\tilde{\theta}_{nd}\right)$. These parameters are known to the airline but unknown to the researcher. The goal of the estimation routine is to recover $\tilde{\theta}_{nd}$ for each date and route from the observed price and quantity data. Given the limitations of the dataset, I need to reduce the dimension of the unknown parameters. To do this, I restrict both observed and unobserved variation in the parameters within and across markets.

The shares of each type, γ_i, are assumed to be the same in all routes and all departure dates. Type-specific mean utilities from flying, μ_i, are proportional to the route distance. The proportionality coefficient in turn linearly depends on the route median income. These coefficients do not

[11] I plan to add more route and time covariates in the later version of the paper.
[12] See, Phillips (2005).

vary with the departure date. Thus,

$$\mu_{ind} = \mu_{1i} + (\mu_{2i} + \mu_{3i} \cdot income_n) \cdot dist_n.$$

The variance of the type-I error (σ_i) that controls intertemporal utility variation within a type is the same in all markets and all departure dates. The probability of having to cancel the trip, $1-\delta_i$, is also the same in all routes but varies with the departure date. It can take two type-specific values: one for regular season and one for holiday seasons. Holiday season departure dates correspond to Thanksgiving, Christmas, New Year's and Spring Break.[13] The probability of canceling a trip is different during these periods as travelers may be more certain about their holiday trips than about their regular trips. If we denote by h_d the holiday season dummy variable, then

$$\delta_{ind} = \delta_i^{\text{holiday}} \cdot h_d + \delta_i^{\text{regular}} \cdot (1 - h_d).$$

The share of new passengers who arrive in period τ, has the following parametric representation:

$$\hat{\lambda}_{i\tau nd} = \lambda(\tau, T, \alpha_i) + \varepsilon_{\lambda\tau nd} = \left(\frac{\tau}{T}\right)^{\alpha_i} - \left(\frac{\tau - 1}{T}\right)^{\alpha_i} + \varepsilon_{\lambda\tau nd},$$

where $\varepsilon_{\lambda 1nd}$ is normalized to 0 and $\varepsilon_{\lambda 2nd}, ..., \varepsilon_{\lambda Tnd}$ are unobserved i.i.d. mean-zero errors. The parameter α_i determines the time when the majority of type-i consumers start searching for a ticket: types with low values of α_i begin their search early, types with high values of α_i arrive to the market only a few days before departure. These parameters are the same for all routes and departure dates. The unobserved error $\varepsilon_{\lambda\tau nd}$ randomly shifts the arrival probabilities. Since the airline observes these errors before it determines its price path, these errors explain a part of the daily variation in observed fares. The sum of the errors does not affect the optimal price path and thus is not identified from the observed fares. For this reason, I normalize the value of the first error to zero.

The opportunity cost of flying an additional passenger comes from a distribution with a mean proportional to the route distance: $\tilde{c}_{nd} = c \cdot dist_n + \varepsilon_{cnd}$ where ε_{cnd} is an i.i.d. unobserved mean-zero error. The unobserved error ε_{cnd} randomly shifts the opportunity cost of flying a passenger each day and in each route and also explains a part of the daily variation in observed fares. This error shifts the entire time path of prices, while $\varepsilon_{\lambda\tau nd}$ affects relative levels of the prices in the path.

[13] See details in Appendix B.

The total number of potential travelers is different for each route and each departure date. I denote by M_n the mean number of travelers on route n and assume that the deviations from these means, the arrival errors $\varepsilon_{\lambda n\tau d}$, and the cost errors ε_{cnd} are jointly independent.

Together, we can divide all demand and cost parameters known to the airline, $\tilde{\theta}_{nd}$, into three groups: estimated coefficients $\theta = (\gamma, \mu, \sigma, \delta, \alpha, c)$ and M_n, errors unobserved to the researcher $\varepsilon_{nd} = (\varepsilon_{\lambda nd}, \varepsilon_{cnd})$, and market specific covariates (h_d, x_n), where x_n denotes route characteristics such as $(dist_n, income_n)$. These restrictions allow me to estimate the coefficients jointly for all markets in my sample.

5.2 Moment Restrictions

To estimate the demand and cost parameters θ, I follow the standard practice of using both price and quantity data. Here I face the nonstandard complication that these data are observed with different frequencies: prices are observed daily, quantities are observed quarterly. Only having quarterly quantity data means that they contain two sources of variation: variation due to different departure dates and variation due to different purchase dates. I use the model of optimal fares to distinguish between these two sources of variation.

5.2.1 Daily prices

Define by p_{tnd} the lowest fare satisfying the advance purchase requirement for period of sale t for route n and departure date d. Since the posted fares should be equal to the optimal fares predicted by the model, the posted fares should satisfy the system of first order conditions should hold stated in Proposition 3:

$$G\left(p_{nd}, \tilde{\theta}_{nd}\right) = 0.$$

Recall that $\tilde{\theta}_{nd} = (\theta, h_d, x_n\, , \varepsilon_{nd})$. Then $G\left(p_{nd}, \theta, h_d, x_n\, , \varepsilon_{nd}\right) = 0$. To construct moment restrictions that correspond to the posted prices, we need to invert the system of equations to derive an expression for the unobserved error term ε_{nd}.

Proposition 4 *There exists a unique mapping $g_P : R^T \times R^{\dim(\theta)} \times R^{\dim(h_d)} \times R^{\dim(x_n)} \rightarrow R^T$, such that for any θ, it holds that $G\left(p_{nd}, \theta, h_d, x_n\, , g_P\left(p_{nd}, \theta, h_d, x_n\right)\right) = 0$.*

Since we assumed that ε_{nd} has zero mean, the moment restrictions that correspond to the observed prices take the following form:

$$\mathbb{E}\varepsilon_{nd} = \mathbb{E}g_p\left(p_{nd}, \theta, h_d, x_n\right) = 0.$$

I use these restrictions as the basis for one set of sample moment conditions.

5.2.2 Monthly traffic

The model predicts the total number of flying passengers for departure date d and route n is equal to $\sum_{i=1}^{I}\sum_{t=1}^{T}\delta_i^{T-t}D_{it}\left(p_{nd}, \tilde{\theta}\right)$. Denote by Q_{nm} the total number of enplaned passengers observed in the data for route n and month m. Thus, the predicted number of enplaned passengers is equal to the actual one if:

$$\sum_{d\in \text{month}(m)}\sum_{i=1}^{I}\sum_{t=1}^{T}\delta_{id}^{T-t}D_{it}\left(p_{nd}, \tilde{\theta}\right) = Q_{nm}.$$

Denote by $g_M\left(p_{nd}, \tilde{\theta}, M_{nm}\right) = \sum_{d\in \text{month}(m)}\sum_{i=1}^{I}\sum_{t=1}^{T}\delta_{id}^{T-t}D_{it}\left(p_{nd}, \tilde{\theta}\right) - Q_{nm}$. Then, a moment restriction that corresponds to the observed number of enplaned passengers is given by:

$$\mathbb{E}g_M\left(p_{nd}, \tilde{\theta}, Q_{nm}\right) = 0.$$

I use this restriction as to define a second set of sample moment conditions.

5.2.3 Quarterly sample of tickets

Denote by r_{lnq} a ticket issued for market n in quarter q and let $p\left(r_{lnq}\right)$ and $f\left(r_{lnq}\right)$ denote the corresponding one-way fare and number of traveling passengers[14]. The quarterly ticket data have several potential sources of measurement error. These data include special fares, frequent flier fares, military and government fares, etc. To reduce the impact of these special fares, I do the following. First, I divide the range of possible prices into $B+1$ non-overlapping intervals[15]: $[p_b, p_{b+1}]$, $b = 0, ..., B$. For each interval, the model predicts the total number of tickets sold during the quarter. Hence, we can calculate the model-predicted probability of drawing a ticket from each interval. Denote by w_{bnq} the probability of drawing a ticket with a price that belongs to interval

[14] I manually removed the taxes to get the published fares. The details are in Appendix B.

[15] I estimate the model using the following 17 price thresholds: \$20, \$50, \$80, \$100, \$120, \$135, \$150, \$170, \$190, \$210, \$220, \$240, \$270, \$300, \$330, \$360, \$410.

$[p_b, p_{b+1}]$ for market n in quarter q. This probability equals:

$$w_{bnq}\left(p_{nd}, \tilde{\theta}\right) = \frac{\sum_{d \in \text{quarter}(q)} \sum_{i=1}^{I} \sum_{t=1}^{T} \delta_{id}^{T-t} D_{it}\left(p_{nd}, \tilde{\theta}\right) \cdot 1\left\{p_{tnd} \in [p_b, p_{b+1}]\right\}}{\sum_{d \in \text{quarter}(q)} \sum_{i=1}^{I} \sum_{t=1}^{T} \delta_{id}^{T-t} D_{it}\left(p_{nd}, \tilde{\theta}\right)},$$

Similarly, we can calculate the relative frequency of observing a ticket within a given price range using the 10% sample of airline tickets.[16] Denote these frequencies as $\hat{w}_{bnq}$ and define $g_W\left(p_{nd}, \tilde{\theta}, r_{nd}\right) = [w_{1nq} - \hat{w}_{1nq}, ..., w_{Bnq} - \hat{w}_{Bnq}]'$.

Assuming that the 10% sample is drawn at random, we can derive the third part of the moment restriction set from the population moment conditions for each price interval:

$$\mathbb{E} g_W\left(p_{nd}, \tilde{\theta}, r_{nm}\right) = 0.$$

To avoid linear dependence of the moment restrictions, I exclude the last interval.

To summarize, we have T moment restrictions based on the daily fare data, one restriction based on the monthly traffic data and B restrictions based on the quarterly ticket data. I use these $T + B + 1 = 5 + 17 + 1 = 23$ moment conditions to estimate 15 parameters that define θ. These parameters are identified from the joint distribution of optimal prices and quantities aggregated to the quarterly level. Appendix D shows how changes in each parameter affect this distribution.

5.3 Estimation Method and Inference

I use a two-step generalized method of moments. The optimal weighting matrix is estimated using unweighted moments. For computational purposes, I optimize the objective function for a monotone transformation of the parameters. This transformation guarantees that the estimates will be positive and, where necessary, less than one. The standard errors are calculated using the asymptotic variance matrix for a two-step optimal GMM estimator.

[16] I treat a ticket with multiple passengers as multiple tickets with one passenger each. If a ticket has a round-trip trip fare, I assume that I observe two tickets with two equal one-way fares. Finally, I only take into account those intervals for which the model predicts non-zero probabilities.

Table 2: Estimates of demand and cost parameters

		Leisure Travelers	Business Travelers
Share of Traveler Type	γ_i	79.71% (0.20%)	20.29% (0.20%)
Mean Utility	μ_i	$\underset{(1.05)}{\$43.63} + \left[\underset{(0.01)}{\$7.11} + \underset{(0.05)}{0.89} income_n\right] dist_n$	$\underset{(19.35)}{\$320.23} + \left[\underset{(4.95)}{\$27.89} + \underset{(1.54)}{2.54} income_n\right] dist_n$
Price sensitivity	σ_i	0.34 (0.007)	2.46 (0.06)
Probability of cancellation regular season / holiday season	$1-\delta_i$	9.95% / 0.79% (0.11%) (0.01%)	12.33% (0.13%)
Arrival process parameter	α_i	0.02 (0.09)	7.85 (1.82)
Opportunity cost of flying	c	\$1.00$\times$ $dist_n$ (\$3.72)	

Note: $income_n$ is in \$ 100,000, $dist_n$ is in 100 miles.

6 Results[17]

6.1 Demand and Cost Estimates

Table 2 presents the optimal GMM estimates of the demand and cost parameters, θ. The estimates suggest that 76% of passengers travel for leisure purposes. Business travelers are willing to pay up to six times more for a seat and they are less price sensitive. If all fares go up by 1%, the total demand of leisure travelers goes down by 1.3%, while the total demand of business travelers goes down by only 0.8%. Business travelers tend to avoid tickets with a cancellation fee as the probability that they have to cancel a ticket is high (12.33% per period of sale).

The dynamics of arrival of each traveler type is depicted by the dotted lines in Figure 3. A significant share of leisure travelers start searching for a ticket at least six week prior to departure. By contrast, 83% of business travelers begin their search in the last week. The bar graph in Figure 3 demonstrates how the number of active buyers changes over time. In the first few periods, the number of active buyers goes down as travelers buy tickets or learn that they will not be able to fly. The arrival of new travelers does not counteract this decrease. Most business travelers are estimated to start searching for tickets a week before departure. As a result, the number of active ticket buyers goes up.

[17]This version does not include DB1B data for the second quarter of 2011. The Department of Transportation hasn't released them yet.

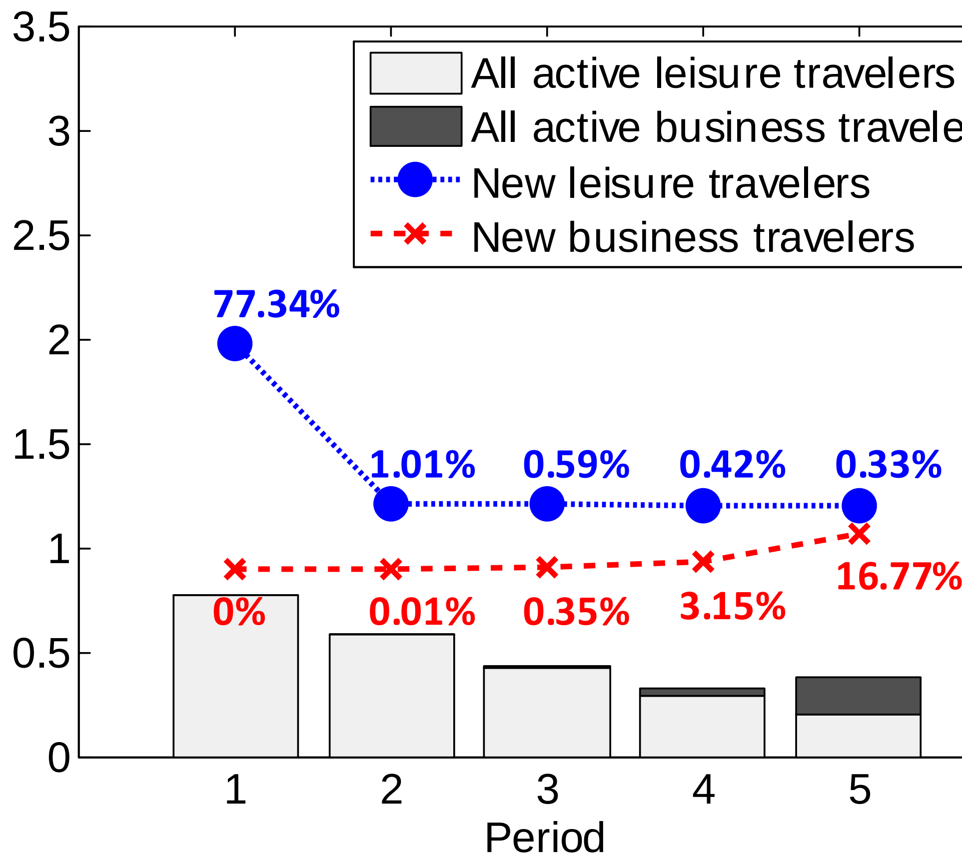

Figure 3: Dynamics of active buyers on a route with median income and distance

6.2 Optimal Price Path and Price Elasticities

To put these estimates into perspective, I use the model of optimal fares to calculate the price path for flights on a route with median characteristics on a non-holiday departure date. Figure 4 shows this path together with the quantities of tickets purchased in each period by leisure and business travelers. The figure shows that leisure travelers usually purchase tickets up until seven days before departure, prior to the moment when most business travelers arrive in the market. When business travelers arrive, the airline significantly increases the price, trying to extract more surplus from travelers who are willing to pay more.

Table 3 presents the estimates of price elasticities evaluated at the optimal price path. The estimates show that in periods 1 and 5 the airline extracts almost the maximum amount of revenue from travelers as the aggregate elasticities are close to one. In both periods, the buyers are almost homogenous. In period 1, the majority of active buyers are leisure travelers. In period 5, the price

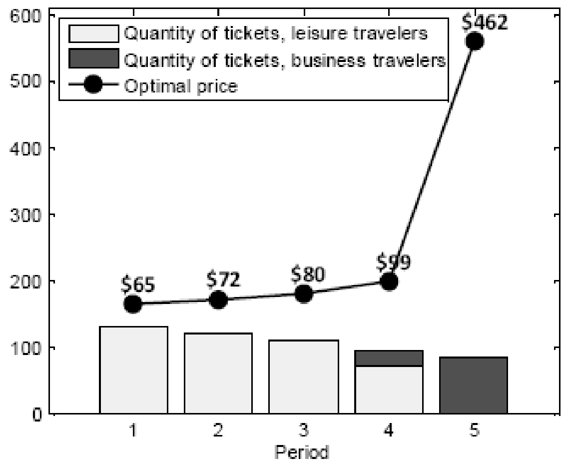

	Allocation			
	Current		Efficient	
	L	B	L	B
Quantity, tickets	0.434	0.107	0.639	0.173
Quantity, seats	0.329	0.104	0.514	0.171
Average price	$77	$382	$4	$4
Expected net utility	$21	$195	$82	$482
Profit	$73		$0	
Social Welfare	$130		$163	

Note: All numbers are divided by the number of potential travelers.

Figure 4: Optimal price path for a route with median distance and income

Table 3: Estimates of price elasticities

	Market Demand in Period:				
Price in Period:	$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$
$t = 1$	-2.634	0.598	0.647	0.562	0.013
$t = 2$	0.549	-6.178	1.596	1.388	0.033
$t = 3$	0.546	1.467	-10.923	2.707	0.072
$t = 4$	0.448	1.207	2.560	-16.538	0.193
$t = 5$	0.034	0.099	0.241	0.695	-2.654

is so high that only business travelers can afford it. By contrast, in periods 3 and 4, the estimates of elasticities indicate that the maximum revenue is not achieved. As we can see from the quantity estimates in Figure 4, both groups are buying tickets at the optimal price in these periods.

6.3 Welfare Estimates

Compared to the efficient supply and allocation of seats, the model's profit-maximizing allocation predicts that travelers and the firm attain 79% of the maximum gains from trade. That the gains are below 100% is due market power distortions and misallocations due to price discrimination. Figure 5 shows the distribution of utilities for two groups of travelers who are able to fly on the day of departure. The first group includes travelers who bought tickets, the second group are travelers who hadn't buy tickets because of high prices. If the allocation were efficient, only travelers who value tickets the most would end up buying it. As we can see from the figure, there is an overlap in the supports of these two distributions. This fact indicates that the optimal price path leads to

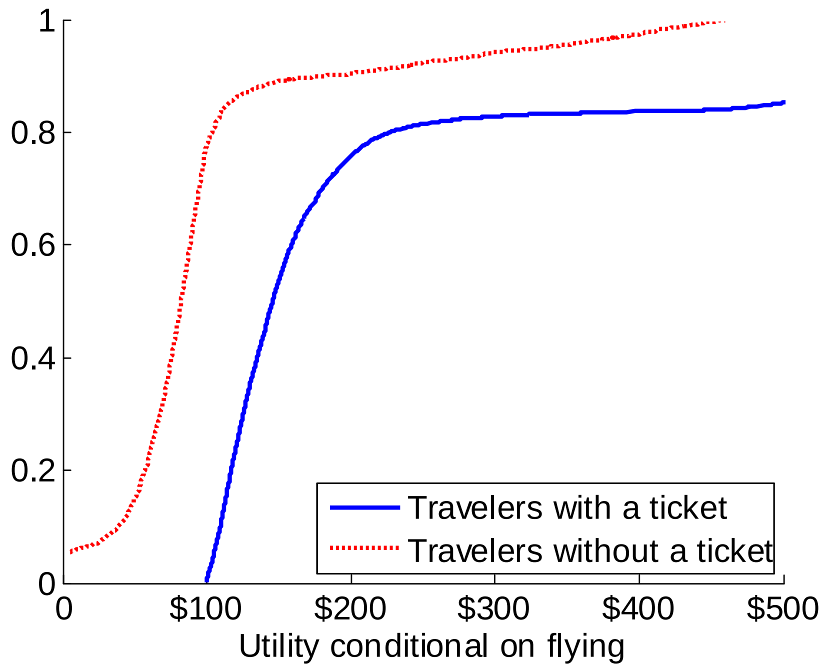

Figure 5: Cumulative distributions of travelers' utilities under the optimal allocation of seats

misallocations of seats.

There exist alternative market designs that can eliminate or at least reduce the misallocations of seats. The next section presents the counterfactual analyses.

7 Counterfactual Simulations

In the counterfactual simulations, I consider three alternative market designs that can eliminate some types of inefficiency caused by intertemporal price discrimination. The first scenario allows costless resale in the presence of market arbitrageurs. Under this assumption, two types of inefficiency would disappear: quality distortions and misallocations among the consumers. On the other hand, the third type of inefficiency, inefficiency in the quantity of production, could increase. In the second scenario, the airline is allowed to sell only fully refundable tickets. This restriction eliminates one type of inefficiency, quality distortions. By doing so, it reduces the firm's ability to price discriminate, and therefore, decreases allocative inefficiency. However, the restriction can increase inefficiency in the quantity of production. The last scenario investigates the welfare properties of the cancellation fee.

7.1 Scenario 1: Costless resale

To study the effects of a potential secondary market, I modify the fare model in the following way. In addition to travelers and the airline, I assume there exists an unlimited number of arbitrageurs. In any period, an arbitrageur can buy a ticket from the airline and then sell it to travelers later. The arbitrageurs are price-takers. Their goal is to maximize the difference between the price at which they buy a ticket and the price they sell a ticket later. Under these assumptions, the optimal price path has to be flat. I prove this statement in two steps.

Lemma 1 *For any optimal sequence of prices, the maximum profit of each arbitrageur is zero.*

Since the maximum profit of each arbitrageur is zero, the optimal price path cannot be increasing. But could it be profitable for the airline to decrease the prices? The next proposition addresses this question.

Proposition 5 *If the optimal price path in a market without resale is increasing, then the optimal price path in a market with costless resale is flat.*

Thus, to calculate the optimal fare in the counterfactual scenario, it is sufficient to consider the profit maximization problem assuming that the price path is flat. The share of type-i buyers who arrive in period τ and purchase a ticket in period t becomes:

$$s_{it\tau} = \frac{\exp\left(\frac{\mu_i - p}{\sigma_i}\right)}{1 + \sum_{k=\tau}^{T} \exp\left(\frac{\mu_i - p}{\sigma_i}\right)} = \frac{\exp\left(\frac{\mu_i - p}{\sigma_i}\right)}{1 + (T - \tau + 1)\exp\left(\frac{\mu_i - p}{\sigma_i}\right)}.$$

This share is the same for all purchase periods t since travelers pay the same price in all periods and can get a full refund if they have to cancel their tickets. The airline's profit is now equal to:

$$\pi\left(p; \tilde{\theta}\right) = (p - \tilde{c}) \sum_{i=1}^{I} \sum_{t=1}^{T} \delta_i^{T-t} D_{it}.$$

The welfare effects of ticket resale are unclear because the ability to resell tickets eliminates the inefficiency in quality of production and the flat optimal price eliminates inefficiency in allocation. However, inefficiency in the quantity of production may go up since the airline is not able to price discriminate.

To quantify the net effect on social welfare, we need to know the actual value of demand and cost parameters $\tilde{\theta}$. Figure 6 shows the optimal price path on the previous route I considered that

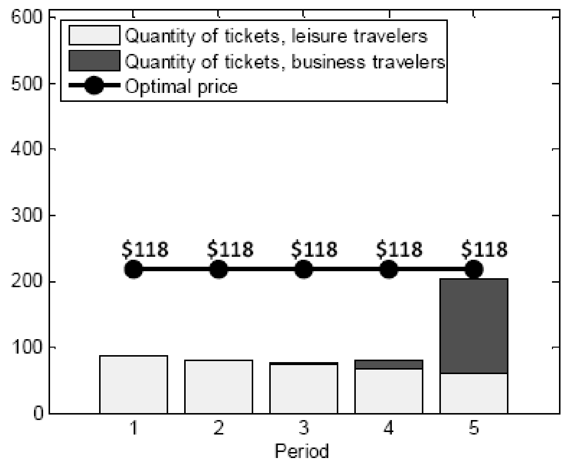

	Allocation			
	Current		Resale	
	L	B	**L**	**B**
Quantity, tickets	0.434	0.107	**0.368**	**0.160**
Quantity, seats	0.329	0.104	**0.296**	**0.158**
Average price	$77	$382	**$118**	**$118**
Expected net utility	$22	$195	**$18**	**$389**
Profit	$73		**$52**	
Social Welfare	$129.6		**$145.6**	

Note: All numbers are divided by the number of potential travelers.

Figure 6: Results of counterfactual simulations for a market with costless resale

has median characteristics and travelers depart on a non-holiday date. If resale were possible, the average price of a ticket bought by leisure travelers would increase from $77 to $118, while the average price of a ticket purchased by business travelers would decrease from $382 to $118. Leisure travelers buy approximately the same amount of tickets in each period of sale. The number of seats occupied by them would correspondingly decrease by 10%. The number of seats occupied by business travelers would go up by 50%. The consumer surplus of business travelers would increase by almost 100%. The airline's profit would decrease by 28%. Overall, social welfare on the average route would increase by 12%.

7.2 Scenario 2: Zero cancellation fee

A zero cancellation fee achieves the socially optimal level of ticket quality. On the other hand, the airline loses one of its screening tools, which makes price discrimination more difficult.

With a zero cancellation fee, the expected value of a refund is equal to $R_{i\tau} = \left(1 - \delta_i^{T-\tau}\right) p_\tau$, changing both individual demand functions and the airline's profit. The share of type-i buyers who arrived in period τ and purchase a ticket in period t now becomes:

$$s_{it\tau} = \frac{\exp\left(\frac{\mu_i - p_t}{\sigma_i}\right)}{1 + \sum_{k=\tau}^{T} \exp\left(\frac{\mu_i - p_k}{\sigma_i}\right)},$$

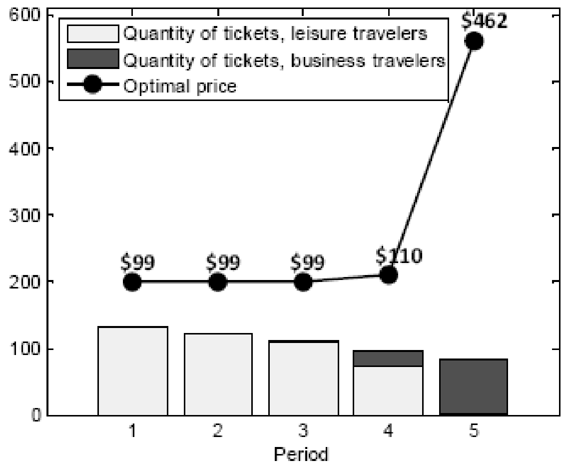

	Allocation			
	Current		**Zero fee**	
	L	B	**L**	**B**
Quantity, tickets	0.434	0.107	**0.434**	**0.102**
Quantity, seats	0.329	0.104	**0.329**	**0.104**
Average price	$77	$382	**$101**	**$384**
Expected net utility	$22	$195	**$22**	**$196**
Profit	$73		**$72**	
Social Welfare	$129.6		**$129.7**	

Note: All numbers are divided by the number of potential travelers.

Figure 7: Results of counterfactual simulations for a market with zero cancellation fee

while the airline's profit is equal to:

$$\pi\left(\mathbf{p};\tilde{\theta}\right)=\sum_{i=1}^{I}\sum_{t=1}^{T}\delta_i^{T-t}\left(p_t-\tilde{c}\right)D_{it}.$$

With a zero cancellation fee, the optimal price path becomes flatter. As a result the inefficiency in allocation goes down but inefficiency in the quantity of production may go up. The net effect on social welfare is theoretically ambiguous and depends on the value of demand and cost parameters $\tilde{\theta}$.

Figure 7 shows the optimal price path on a route with median distance and income departing on a non-holiday date. With zero cancellation fee, the difference between average prices paid by business and leisure travelers would go down from $305 to $273. This decrease is mainly caused by the fact that the average price that leisure travelers pay goes up. The reason why leisure travelers would be willing to accept higher prices is the better quality of airline tickets. The consumer surplus of both groups would go up slightly while the airline's profit would go down. Overall, social welfare would increase but by a small amount (less than 1%). This result is not too surprising as the airline does not really need to separate business and leisure travelers as most business travelers are estimated to arrive in the last periods.

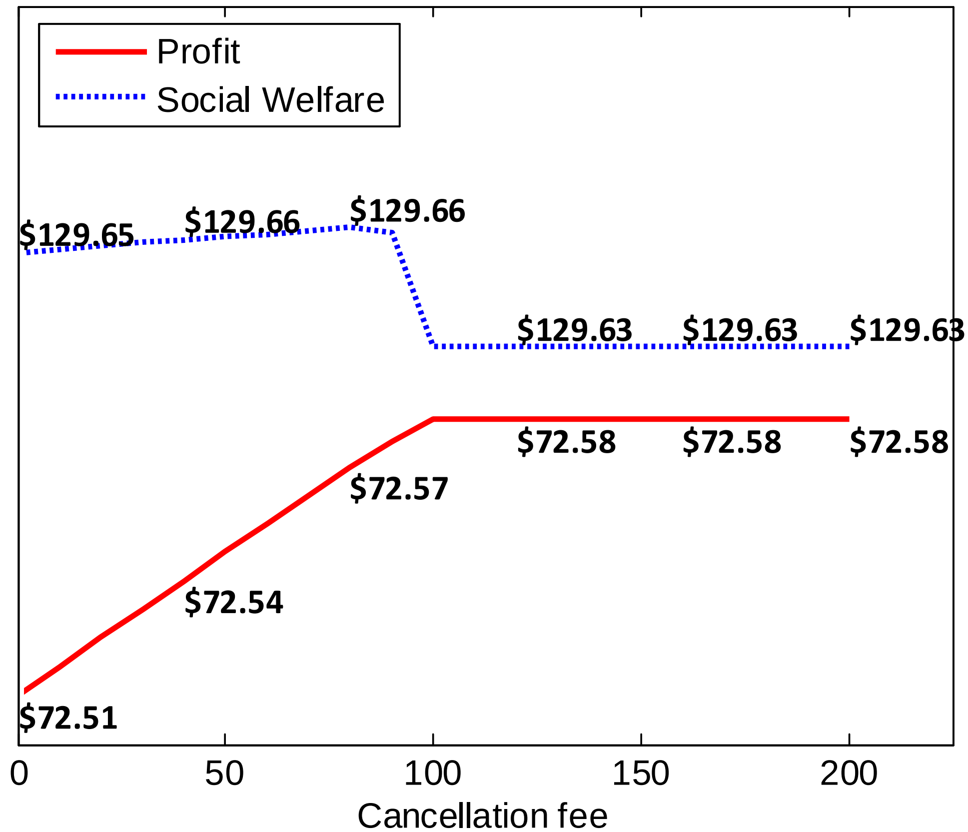

Figure 8: Profit and social welfare as functions of cancellation fee

7.3 Scenario 3: Endogenous cancellation fee

The cancellation fee has two effects on social welfare. Directly, it affects the quality of production. Indirectly, it also affects the allocation and supply of tickets as it changes the airline's ability to price discriminate over time. The profit-maximizing cancellation fee is positive. Intuitively, an airline faces the following trade-off: a higher change fee decreases the quality of a ticket but it also increases the firm's ability to price discriminate.

In practice, airlines impose the same cancellation fee for restricted tickets on all domestic markets they serve. For markets in my dataset, this fee often exceeds a one-way fare, making tickets effectively non-refundable. This fact suggests that there may exist a lower cancellation fee that would increase both consumer surplus and the airline's profit in these markets.[18] To investigate

[18] The next version of the paper will show how the maximum profit changes if the airline charges the cancellation fee as a percentage of the ticket price.

this conjecture, I solved for the optimal price paths for different levels of the cancellation fee. Figure 8 shows the values of maximum profit and social welfare calculated for different levels of the cancellation fee. As the level of the cancellation fee goes up, both social welfare and the airline's profit increase. When the cancellation fee reaches \$80, social welfare starts decreasing, while profit continues to grow until the cancellation fee reaches \$100. After this level of the cancellation fee, both social welfare and profit are unchanged. This pattern for social welfare and profit suggests three conclusions. First, in the presence of market power distortions, both profit-maximizing and socially optimal levels of the cancellation fee are positive. Second, the profit-maximizing level of the fee exceeds its socially optimal level. Third, the current level of cancellation fee (\$150 for most airlines) exceeds the socially optimal level but does not contradict the hypothesis that airlines are profit-maximizing.

8 Conclusion

In this paper, I developed an empirical model of optimal fares and estimated it using new data on daily ticket prices from domestic monopoly markets. The estimates of demand and cost parameters for monopoly routes allowed me to quantify the costs and benefits of intertemporal price discrimination. I found that intertemporal price discrimination results in a lower ticket quality for leisure travelers, higher prices for business travelers, lower supply of tickets for business travelers, lower overall supply and misallocations of tickets among travelers. On the other hand, the benefits of intertemporal price discrimination for leisure travelers are lower prices and higher supply.

I also found that free resale of airline tickets would reduce airlines' ability to price discriminate over time. As a result, business travelers would win from resale and leisure travelers would lose, even though the quality of tickets would improve. Overall, the short-run effect of ticket resale on social welfare is positive. However, since the airline's profit goes down, it may choose to exit from the market in the long run. The effect of the cancellation fee on social welfare is small. Prices would go up mainly due to an increase in ticket quality. Finally, I found that when firms have market power, both socially efficient and profit maximizing levels of the cancellation fee are positive. The current level of the cancellation fee exceeds the socially optimal level but does not contradict the profit-maximizing behavior of airlines.

References

Aguirre, I., S. Cowan, and J. Vickers (2010): "Monopoly Price Discrimination and Demand Curvature," *The American Economic Review*, 100(4), 1601–1615.

Bly, L. (2001): "Web Site for Resale of Air Tickets Shows Promise, but It's Still a Fledgling," *Los Angeles Times. June 3, 2001.*

Board, S. (2008): "Durable Goods Monopoly with Varying Demand," *Review of Economic Studies*, 75(2), 391–413.

Board, S., and A. Skrzypacz (2010): "Revenue Management with Forward-Looking Buyers," Discussion paper, Stanford University.

Borenstein, S., and N. L. Rose (1994): "Competition and Price Dispersion in the U.S. Airline Industry," *Journal of Political Economy*, 102(4), 653–683.

Chen, J., S. Esteban, and M. Shum (2010): "How much competition is a secondary market?," *Working Paper.*

Curtis, W. (2007): "World of Wonders," *The New York Times. December 22, 2007.*

Deneckere, R. J., and R. Preston McAfee (1996): "Damaged Goods," *Journal of Economics & Management Strategy*, 5(2), 149–174.

Elliot, C. (2011): "Ridiculous or not? Your airline ticket isn't transferrable," *Consumer advocate Christopher Elliott's site. March 16, 2011.*

Esteban, S., and M. Shum (2007): "Durable-goods oligopoly with secondary markets: the case of automobiles," *The RAND Journal of Economics*, 38(2), 332–354.

Gavazza, A., A. Lizzeri, and N. Roketskiy (2010): "A Quantitative Analysis of the Used Car Market," Discussion paper, mimeo, New York University.

Gerardi, K. S., and A. H. Shapiro (2009): "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry," *Journal of Political Economy*, 117(1), 1–37.

Gershkov, A., and B. Moldovanu (2009): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*, 1(2), 168–198.

Hendel, I., and A. Nevo (2011): "Intertemporal Price Discrimination in Storable Goods Markets," *National Bureau of Economic Research Working Paper Series*, No. 16988.

Hoerner, J., and L. Samuelson (2011): "Managing Strategic Buyers," *Journal of Political Economy*, 119(3), 379–425.

Leslie, P. (2004): "Price Discrimination in Broadway Theater," *The RAND Journal of Economics*, 35(3), 520–541.

Leslie, P., and A. Sorensen (2009): "The Welfare Effects of Ticket Resale," *National Bureau of Economic Research Working Paper Series*, No. 15476.

McAfee, R. P., and V. te Velde (2007): "Dynamic Pricing in the Airline Industry," *Handbook on Economics and Information Systems*, 1.

Mussa, M., and S. Rosen (1978): "Monopoly and product quality," *Journal of Economic Theory*, 18(2), 301–317.

Nair, H. (2007): "Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games," *Quantitative Marketing and Economics*, 5, 239–292.

Phillips, R. L. (2005): *Pricing and revenue optimization*. Stanford University Press.

Robinson, J. (1933): "The Economics of Imperfect Competition," .

Schmalensee, R. (1981): "Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination," *The American Economic Review*, 71(1), 242–247.

Shepard, A. (1991): "Price Discrimination and Retail Configuration," *Journal of Political Economy*, 99(1), 30–53.

Stavins, J. (2001): "Price Discrimination in the Airline Market: The Effect of Market Concentration," *The Review of Economics and Statistics*, 83(1), 200–202.

Stokey, N. L. (1979): "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, 93(3), 355 –371.

Sweeting, A. (2010): "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets," Discussion paper, mimeo.

Talluri, K., and G. Van Ryzin (2005): *The theory and practice of revenue management*, vol. 68. Springer Verlag.

Varian, H. R. (1985): "Price Discrimination and Social Welfare," *The American Economic Review*, 75(4), 870–875.

Verboven, F. (1996): "International Price Discrimination in the European Car Market," *The RAND Journal of Economics*, 27(2), 240–268.

Villas-Boas, S. B. (2009): "An empirical investigation of the welfare effects of banning wholesale price discrimination," *The RAND Journal of Economics*, 40(1), 20–46.

9 Appendix

[TBA]