The Welfare Effects of Intertemporal Price Discrimination:

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

John Lazarev[*]

This version: January 7, 2013

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 in the model 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.

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 also would 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 customers willing to pay more 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 public 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 seat on a given flight from variation for similar seats on different days of departure. McAfee and te Velde (2007) is one of the few attempts to

---

[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), Schwartz (1990), Aguirre et al (2010), and others have analyzed these welfare effects in varying degrees of generality.

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

use airline data to analyze intertemporal price discrimination. They had a sample of price paths, but they did not have access to the corresponding quantities of seats sold. 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 too complicated to estimate. Among other difficulties, one would have to deal with the multiplicity of equilibrium predictions and account for multimarket contact the presence of which is well documented in the industry (see e.g. Evans and Kessides (1994)). 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 identify and segregate different consumer groups, but is able to charge different prices in different periods of sale. Aggregate demand uncertainty is not modeled explicitly.[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

---

[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.

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 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. More than 90% 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 significantly 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%,

---

[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.

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 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, the third scenario compares the welfare properties of intertemporal and third-degree price discrimination. Third degree price discrimination implies that the airline can identify the customers' types and is able to set different prices to different types. By varying the price over time, the airline captures more than 90% of the profit that it would receive if third degree price discrimination was possible. Surprisingly, the estimates show that some customer groups would prefer third-degree price discrimination to intertemporal price discrimination. Total social welfare is also higher under third degree price discrimination.

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. In Section 5, I show 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. In Section 7, I formally describe 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 any 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 substantial 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 1, 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 Tuesday, March 1st, 2011. American Airlines is the only carrier that serves this route; there are

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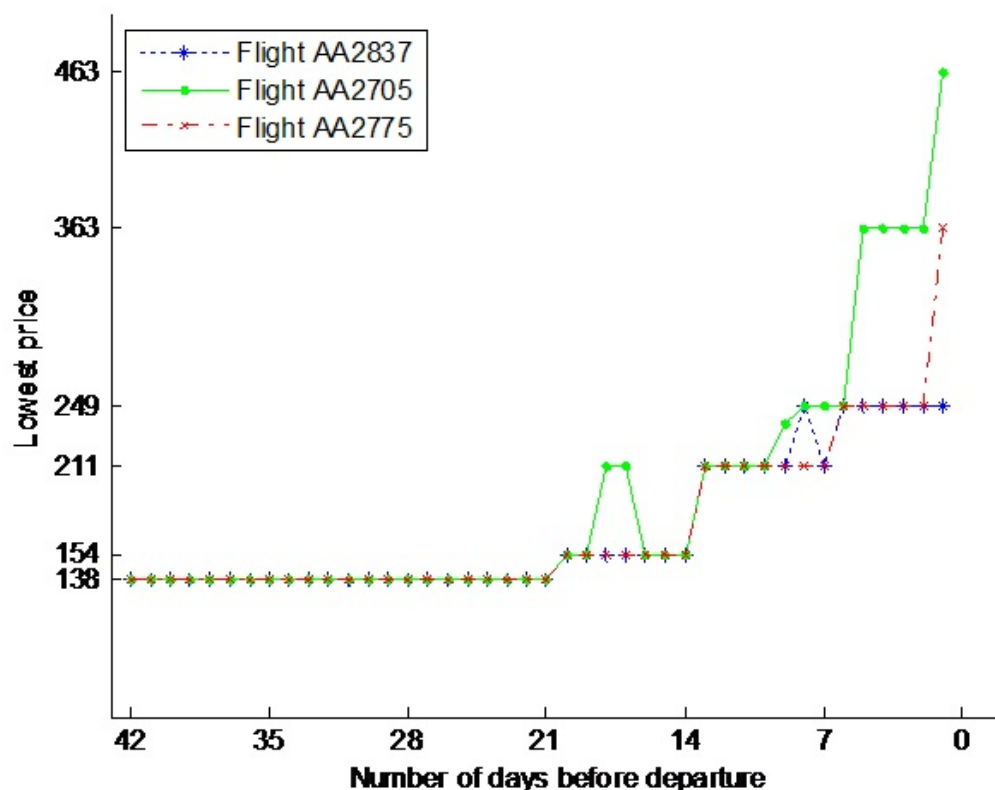
Figure 2: Example Price Path. Route: Dallas, TX - Roswell, NM. Departure Date: 03/01/11

three flights available during that day.

The behavior of ticket prices depicted is representative of monopoly markets in my data. There are three main stylized facts. 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 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 1 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 rivals very closely: pricing departments respond to competitors' price moves very quickly, often responding 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

To calculate the effect of intertemporal price discrimination on consumer welfare, we need to estimate consumers' demand functions. The demand system is estimated using assumptions about pricing and the supply side. To recover consumers' preferences (or, to be more precise, the airline's expectations about consumers' preferences), I develop a model that shows how a set of parameters reflecting travelers' preferences transforms into a path of profit-maximizing fares.[6]

A theoretical model that is able to generate the stylized facts listed in Section 2 has to include the decision problems of both the pricing and revenue-management departments. The solution of the pricing department's problem is a finite set of fares that include advance purchase requirements. To construct an optimal set of fares, the pricing department has to calculate the value of the airline's expected profit for each possible set of fares. This value, in turn, depends on the strategy

---

[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.

of the revenue management department that takes the set of fares as given and updates availability of each booking class in real time. Another complication comes from the fact that the airline has to take into account not only direct passengers that travel on a particular route but also passengers for whom this route is only a part of their trip. I will call them "direct passengers" and "connecting passengers", respectively. The model is initially formulated for a representative origin and destination and a representative departure date.

3.1 Airline's problem

Consider a representative market that is defined by three elements: origin, destination and travel date. The airline is the only producer in the market. It can offer up to C seats on its flights from the origin to the destination. It flies both direct and connecting passengers. For direct passengers, the origin is the initial point of their trip and the destination is the final point of their trip. For connecting passengers, this flight is only a part of their trip.

The airline is selling tickets during a fixed period of time. Advance purchase requirements divide this period into T periods of sale. At the beginning of the first period of sale, the airline's pricing department sets a menu of fares for this market $\mathbf{p} = (p_1, ..., p_T)$ and for all markets that connecting passengers fly $\mathbf{p}_j = (p_{j1}, ..., p_{jT})$. The price p_t is the price of the cheapest fare that satisfies the advance purchase requirement for period of sale t. In the empirical application, advance period requirements observed define five periods of sale: 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.

The revenue management department at each moment of time decides which of the fares that satisfy the advance purchase requirements to offer for purchase based on the information ξ_t. Denote by $\tilde{D}_t(\mathbf{p}, \xi_t)$ the number of tickets that the airline sells at price p_t. Not all passengers that bought tickets will end up flying. Denote by $\tilde{Q}_t(\mathbf{p}, \xi_T)$ the number of seats that that will be occupied by passengers who bought tickets at price p_t. Both $\tilde{D}_t$ and $\tilde{Q}_t$ are the solutions of the revenue management department's problem.

The airline's revenue comes from selling tickets and collecting cancellation fees. If a traveler needs to cancel a ticket, she has to pay a cancellation fee 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. The airline's operational cost, $\varphi(\cdot)$, depends on the total number of enplaned passengers. Thus,

the airline's profit takes the following form:

$$\pi = \tilde{R} + \sum_j \tilde{R}_j - \varphi \left(\tilde{Q} + \sum_j \tilde{Q}_j \right),$$

where

$$
\begin{aligned}
R &= \sum_{t=1}^{T} \left(p_t \tilde{Q}_t + \min\left(f, p_t\right)\left(\tilde{D}_t - \tilde{Q}_t\right)\right) \text{ revenue from direct passengers,} \\
R_j &= \sum_{t=1}^{T} \left(p_{jt} \tilde{Q}_{jt} + \min\left(f, p_{jt}\right)\left(\tilde{D}_{jt} - \tilde{Q}_{jt}\right)\right) \text{ revenue from connecting passengers,} \\
\tilde{Q} &= \sum_{t=1}^{T} \tilde{Q}_t \text{ the number of seats occupied by direct passengers,} \\
\tilde{Q}_j &= \sum_{t=1}^{T} \tilde{Q}_{jt} \text{ the number of seats occupied connecting passengers from market } j.
\end{aligned}
$$

The pricing department chooses menus of direct fares $\mathbf{p}$ and connecting fares $\mathbf{p}_j$ to maximize the expected value of the profit function subject to the capacity constraint. Formally, the profit maximization problem takes the following form:

$$\max_{p, p_j} \mathbb{E}_0 \pi \text{ s.t. } \tilde{Q} + \sum_j \tilde{Q}_j \leq C.$$

The expectation is taken with respect to all information available at the beginning of the first period of sale.

I will simplify the problem in three steps. First, the constrained optimization problem can be written as unconstrained using the method of Lagrange multipliers. Let $\phi\left(C\right)$ denote the value of the Lagrange multiplier that corresponds to the capacity constraint. Then the unconstrained profit function takes the following form:

$$\pi = \tilde{R} + \sum_j \tilde{R}_j - \varphi \left(\tilde{Q} + \sum_j \tilde{Q}_j \right) - \phi\left(C\right)\left[\tilde{Q} + \sum_j \tilde{Q}_j - C \right].$$

The last two components of the profit function represent the economic cost of the airline. The $\varphi\left(\cdot\right)$ term is the operational cost, the $\phi\left(\cdot\right)$ term is the shadow cost of capacity. Denote by $\tilde{c}$ the value of the marginal economic cost evaluated at the profit-maximizing level. Then, the solution

11

of the original profit maximizing problem coincides with the solution of the following problem:

$$\max_{\mathbf{p},\mathbf{p}_j} \mathbb{E}_0 \left[R + \sum_j R_j - \tilde{c} \cdot \left(\tilde{Q} + \sum_j \tilde{Q}_j \right) \right].$$

The last problem is separable with respect to $\mathbf{p}$ and $\mathbf{p}_j$, i.e.

$$\mathbb{E}_0 \left[R + \sum R_j - \tilde{c} \cdot \left(\tilde{Q} + \sum_j \tilde{Q}_j \right) \right] = \mathbb{E}_0 \left[R - \tilde{c}\tilde{Q} \right] + \sum_j \mathbb{E}_0 \left[R_j - \tilde{c}\tilde{Q}_j \right].$$

Thus, if the optimal value of the expected marginal cost $\tilde{c}$ is known, then it is sufficient to solve the profit-maximization problem for direct passengers without looking at the fares set for connecting passengers or knowing the value of the capacity constraint. The value $\tilde{c}$ can be interpreted in two ways. First, it reflects the expected marginal revenue of adding an additional unit of capacity to the market. Second, it is equal to the marginal revenue of flying connecting passengers.

Finally, consider the profit-maximization problem for direct passengers:

$$\max_{\mathbf{p}} \mathbb{E}_0 \left[\tilde{R} - \tilde{c}\tilde{Q} \right] = \max_{\mathbf{p}} \mathbb{E}_0 \left[\sum_{t=1}^{T} p_t \tilde{Q}_t + \min\left(f, p_t\right) \left(\tilde{D}_t - \tilde{Q}_t \right) - \tilde{c}\tilde{Q}_t \right].$$

By the law of iterated expectations, we can rewrite this problem as:

$$\max_{\mathbf{p}} \sum_{t=1}^{T} \left[\left(p_t - \tilde{c} \right) Q_t + \min\left(f, p_t\right) \left(D_t - Q_t \right) \right],$$

where $Q_t = \mathbb{E}_0 \tilde{Q}_t$ and $D_t = \mathbb{E}_0 \tilde{D}_t$. The function D_t is the expected number of tickets that will be sold at price p_t if the pricing department offers the menu of fares $\mathbf{p}$ and then the revenue management department behaves optimally given this menu. The function Q_t is the corresponding expected number of occupied seats.

To calculate the welfare effects of intertemporal price discrimination, we need to know how the quantity of sold tickets and the number of occupied seats respond to changes in the menu of fares and the cancellation fee. In other words, we need to know the elasticities of demand with respect to the prices of all available fares and the cancellation fee. Three limitations of the data do not allow us to estimate these elasticities directly. The number of occupied seats for each fare p_t is not available for each individual flight or departure date. The data include only a 10% random sample of the quantity data aggregated to the quarterly level. Second, the data do not record tickets that were sold but later canceled. Third, it would be hard to find a source of exogenous variation that

comes from the supply side and would affect the components of the fare menu independently of each other. The form of the profit function suggests that any variation in the cost function affects the entire menu of fares in a very specific way. From the pricing department's point of view, the value of the expected marginal cost of flying an additional passenger is the same in all periods of sale. Finally, there is almost no variation in the cancellation fee in the data. Almost all airlines charged \$150 in all domestic markets.

Given these limitations, I follow a different approach. I assume that the market demand defined by Q_t and D_t reflects the optimal decision of strategic consumers whose preferences with respect to the price and time of purchase depend on a vector of demand parameters $\tilde{\theta}$. The vector of demand parameters $\tilde{\theta}$ determines the level of consumer heterogeneity, their willingness to pay for an airline ticket, their aversion of the imposed cancellation fee. The airline's pricing department knows the value of $\tilde{\theta}$ and chooses a menu of fares $\mathbf{p}$ to maximize the airline's profit defined by functions Q_t and D_t that in turn depend on $\tilde{\theta}$ and $\tilde{c}$. Using daily price data and quarterly aggregated quantity data, I will recover these parameters assuming that the observed prices maximize the airline's profit for these parameters.

A natural way to model the demand functions D_t and Q_t would be first to model the revenue-management problem explicitly. For each realization of aggregate demand uncertainty, the solution of the revenue-management problem, $\tilde{D}_t$ and $\tilde{Q}_t$, defines the optimal allocation of tickets and seats. To find this allocation, we would need to find a stochastic process that defines the evolution of prices. This process would have to be consistent with the travelers expectations and maximize the airline's profit given these expectations at the same time. So, we would also need to model travelers' expectations explicitly as well as how much the travelers know about the remaining capacity on a given flight in each period of sale. All these assumptions will not be possible to test because the data contain no information on how far in advance a particular ticket was bought. In fact, no flight-level quantity data are publicly available.

Since the object of interest are the expected quantities D_t and Q_t, I propose a simpler model of traveler behavior. This model gives a closed form solution that approximates the expected demand functions D_t and Q_t. This approximation captures all relevant features of the data available to us and thus allow us to make counterfactual predictions under the assumption that the process that defines aggregate demand uncertainty does not change in the counterfactual scenario.

13

3.2 Demand System and Consumer Welfare

This subsection describes how the vector of demand parameters $\tilde{\theta}$ determines the relationship between the expected quantities of sold tickets D_t, the occupied seats Q_t, and the menu of offered fares $\mathbf{p}$. It can be viewed as a micro model of the market demand functions $Q_t\left(\mathbf{p}; \tilde{\theta}\right)$ and $D_t\left(\mathbf{p}; \tilde{\theta}\right)$. Since these functions by construction represent expected quantities, the model does not allow any demand uncertainty at the market level.

Types, Arrival and Exit   The population of potential direct passengers of size $\tilde{M}$ 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:

$$u_{\iota i t} \equiv \mu_i + \sigma_i \left(\varepsilon_{\iota i t} - \varepsilon_{\iota i 0} \right), \tag{1}$$

where μ_i is type-i's mean utility from flying on this route measured in dollar terms, $\varepsilon_{\iota i t}$ are i.i.d. Type-1 extreme value terms that shift traveler ι's utility in each period, and σ_i is a normalizing coefficient that controls the variance of $\varepsilon_{\iota i t}$. The error term $\varepsilon_{\iota i t}$ reflects idiosyncratic customers' preferences with respect to the time of purchase. They may reflect customers' tastes with regard to other characteristics of restricted fares or their idiosyncratic level of uncertainty about their travel plans. The errors represent the consumer tastes that the airline and researcher do not observe.

---

[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.

This coefficient σ_i captures the slope of the demand curve and hence the price sensitivity across the population of type-i travelers: the lower the coefficient, the less sensitive are type-i travelers. The traveler learns all components of their utilities defined in equation (1) at the beginning of the period she arrived in the market.[8]

After purchase, the traveler can cancel a ticket. If she cancels a ticket in period 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 $\max(p_t - f, 0)$. 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 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 if the airline increases 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:

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

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

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

Suppose the traveler decides to wait 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 $\delta_i^{T-\tau'}\left[\mu_i + \sigma_i\left(\varepsilon_{\iota i\tau'} - \varepsilon_{\iota i0}\right)\right] + \rho_{i\tau'} - p_{\tau'}$. In this case, her expected utility is equal to

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

---

[8]An alternative assumption would be for travelers to learn a component of $\varepsilon_{\iota i}$ before each period of sale.Under this assumption each customer would compare the current value of the term with its expected future values. Under the original assumpton each customer would compare this value with its actual future values. Qualitatively we would receive the same results. However, the demand function will not have a closed form solution.

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

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

$$\text{for all } \tau < \tau' \leq T \text{ and}$$

$$\delta_i^{T-\tau}\left[\mu_i + \sigma_i\left(\varepsilon_{\iota i\tau} - \varepsilon_{\iota i0}\right)\right] + \rho_{i\tau} - p_\tau \;>\; 0.$$

These inequalities can be rewritten in a more convenient way:

$$\frac{\delta_i^{T-\tau}\mu_i + \rho_{i\tau} - p_\tau}{\sigma_i\delta_i^{T-\tau}} + \varepsilon_{\iota i\tau} \;>\; \frac{\delta_i^{T-\tau'}\mu_i + \rho_{i\tau'} - p_{\tau'}}{\sigma_i\delta_i^{T-\tau'}} + \varepsilon_{\iota i\tau'} \text{ for all } \tau < \tau' \leq T \text{ and} \qquad (2)$$

$$\frac{\delta_i^{T-\tau}\mu_i + \rho_{i\tau} - p_\tau}{\sigma_i\delta_i^{T-\tau}} + \varepsilon_{\iota i\tau} \;>\; \varepsilon_{\iota i0}.$$

Market demand for airline tickets 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

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

Consider 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:

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

Thus, the vector of demand parameters $\tilde{\theta}$ includes the following parameters: shares of each customer type γ_i, the mean utilities μ_i, the price sensitivity σ_i, the probability of cancellation δ_i, the arrival parameters λ_{it}.

Number of occupied seats The probability of not cancelling a trip for traveller of type i who bought a ticket in period t by the time of departure is given by δ_i^{T-t}. Thus the number of occupied seats is equal to

$$Q = \sum_{t=1}^{T} Q_t, \text{ where } Q_t = \sum_{i=1}^{I} \delta_i^{T-t} D_{it}.$$

Welfare[9] 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,

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

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

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

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 an efficient supply and allocation of seats. By the First Welfare Theorem, the allocation of seats will be efficient only if all consumers take the same prices into account. If it is not the case, then there could be two customers who would be willing to trade with each other right before departure. The reason why the customer who wants to buy the ticket now didn't buy it before was his higher probability of cancellation. Therefore, there is always some positive probability that the ex-post allocation is not efficient, therefore any price path with a positive cancellation fee is not efficient.

Thus, there are three conditions for efficient supply and allocation of seats. First, the price path has to be flat. Second, it has to equal to the value of the marginal costs $\tilde{c}$. Third, the cancellation fee has to be zero. If the cancellation fee is positive, then the expected value of the refund is

---

[9]Given the data limitations, I can only estimate the welfare effects of intertemporal price discrimination on direct passengers.

different for customers of different types. This fact implies that even thought the airline offers the same menu of fares to all customers, the effective ex-post price is different for different customer types.

These conditions illustrate two impediments to the efficient supply and allocation of seats: market power and dynamic pricing. First, if 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.* A positive cancellation fee makes a ticket less attractive to travelers. For this reason, I refer to it as a measure of ticket quality. A positive cancellation fee thus implies *inefficiency in the quality of production.* 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.

3.3 Optimal Price Path

A price path $\mathbf{p}$ is called *optimal* if it maximizes the airline's profit $\pi(\mathbf{p})$:

$$\pi(\mathbf{p}) = \sum_{t=1}^{T} \left[(p_t - \tilde{c}) Q_t + \min(f, p_t)(D_t - Q_t) \right]$$

Denote by $\mathbf{p}^* \left(\tilde{\theta}, c \right)$ the optimal price path as a function of the demand parameter $\tilde{\theta}$ and the cost parameter $\tilde{c}$.

Except for a knife-edge realization of the demand and cost parameters, the optimal price path implies intertemporal price discrimination, i.e. prices differ in different periods. Furthermore, in practice, airlines often impose a positive cancellation fee for lower fares. 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.

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 evaluate the welfare losses associated with each type of inefficiency, we need to know the estimates of the demand parameter $\tilde{\theta}$ and cost parameter $\tilde{c}$. I will estimate these parameters using a sample of optimal price paths and corresponding quantities.

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). 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	$205	$236
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

Fares are distributed by the Airline Tariff Publishing Company[10] (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.[11] 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.[12] 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.

I consider three quarters of departure dates between October 1, 2010 and June 30, 2011. Besides

---

[10]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 Transport and the Canadian Transportation Agency to distribute data for airlines operating in the region.

[11]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.

[12]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.

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 issued in the U.S. 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.

5   Estimation

5.1   Econometric Specification

My empirical model allows for two types of travelers. I refer to the first type as leisure travelers (L), and to the second type as business travelers (B). Leisure travelers are highly price sensitive customers who are willing 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.[13] 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 parameters $\tilde{\theta}_{nd}$ and the cost parameter $\tilde{c}_{nd}$ determine the optimal price path $\mathbf{p}^* \left(\tilde{\theta}_{nd}, \tilde{c}_{nd} \right)$. These parameters are known to the airline but unknown to the researcher. The goal of the estimation routine is to recover $\tilde{\theta}_{nd}$ and $\tilde{c}_{nd}$ 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.[14] 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

---

[13]See, Phillips (2005).

[14]This assumption is very restrictive. I have estimated the model on different subsamples of routes for different airlines. The results are qualitatively similar.

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. 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 $\varepsilon_{\lambda 1 nd}$ is normalized to 0 and $\varepsilon_{\lambda 2 nd}, ..., \varepsilon_{\lambda Tnd}$ 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 value of the expected marginal costs $\tilde{c}_{nd}$, by construction, is equal to the derivative of the total economic costs evaluated at the profit-maximizing level of the total quantity of occupied seats. The economic costs include both the operational costs and the shadow costs of capacity. If the total quantity of occupied seats were available, then the most natural way to estimate $\tilde{c}$ would be as nonparametric function of the total quantity. I do not observe this quantity, so I estimate the average value of the marginal costs by assuming that $\tilde{c} = c + \varepsilon_{cnd}$ where ε_{cnd} is a

mean-zero deviation of the actual value from its mean. 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. It captures factors that affect both the operational costs (such us distance, capacity, etc.), and the shadow cost of the capacity constraint (the demand of connecting passengers etc.). This error shifts the entire time path of prices, while $\varepsilon_{\lambda \tau nd}$ affects relative levels of the prices in the path.

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 $\varepsilon_{\lambda \tau nd}$, and the cost errors ε_{cnd} are jointly independent.

Together, we can divide all demand and cost parameters known to the airline into three groups: estimated coefficients $\theta = (\gamma, \mu, \sigma, \delta, \alpha)$, c, and M_n, errors unobserved to the researcher $\varepsilon_{nd} = (\varepsilon_{\lambda nd}, \varepsilon_{nd})$, 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 parameter θ and cost parameters c, I follow the standard practice of using both price and quantity data. However, 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:

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

23

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}. It turns out that there exists a unique mapping $g_P : R^T \times R^{\dim(\theta)} \times R^{\dim(h_d)} \times R^{\dim(x_n)} \to R^T$, such that for any θ, it holds that $G(p_{nd}, \theta, h_d, x_n, g_P(p_{nd}, \theta, h_d, x_n)) = 0$. The proof of this statement follows from the fact that the system of first order conditions is triangular and linear with respect to the errors. The first equation includes only ε_{cnd}, the second equation includes ε_{cnd} and $\varepsilon_{\lambda 2nd.}$, etc. Thus, we can invert the system by the substitution method: derive the value of ε_{cnd} from the first equation and plug it into the second one, etc.

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

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

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

5.2.2 Monthly traffic

The model predicts the expected total number of direct passengers for departure date d and route n is equal to $\sum_{t=1}^{T} Q_{ndt}\left(p_{nd}, \tilde{\theta}\right)$. In the data, we observe the actual number of flying passengers. Denote by $Q_{nm}^{traffic}$ the total number of enplaned direct passengers observed in the data for route n and month m. Thus, the predicted number of enplaned passengers is equal to

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

Denote by $g_M\left(p_{nd}, \tilde{\theta}, M_{nm}\right) = \sum_{d \in \text{month}(m)} \sum_{i=1}^{I} \sum_{t=1}^{T} \delta_{id}^{T-t} D_{it}\left(p_{nd}, \tilde{\theta}\right) - Q_{nm}^{traffic}$. This error comes from the fact that the revenue-management department due to the stochastic nature of the demand cannot perfectly implement the plan designed by the pricing department. Sometimes it allocates more seats to a certain class, sometimes less. The goal of the revenue management department, however, is to get as close to the target level as possible. Therefore, it is not unreasonable to assume that the variance of the error is bounded and its expected value is equal to zero. Then, a moment restriction that corresponds to the observed number of enplaned passengers is given by:

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

I use this restriction as to define the 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 $p(r_{lnq})$ and $f(r_{lnq})$ denote the corresponding one-way fare and number of traveling passengers.[15] 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:[16] $[p_b, p_{b+1}]$, $b = 0, ..., B$. 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 $[p_b, p_{b+1}]$ for market n in quarter q. This probability equals:

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

Similarly, we can calculate the relative frequency of observing a ticket within a given price range using the 10% sample of airline tickets. 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. 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.

---

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

[16]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.

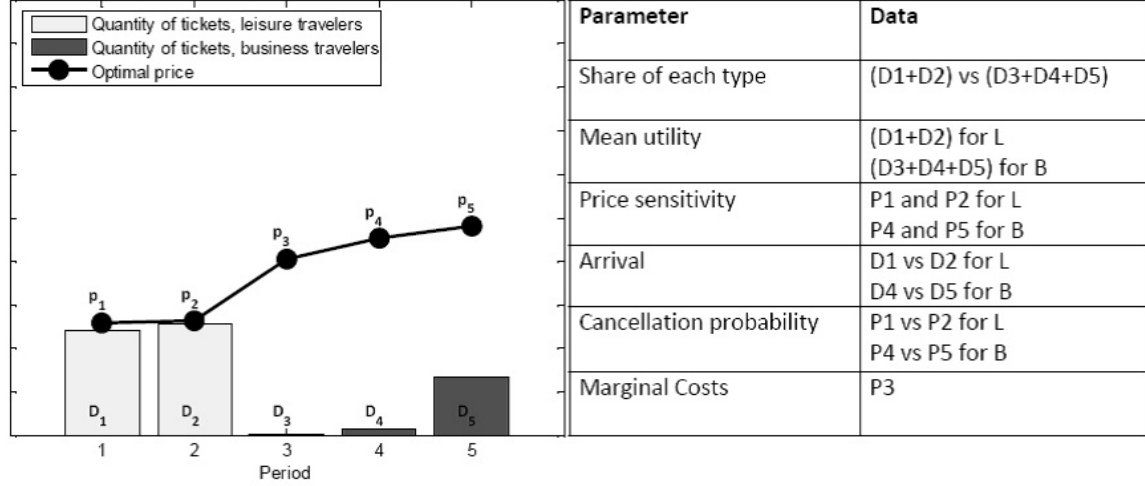| Parameter	Data
Share of each type	(D1+D2) vs (D3+D4+D5)
Mean utility	(D1+D2) for L
	(D3+D4+D5) for B
Price sensitivity	P1 and P2 for L
	P4 and P5 for B
Arrival	D1 vs D2 for L
	D4 vs D5 for B
Cancellation probability	P1 vs P2 for L
	P4 vs P5 for B
Marginal Costs	P3

Figure 3: Identification

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.

5.4 Identification

Section 5.2 established 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 the 15 parameters that define θ and c. These parameters are identified from the joint distribution of daily optimal prices and quantities aggregated to the quarterly level. To show identification formally, I would need to prove that the T moment restrictions can be satisfied only under the true parameter θ_0. This fact is rarely possible to prove without knowing the true distribution of the data.

To gain intuition on what properties of the joint distribution identify each component of the parameter θ, I performed two simulation exercises using the model of optimal fares. The first exercise shows how a change in each component of the demand and cost parameter θ affects the

profit maximizing vectors of prices and quantities. The second exercise does the opposite. After changing a component of the price-quantity vector, I find a vector of parameters θ under which the new price-quantity vector would maximize the airline's profit. Based on these results, I can provide an intuitive explanation for how the joint distribution of the data may identify the parameters of the model. The explanation is, by all means, heuristic as we should keep in mind that whenever we change one parameter of the model, all components of the profit-maximizing prices and quantities will necessarily change.

Consider a representative market. The solid line in Figure 3 shows a typical price path that we observe in the data. For the sake of argument, suppose we also observe the corresponding quantities of sold tickets for this departure day. These quantities are depicted by the bar graph on Figure 3. Thus, we know two profit maximizing vectors $p = (p_1, p_2, p_3, p_4, p_5)$ and $q = (q_1, q_2, q_3, q_4, q_5)$. From these vectors, we need to infer the following demand and cost parameters: a share of each type γ, the mean utilities μ_i, the within-type heterogeneity parameter σ_i, the probability of cancellation δ_i, the arrival parameters α_i, and the cost parameter c.

The behavior of the typical price path can be described as follows. In the first two periods, the price rises but at a relatively slow level. Then in period 3 or 4, the price jumps up and continues to increase but, again, with a slower speed. To understand this behavior, consider the tradeoff that the airline has. Recall that it faces two heterogeneous groups of customers with different marginal willingness to pay: business travelers are willing to pay more than leisure travelers. Therefore, the airline can charge a high price and receive a low quantity as most leisure travelers cannot afford to fly. Alternatively, it can charge a low price but receive a high quantity. The price path suggests that it should be profit maximizing for the airline to charge a low price in the first periods and then switch to a high price.

Having this intuition in mind, we can infer that most customers buying early are leisure (type 1) travelers, while customers who are buying later, at a higher price, are business (type 2) travelers. The exact level of the prices in early periods is determined by the elasticity of leisure travelers, while the price level in later periods is determined by the elasticity of business travelers. The elasticity of each group in turn depends on the price-sensitivity parameter σ_i. Similarly, the quantities sold in early periods reveal information about the mean utility of leisure travelers (μ_L), while the quantities sold in later periods depend on the mean utility of business travelers (μ_B). By comparing the sum

27

of quantities sold in early periods with the total sum of quantities, and taking into account the profit maximizing conditions, we can infer the share of leisure type (γ).

The increase in prices in period 2 compared to period 1 is determined by the probability of cancellation. After the first period, customers became more certain about their travel plans since there are fewer periods during which they can learn that they won't be able to fly. As a result, they are willing to pay more. The airline realizes this change and increases price. Since most customers who are buying tickets in the first two periods are leisure travelers, the change in these two prices identifies the probability of cancellation for leisure travelers (δ_L). Similarly, the probability of cancellation for business travelers (δ_B) is identified from the change in the last two prices. Further, if no new customers arrived in period 2, the profit-maximizing quantities in period 1 and 2 would be the same. Customers with a high first-period shock ε_{Li1} would buy in period 2, customers with a high second-period shock ε_{Li2} would buy in the second period. The picture suggests that this is not the case. The reason why the quantity in period 2 is higher is the arrival of new customers. For the same reason, quantities in period 4 and 5 are also different. Thus, the exact difference between the two quantities reveals the value of the arrival parameter α_i.

Finally, the period in which the price jump occurs identifies the value of the cost parameter c. Intuitively, in the equilibrium, the marginal revenue that the airline receives from business travelers should be equal to the marginal revenue it receives from leisure travelers and both should be equal to the value of marginal cost. If the costs are high, then the marginal revenue the airline receives from leisure travelers has to be higher. Therefore, fewer leisure travelers will be served in the equilibrium, so the airline has to switch to business travelers sooner. If the costs are low, then the marginal revenue from leisure travelers has to be low, so the airline will offer the lower price longer.

If the menus of fares are the same for all travel dates within a quarter, we can just divide the quarterly aggregated quantities by the number of travel dates and apply this intuition directly. Suppose that the menus of fares are the same except for one travel date, say, Thanksgiving. Then, this travel date has its own menu of fares, at least one price of which is different from the rest. We can look at the quantity that is associated with this price, and based on it and the model of optimal fares, deduce the quantities for other fares from these menus. After subtracting these quantities from the aggregated data, we are back in the original setting when the fares are the same for the remaining travel dates. This intuitive explanation suggests that the aggregated quantity

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	$\$43.63 + \left[\underset{(1.05)}{\$7.11} + \underset{(0.01)}{0.89} \underset{(0.05)}{income_n}\right] dist_n$	$\$320.23 + \left[\underset{(19.35)}{\$27.89} + \underset{(4.95)}{2.54} \underset{(1.54)}{income_n}\right] dist_n$
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)
Marginal cost	c	$4.00 ($12.36)	

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

data provide us with informative moment conditions.

6   Results

6.1   Demand and Cost Estimates

Table 2 presents the two-step GMM estimates of the demand and cost parameters. Based on these estimates and the model of optimal fares, I calculate that 76% of passengers travel for leisure purposes. Business travelers are willing to pay up to six times more for a seat on the average route in my data sample and they are less price sensitive. If fares in all periods 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 0.8%. Business travelers tend to avoid tickets with a cancellation fee as the probability that they have to cancel a ticket is high.

The dynamics of arrival of each traveler type for the estimate of the arrival process α_i is depicted by dotted lines in Figure 4. 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. The bar graph in Figure 4 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. A week before departure, most business travelers start searching for tickets, and the number of active ticket buyers goes up.
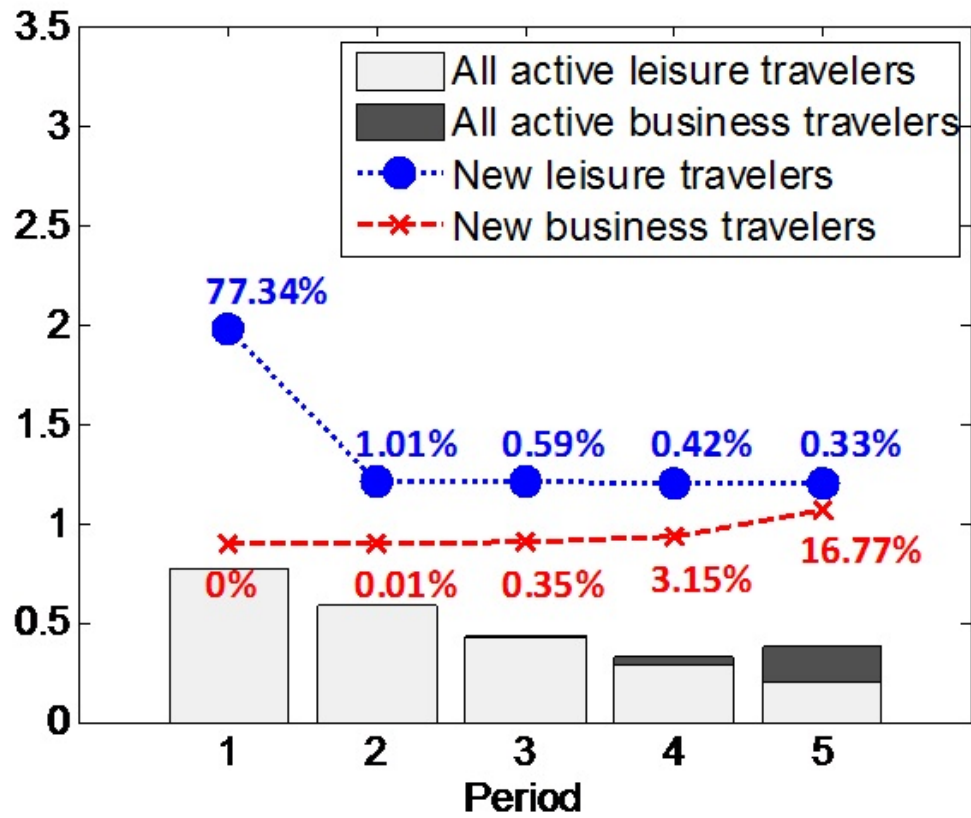
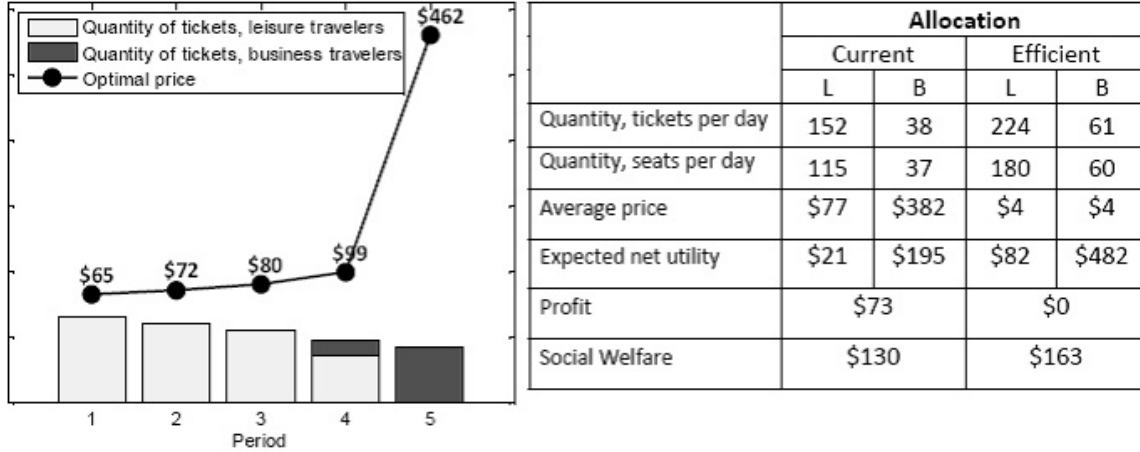Figure 4: Dynamics of active buyers on a route with median income and distance

	Allocation			
	Current		Efficient	
	L	B	L	B
Quantity, tickets per day	152	38	224	61
Quantity, seats per day	115	37	180	60
Average price	$77	$382	$4	$4
Expected net utility	$21	$195	$82	$482
Profit	$73		$0	
Social Welfare	$130		$163	

Figure 5: Optimal price path for a route with median distance and income

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 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 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 5, both groups are buying tickets at the optimal prices in these periods.

6.3 Welfare Estimates

Compared to the efficient supply and allocation of seats, the model's profit-maximizing ticket allocation predicts that travelers and the firm attain 79% of the maximum gains from trade. That the

31

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

gains are below 100% is due market power distortions and misallocations due to price discrimination. Figure 6 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 didn't buy tickets because of high prices. If the allocation was efficient, only travelers who value a ticket more 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 misallocations of seats.

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 inefficiencies 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 considers the case of direct price-discrimination when the airline can perfectly identify customers' types and set prices contingent on them.
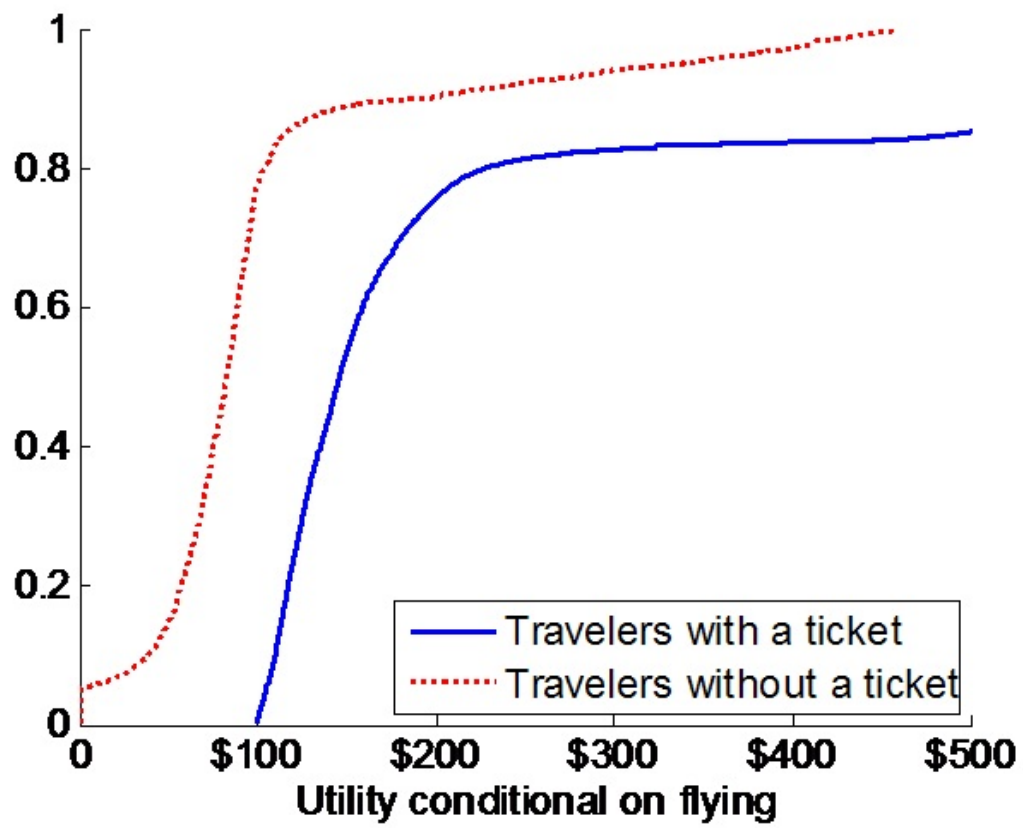
Figure 6: Distributions of travelers' utilities under the optimal allocation of seats

7.1 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. To see that, first, note that for any optimal sequence of prices, the maximum profit of each arbitrageur is zero. Indeed, if an arbitrageur is able to extract some profit then the airline can repeat her actions and increase its profit, which would violate the condition of profit-maximization. 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? Only if it did so without resale. Thus, if the price path without resale is increasing, then the optimal price path in a market with costless resale is flat.

To calculate the optimal fare in this 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 equal to:

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

Since the value of the expected marginal costs is identified only at the profit-maximizing level, we need to make an assumption about its value in the counterfactual scenario. I will make two alternative assumptions. In the first case, I assume that the expected value of the marginal costs is flat. This assumption corresponds to an ideal situation in which the airline is able to adjust its capacity continuously. The value of $\tilde{c}$ will represent the minimum expected value of the average costs, which is the value of the expected marginal costs evaluated at the minimum efficient scale. In the second case, I assume that the graph of the marginal costs is a vertical line, i.e. the airline
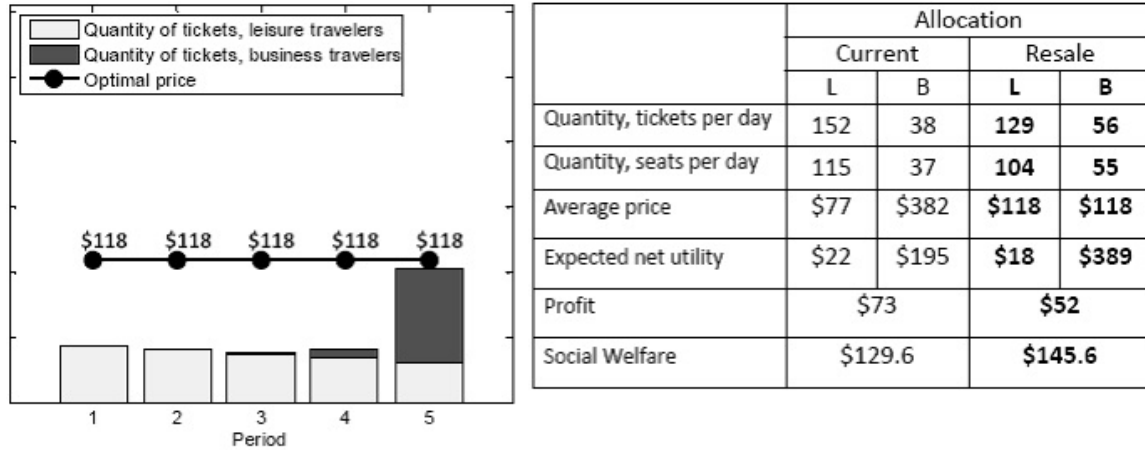
Figure 7: Resale (constant marginal costs)

cannot adjust their capacity.

In both cases, 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, I again use the value of demand parameters that correspond to a route with median characteristics and a non-holiday travel date.

Figure 7 shows the optimal price path for the first case in which the expected marginal costs are fixed. 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 \$318 to \$118. The effect on the business traveler is unambiguous: they pay a lower price and buy a higher quality product. The effect on the leisure travelers is theoretically ambiguous. The price for them increases for two reasons. First, they compete against customers who are willing to pay more. Second, they are willing to pay more for a higher quality product. The estimates suggest that the first effect dominates: their consumer welfare goes down by 20%. 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% and the consumer surplus of business travelers increases by almost 100%. The airline's profit decreases by 28%. Overall, social welfare on the average route increases by 12%. The decrease in the airline's profit may force the airline to exit from the market,

Quantity of tickets, leisure travelers
Quantity of tickets, business travelers
Optimal price

$124 $124 $124 $124 $124

Period 1 2 3 4 5

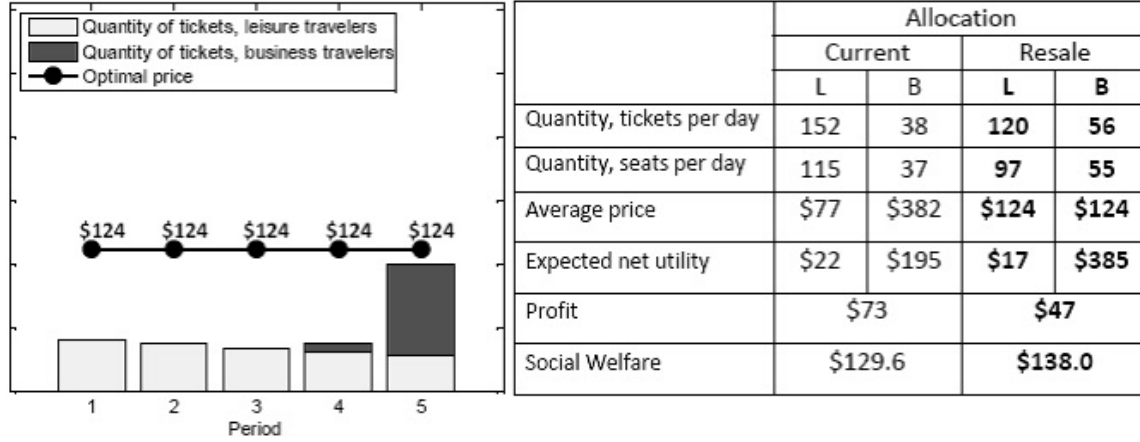| | Allocation | | | |
	Current		Resale	
	L	B	L	B
Quantity, tickets per day	152	38	120	56
Quantity, seats per day	115	37	97	55
Average price	$77	$382	$124	$124
Expected net utility	$22	$195	$17	$385
Profit	$73		$47	
Social Welfare	$129.6		$138.0	

Figure 8: Resale (fixed capacity)

which will decrease the social welfare to zero. Since the fixed costs of the airline are not identified without observing any variation in entry-exit behavior, I cannot evaluate how plausible such an outcome may be.

In the first case, the total number of occupied seats goes up. Therefore, to consider the case in which the airline cannot adjust their capacity, I increased the value of the marginal costs until the number of occupied seats in the counterfactual scenario is equal to its initial level. Figure 8 shows that qualitatively the welfare effects of intertemporal price discrimination remain the same. The average price goes up even more, the median price goes down. The airline's profit decreases even further. The gains for the business travelers outweighs the losses of leisure travelers and the airline. In this counterfactual, the inefficiency in production is fixed since the total quantity remains the same. The increase in the social welfare (+6%) comes from elimination inefficiency in allocation of seats caused by intertemporal price discrimination.

7.2 The role of 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. 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.
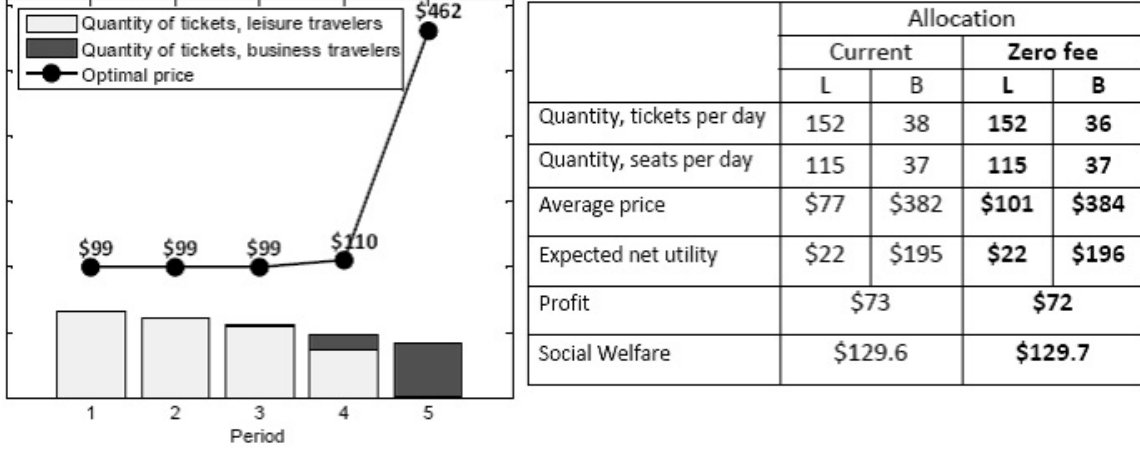
36

	Allocation			
	Current		Zero fee	
	L	B	L	B
Quantity, tickets per day	152	38	152	36
Quantity, seats per day	115	37	115	37
Average price	$77	$382	$101	$384
Expected net utility	$22	$195	$22	$196
Profit	$73		$72	
Social Welfare	$129.6		$129.7	

Figure 9: Zero cancellation fee

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:

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

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.

Figure 9 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 relatively small amount (less than 1%). This result is not too surprising as
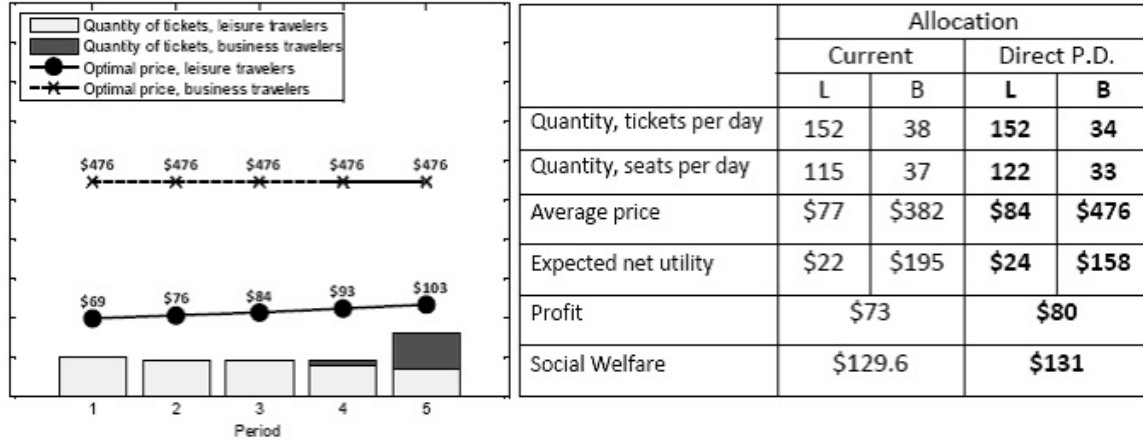
Figure 10: Third degree price discrimination

the airline does not really need to separate business and leisure travelers, as most business travelers are estimated to arrive later than leisure travelers.

This counterfactual assumes that the time when travelers start searching for the ticket is exogenous and therefore does not depend on the value of the cancellation fee. The exogeneity of customers' arrival to the market is the reason why the airline is able to price discriminate. This assumption, however, may not hold in reality. If there is no cost associated with booking tickets early, business travelers might start arriving to the market early and book preemptively. This assumption quickly brings us to the case of costless resale.

7.3 Direct price discrimination

The last counterfactual evaluates the effectiveness of the intertemporal price discrimination strategy. Suppose the airline can recognize a customer type and charge different prices to different customer types. Then there will be two price paths: one for business travelers, another for leisure travelers. The airline will not impose a cancellation fee to separate customers within its type, since there is no within type variation in the value of the cancellation probability. Therefore, in this counterfactual I set the cancellation fee to zero. Figure 10 presents the optimal price paths and the corresponding quantities of sold tickets.

By using intertemporal price discrimination, the airline captures more than 90% of the profit that it could achieve if type-specific prices were possible. Surprisingly, leisure travelers would

prefer to see type-specific prices. There are two reasons for that. First, the airline does not have to damage the product by imposing a cancellation fee. Second, leisure travelers do not compete directly or indirectly with business travelers. As the result, the airline can offer a lower price to leisure travelers, not fearing to lose the price margin on business travelers. Business travelers lose from third-degree price discrimination but their loss is smaller than the total gain of leisure travelers and the airline.

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 are lower prices and higher supply for leisure travelers.

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. The estimated increase in prices is mainly caused by an increase in ticket quality, which does not affect social welfare. Finally, I found that intertemporal price discrimination allows the airlines to achieve more than 90% of the profit that third degree price discrimination would generate.

The study focuses on the set of monopoly markets. There are two potential difficulties with generalizing its results to more competitive markets. First, one may worry about special characteristics of isolated monopoly markets. As the result, the estimated demand parameters may not be representative of the entire industry. Unless the difference between monopoly markets and the rest of the industry is solely caused by the number of potential travelers, this is a valid concern. The second problem is the impact of competition. Dynamic oligopoly models do not generally have

a unique equilibrium prediction. As a result, in may be very difficult to compare equilibria with and without price discrimination. In particular, if resale were allowed, we will have to consider an equilibrium in which a travel agency buys all tickets from the competing airlines at the beginning of sale and then acts as a monopoly in the secondary market. Whether this outcome is plausible is a question for future research.

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.

EVANS, W. N., AND I. N. KESSIDES (1994): "Living by the "Golden Rule": Multimarket Contact in the U.S. Airline Industry," *The Quarterly Journal of Economics*, 109(2), pp. 341–366.

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.

SCHWARTZ, M. (1990): "Third-Degree Price Discrimination and Output: Generalizing a Welfare Result," *The American Economic Review*, 80(5), 1259–1262.

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.