Service Operations in the Presence of Strategic Consumer Behavior

by

Katherine Ashley

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Pnina Feldman, Chair
Professor Candace Yano
Professor Zuo-Jun Shen

Fall 2016

Service Operations in the Presence of Strategic Consumer Behavior

Abstract

Service Operations in the Presence of Strategic Consumer Behavior

by

Katherine Ashley

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Pnina Feldman, Chair

This dissertation analyzes pricing and information provision in the presence of strategic consumer behavior, with a focus on service industries. Using both analytical and empirical methods, I analyze two tools that firms may use to match supply with demand when customers are strategic: reservation no-show fees and inventory announcements.

First, I study a common problem in the service sector: reservation no-shows. Many firms allow customers to make reservations for service at a specified time in the future. Reservations are valuable to consumers because they insure against the service not being available, but are costly for the firm because valuable capacity may be held for customers who fail to show up. No-show fees, payable only if a reservation-holder fails to show up in the service period, are one tool that firms can use to mitigate the loss from wasted capacity. I analyze the performance of a no-show fee pricing scheme for a capacity-constrained service firm in a market with both forward-looking customers and a walk-in demand segment. I characterize three types of optimal policies that emerge in equilibrium and examine the social welfare implications of charging an optimal fee. In equilibrium, social welfare may strictly increase under a no-show fee policy; that is, the utility gain from using a no-show fee to allocate capacity more efficiently may be greater than the cost it imposes on consumers.

In the second half of the dissertation, I analyze the information content and sales impact of inventory announcements, i.e. messages such as "3 seats left at this price." I use an expansive original dataset to study the impact of these announcements in the market for airline tickets. My data include itinerary listing and seat map information for thousands of flights on three major U.S. carriers, with travel dates in the winter and spring of 2015. I measure the "informativeness" of announcements, which I define as their ability to improve customer predictions of future price changes, and estimate the impact that these messages have on ticket sales. I find that announcements contain significant information about future price changes (both direction and magnitude), and are most informative early in the booking horizon. To measure the customer response, I use an instrumental variables model to address potential endogeneity between sales and announcements. My findings provide evidence that

for at least some airlines, consumer ticket purchases increase immediately after new inventory messages, a result that is consistent with rational consumer behavior.

For Mom, Dad, and Liza

Contents

List of Figures

List of Tables

Acknowledgments

First and foremost, I extend heartfelt thanks to my advisor, Professor Pnina Feldman, for her guidance, patience, and support during my work on this dissertation. Pnina's approach to identifying interesting and important research questions will serve as a model for me throughout my academic career, and I feel privileged to have learned from her during my doctoral studies.

I am grateful to the other members of my committee, Candace Yano and Z. Max Shen, for their input and guidance on this dissertation. To the many other Berkeley faculty who have shaped my development as a researcher, thank you for your generosity in helping the next generation of academicians.

My journey into academia began at Northwestern, and I am grateful to faculty members including Ian Savage, Janice Eberly, and Robert Coen for their advisorship during my undergraduate studies and the graduate school admissions process.

To my Ph.D. colleagues, past and present: it is no exaggeration to say that I wouldn't have gotten through this without you. Sanaz, the charm to my wit, thank you for being an incredible officemate, friend, and human. I am so excited to see you take on the illustrious academic career that lies ahead of you. Eliot, thank you for limitless emotional support over the years and for knowing even more 30 Rock quotes than I do. Moshe, thank you for being a loyal and generous friend and an excellent roommate. Liz, thank you for being the perfect partner in spontaneous travel adventures. Jenn, thank you for those much-needed yoga and wine (and cheese) girls' nights. Mark, thank you for introducing me to business academia, and for a decade of friendship and good conversation.

To my students at Haas, thank you for inspiring me and giving meaning to what I do.

To the Strawberry Canyon Track Club, especially Coach Carl Rose, thank you for helping me discover a lifelong passion outside of academia. To Kari, thank you for hundreds of miles on the Fire Trail.

To everyone else who has been a part of my experience here, thank you for sharing this challenging and transformative time in my life.

To Haley, Becca, and Hannah, thank you for your lifelong friendship no matter how many miles are between us. To Meaghan, I (still) owe you a half Ironman! Thank you for being a wonderful and understanding friend.

To my parents and my extended family, thank you for loving me unconditionally and for raising me to be hard-working, independent, and intellectually curious. To Bud, I miss you.

To Liza, thank you for being my very first student (not that you had any say in the matter) and the best little sister in the world. I am so proud of the brilliant and compassionate young woman you've become.

And finally, to Owen. Thank you for all of your love and support through the final stretch of preparing this document. I can't wait to start the post-dissertation chapter of life with you.

Chapter 1

Introduction

Briefly stated, operations management is the science of matching supply with demand. Economic theory predicts that prices will naturally rise when demand outpaces supply and fall when supply exceeds demand, but in practice, managers often need to intervene in the process to avoid losing value and wasting resources. Even with macroeconomic forces working to align supply and demand, costly mismatches can exist in the short term, and may be exacerbated and perpetuated by constraints on capacity or price levels. Furthermore, some of the mechanisms that emerge to correct imbalances, such as secondary markets, transfer value to agents other than the firm that provides a good or service. For these reasons, effectively matching supply with demand is a task that is both difficult and crucially important for firms to master (Cachon and Terwiesch 2011).

In much of the existing operations management literature, customer arrivals are modeled as exogenous processes, with each individual making a simple decision to buy if his valuation exceeds the price and leave the market otherwise. While these assumptions are appropriate in some settings, in other contexts they do not accurately capture reality. Consumers today have access to an overwhelming amount of information. Before making a purchase, they can read reviews from other customers, research alternative options, check for a lower price from another seller, and perhaps even engage with the firm through social media channels. In settings where such behavior is commonplace, it no longer makes sense to assume that customers arrive according to an exogenous process, or that their purchase decisions are determined by a simple comparison of price and instantaneous utility. There are myriad opportunities for customers to be strategic in their interactions with firms, and previous research in this area has shown that failing to account for strategic consumer behavior can be disastrous for revenues (Shen and Su 2007).

In this dissertation, I analyze two different tools that firms may use to match supply and demand when customers behave strategically: no-show fees and inventory announcements. I frame my studies of operational decision-making in the service sector, an important and growing segment of the world economy, and use both analytical and empirical techniques to develop theory and validate theoretical predictions with data.

First, I analyze a monopolist service provider that may choose to offer reservations for service. Reservations play an important role in many service industries: when capacity is limited and demand is uncertain, reservations are valuable to customers because they provide insurance against the risk faced by customers who walk in during the service period: a reservation eliminates or decreases the risk that a customer who desires service will be unable to obtain it. Meanwhile, offering reservations can be costly for the firm: if customers reserve capacity but fail to show up, the firm forfeits revenue that could have been earned by selling that capacity to other customers.

I analyze the choice of a reservation no-show fee by a monopolist service provider facing customers who act strategically, both in making reservations and in deciding whether to keep them. I characterize the fee policy that emerges in equilibrium and show that despite consumers' aversion to fee-based reservation policies, both customers and firms can benefit when the firm uses a no-show fee pricing scheme. The no-show fee policy results in a more efficient allocation of scarce capacity, and serves as a partial antidote to strategic behavior.

Next, I present an empirical study of inventory announcements–messages such as "3 seats left at this price"–in the market for airline tickets. Although inventory announcements have become increasingly prevalent in both retail and service settings, there is little empirical evidence of their informational value to customers, nor of the impact they have on sales. Messages of this type suggest scarcity, but it is not clear how much information they contain, as their content is not directly verifiable by consumers.

I use an expansive original dataset with data from three different airlines to answer two central questions related to airline inventory announcements. First, how informative are airline inventory announcements about future price changes? I define "informativeness" as the ability of inventory announcements to improve customer predictions of future price changes: both direction and magnitude. Using a multinomial logit regression model, I find that itineraries displaying announcements are much more likely to change price than are comparable itineraries without announcements. Both price increases and price decreases are more likely after an inventory announcement, and although the relative probability of increase is higher after an announcement in most cases, there are times in the booking horizon (close to departure) when an announcement can signal a higher likelihood of price decrease. Inventory announcements also predict price changes that are larger in magnitude, compared to price changes that occur when no announcement is shown.

My results imply that the implications of an announcement will depend on customer preferences. For a risk-averse or risk-neutral customer who is certain she will travel, buying now becomes more attractive relative to waiting if an inventory message is shown. However, for a customer who has other options, or wants to travel only if he can find a good deal, an announcement could signal that it is worth waiting to see if the price drops–at significant risk of a price increase, of course. Furthermore, we find that the informativeness of inventory messages varies over the booking horizon, with the highest information content further from the travel date.

The final question I analyze is how announcements affect customer purchase decisions. To correct for bias resulting from the potential endogeneity of announcements and sales, I use an instrumental variables two-stage least squares (2SLS) model. I find a positive immediate sales response to announcements for two out of the three carriers in my study, consistent with rational consumer behavior given the information content of announcements. This result is robust to the inclusion of various market- and itinerary-level explanatory variables.

My dissertation research analyzes two methods that service-sector firms use to influence the behavior and purchase decisions of strategic customers: charging no-show fees and making announcements about inventory levels. My analysis provides managerial insights for firms that may consider using these tools in their practice of revenue management.

Chapter 2

Reservation No-Show Fees

2.1 Introduction

Reservations play an important role in many service industries. Restaurants, healthcare providers, and fitness studios are just a few examples of firms that may offer reservations for service at a fixed time in the future. When capacity is limited and demand is uncertain, reservations are valuable to customers because they provide insurance against the risk faced by customers who walk in during the service period: a reservation eliminates or decreases the risk that a customer who desires service will be unable to obtain it. Meanwhile, offering reservations can be costly for the firm. While there are some operational benefits associated with offering reservations, there is also a significant cost: if customers reserve capacity but fail to show up, the firm loses out on revenue that could have been earned by selling that capacity to other customers. This is because reservations confer a *right* to purchase service, but not an *obligation*: like the owner of a financial call option, a reservation holder may choose not to exercise his right to buy. For a traditional free reservation, there is no penalty for deciding not to show up in the service period. This is especially problematic because in many service industries, many costs for a given service period are incurred before no-shows take place. For example, restaurants purchase perishable ingredients and set staffing levels ahead of time, so the revenue loss from a no-show party is not offset by a comparable reduction in costs. As a result, the no-show problem is a serious concern for firms that offer reservations.

The fact that reservations are valuable to customers and costly for firms suggests that a possible solution lies in charging for reservations. However, attempts to implement reservation pricing schemes have a spotty history in hospitality industries. A recent New York Times article reports that while reservation charges have become more common in recent years, they are still met with resistance from many customers. Even firms are often not enthusiastic about charging no-show fees, which can make the reservation feel more like an obligation than an enjoyable experience (Wells 2015). Some service firms have implemented reservation fee policies only to modify or abolish them soon after. Volver, a high-end tapas restaurant in Philadelphia, used a ticketing policy (collecting the entire price of its prix fixe menu in a nonrefundable payment at the time of reservation) when it first opened in 2014, but has since scaled back to a more modest no-show fee with free cancellation up to 48 hours in advance (Szkaradnik 2014).

In this paper, we analyze the optimal no-show fee policy for a firm that serves both a population of consumers who consider reserving in advance and a walk-in demand segment. We identify two distinct mechanisms through which no-show fees can benefit the firm: behavioral influence and revenue generation. We characterize the market conditions under which various types of equilibrium emerge and quantify the impact on customer welfare in each case. We show that no-show fees can strictly increase consumer surplus when they produce a better allocation of scarce capacity than would occur without fees, reducing the deadweight loss that results from wasted capacity saved for no-shows.

The rest of the paper is structured as follows. Section 2.2 reviews relevant literature. Section 2.3 describes the model fundamentals in greater detail. Section 2.4 analyzes the free reservations policy that we use as a baseline for comparison. Section 2.5 analyzes the primary no-show fee model, describes possible equilibrium outcomes, and characterizes the impact of an optimally designed no-show fee policy on firm revenues and consumer surplus. Section 2.6 concludes with a summary of our results and contributions.

2.2 Related Literature

Our paper relates to several branches of existing research. Broadly speaking, our research fits into the revenue management literature, of which Talluri and van Ryzin (2004) provide a comprehensive review. More recently, Guillet and Mohammed (2015) review revenue management research in the hospitality and tourism industries. Within revenue management, there are a number of papers that directly address aspects of the reservations problem. Png (1989) studies a joint pricing and capacity choice problem for a monopolist firm selling to customers who are initially uncertain about their valuations for service. The optimal pricing policy is a reservation scheme which specifies payment amounts that differ with respect to when the customer uses the reservation and is able to obtain a unit of capacity, uses the reservation and is unable to obtain a unit of capacity, and does not use the reservation. In a key difference from our paper, the optimal policy in Png's setting is to charge nothing when a customer does not keep the reservation. Çil and Lariviere (2013) consider reservations in a market with two segments: one segment makes reservations in advance (and will not patronize the firm without one), the other walks in at the time of service. They find that contrary to the result from Littlewood (1972), it may not be optimal to save capacity for late-arriving customers even when that segment is more valuable. In our paper, advance consumers may visit the firm with or without a reservation, and we consider the use of no-show fees to influence their reservation and purchase decisions.

In Georgiadis and Tang (2014), customers are differentiated along two dimentions, valuation and probability of showing up. The optimal no-show fee reservation policy may target some of the segments and price out others. Our model also features consumer heterogeneity, but only in the second period after reservation decisions have been made. Osadchiy and Vulcano (2010) model a firm that sells a good on a spot market at full price, and also accepts binding reservations to purchase the item at a discount at the end of the selling season (if there are units available). Their model is based on a continuous time setting in which customers know their valuations ahead of time. The authors show that an equilibrium exists and that the seller can increase revenues by using a binding reservation scheme. The reservations with no-show fees in our paper have some similarities to binding reservations, in that they may cause customers to purchase service when they would not have otherwise. However, valuation uncertainty in the advance period is integral to our analysis because it drives the endogenous no-show behavior that is central to our model.

Within the literature on reservations, several papers specifically address hospitality industries that fit our setting. Kimes et al. (1998) apply revenue management principles to the restaurant industry and present a unified framework of strategies that restaurants can use to manage demand. Reservations are discussed as one method of reducing arrival uncertainty, though they do not completely eliminate the uncertainty due to customer no-shows and late arrivals. The authors mention deposits and credit card-guaranteed reservations as ways to further reduce arrival uncertainty, but do not analyze these fees. Bertsimas and Shioda (2003) study restaurant revenue management in a queueing framework, in which arriving parties may be of different sizes and have different service times. In their paper, the firm optimizes the number of reservations to accept and the seating decisions for walk-in parties; they do not study the firm's pricing decisions or model strategic customer behavior. Oh and Su (2012) also model restaurant reservations in a queueing framework, in which the benefit of holding a reservation is a reduction in waiting time for service. Their results indicate that the firm should always charge a no-show fee equal to the service price, while also offering a discount to reservation customers compared to walk-ins. By contrast, we consider a setting with limited capacity and find that a range of equilibrium no-show fees are possible.

Alexandrov and Lariviere (2012) also study the reservations problem in a restaurant setting, and characterize the conditions under which offering free reservations is more profitable for the firm than not offering reservations. In their model, offering reservations stimulates demand when the realized market size is small, but hurts the firm when market size is large because some capacity is allocated to customers who fail to keep their reservations. This occurs because customers face positive travel and denial costs, and having a reservation makes the trip to the restaurant worth it by guaranteeing that capacity will be available. In our model, reservations themselves do not have a demand stimulating effect because customers do not face travel or denial costs. Instead, we focus on the behavioral changes incentivized by no-show fees, which can reduce wasted capacity by both stimulating demand and encouraging early cancellation instead of simply not showing up.

Our paper also relates to the broad literature on advance selling. Two key features of advance selling models are (1) purchase and consumption of a good or service occur at distinct points in time, and (2) customers are more homogeneous in advance than they are in the period when the delivery of goods or services takes place. In particular, a seller can get customers to make a payment in advance, while they are still uncertain about their valuations for the good or service being sold. Gale and Holmes (1993) show that a monopolist seller can use advance purchase discounts to segment a market in which some consumers are more certain about their future valuation for a service than are others. DeGraba (1995) shows how a monopolist seller can induce consumers to purchase early, while they are still uninformed about their valuations, by making fewer units available than there are customers in the market. Our paper also considers a monopolist with limited capacity, but our primary model focuses on the firm's use of fees to allocate scarce capacity between two customer populations, rather than market segmentation or capacity choice. Xie and Shugan (2001) and Shugan and Xie (2000) study the marketing implications of advance selling in a service context; they show that when customers are uncertain about their future valuations, a monopolist service provider can do substantially better (under a wide variety of market conditions) by selling services in the advance period. In our model, a form of advance selling may arise in equilibrium (in the sense that the firm may get advance customers to commit to a no-show fee 'deposit'), but under some market conditions the firm chooses not to sell in advance. Cachon and Feldman (2011) compare offering subscriptions and charging on a per use basis in a service system. Viewing subscriptions as a form of advance selling, they find that the firm may benefit from selling subscriptions despite their limited ability in controlling congestion. Their model applies to repeated-purchase contexts, rather than the single-purchase setting that we study in our paper. Cachon and Feldman (2016) study advance selling under two forms of competition. They show that when firms compete, advance selling always arises in equilibrium but lowers profits relative to selling only in the spot period. However, when a firm competes against customers who purchase in advance and then resell, the benefit of advance selling is restored; resale can be beneficial for competitive firms but not for monopolists. Our analysis provides some of the same intuition as these papers: the firm can take advantage of the fact that customers are homogeneous in the advance period to get them to commit to making a payment before they learn their valuations. However, the consideration of no-show fee policies and the fact that payment may occur in both periods make our paper distinct from the pure advance selling literature.

Finally, in our model, reservations give customers the opportunity (but not the obligation) to purchase the firm's service at a specified price. In this way they are equivalent to real options on the firm's capacity. In the operations literature, Gallego and Sahin (2010) study the use of real options as a revenue management tool for a firm that faces strategic customers whose valuations become more certain over time. The authors show that selling options on capacity (equivalent to charging a non-refundable deposit, assuming the time value of money can be ignored) allows the firm to do strictly better than charging a lower fare in the advance period and a higher fare in the spot period (a practice frequently used by airlines to price discriminate between leisure and business travelers). In fact, selling real options is socially optimal and efficient. In our paper, the no-show fee scheme is equivalent to choosing how much of the service price should be taken as a nonrefundable deposit; it does not allow the firm to adjust the total amount paid by customers who purchase service. Our paper contributes to the existing literature by focusing directly on the problem of mitigating no-show behavior in a setting with capacity constraints, stochastic market size, valuation uncertainty, and strategic consumer behavior. We identify two distinct mechanisms through which no-show fees create value for the firm–behavioral influence and revenue generation–and characterize when each one is in effect. We also analyze the impact of no-show fees on customer and social welfare and show that no-show fees are not always bad for the consumer. This is a useful finding for service firms that face backlash against no-show fee policies; in some cases, simply communicating the broader implications of no-show fees can help to allay customer concerns.

2.3 Model Description

A single firm with fixed capacity κ offers a service to consumers. There are two periods: an advance period, in which customers can make reservations for future service, and a service period in which customers purchase and consume the service. The service has a price p, which is exogenous in our model. Under some reservation policies that we consider, the firm may also charge a no-show fee. When the firm charges a no-show fee, customers who reserve agree to pay a fixed amount η if they do not show up to purchase service in the second period. A customer who reseserves and keeps his reservation (i.e., visits the firm and pays p for service) does not have to pay the no-show fee. The no-show fee is fixed before reservations take place and cannot be changed.

There are two distinct customer populations, which arrive at different points in time. At the beginning of the first period, $A < \kappa$ customers arrive to the market; we refer to this segment as the *advance* population. Advance customers are strategic (forward-looking), so although they are present in time to make reservations, they may decide to wait until the service period if doing so offers higher expected utility than making a reservation. We assume that the firm knows the size of this population ex ante, e.g. from market research and/or data from its website and reservation systems.[1] At the beginning of the second period, an uncertain number of new customers arrive to the market; we refer to these consumers as the *walk-in* population. The number of walk-ins is a random variable X with cumulative distribution function $F(\cdot)$ on $[0, \infty)$. These consumers may not have been aware of the firm in advance, but they are nearby during the second period and desire service. The size of the walk-in population is influenced by short-term factors that are uncertain until the day of service, such as weather, traffic conditions, and unpredictable events in the surrounding area. In both segments, each consumer who visits the firm demands exactly one unit of capacity. We assume that both segments have the same outside option value, which is normalized to zero.

[1] In some service settings, a customer may need to create an account online before making a reservation; the firm can infer the advance population size from the number of registered users in its systems. Even if the firm does not know the size of the advance market perfectly, the number of customers in this population is not influenced by idiosyncratic events on the day of service (because reservations happen in advance), making it easier to predict than the number of walk-ins.

Advance customers are uncertain about how much they will value the service in the second period. Specifically, each advance customer knows that he will have value V for service, where V can take on one of two values, v_l or v_h, where $v_h > p > v_l$.[2] For v_l and v_h satisfying these criteria, we can think of the advance population as 'potential customers' who may or may not want the service at the time it takes place, depending on whether their realized valuations are above or below the service price. A known fraction $\alpha \in [0, 1]$ of the advance population will end up being high types, and each customer attaches a prior probability α to being a high type, and $1 - \alpha$ to being a low type. The advance customer population is homogeneous in advance, but heterogeneous in the second period once valuations are realized. By contrast, the walk-in customers who arrive in the second period all have high valuation v_h for service. However, the firm does not know in advance how many of these customers will be in the market: capacity allocated to reservation-holders may go unused if some customers are no-shows, but saving capacity for walk-ins is risky because the realized population size could be small.[3]

In our primary model, capacity that is held for reservation-holders who don't show up cannot be reoffered to walk-ins; the no-shows become known too late for the firm to offer the capacity to other customers. Our main results are qualitatively unchanged under partial reoffering, i.e., some fraction of unused reservation capacity can be made available to walk-in customers. The higher the reoffering percentage, the less costly no-shows are for the firm; in the limit, with perfect reoffering, no-shows have no impact on firm revenues at all. The attention that no-shows have received from firms in recent years indicates that wasted reservation capacity is an issue, so an assumption of imperfect reoffering is appropriate for most service settings. The no-reoffering case isolates the impact of customer no-shows, independent of firm reoffering policies, and focuses attention on the underlying mechanics of this setting.

The timeline of firm and customer actions in our primary model is as follows. First, the firm announces the pricing and reservation policy, $\{p, \eta, k\}$, where p is the service price, η is the no-show fee, and k is the quantity of reservations offered. For our main analysis in Sections 2.4 and 2.5, we consider the case where all capacity is offered through reservations (i.e. $k = \kappa$). We assume that the firm knows the consumer valuation and population size parameters when setting the reservation fee. After observing prices and the firm's reservation policy, advance customers decide whether to request a reservation or wait until the service period. The full timeline of events is illustrated in Figure 2.1.

[2] Although we do not impose an explicit lower bound on the low valuation here, for advance customers to find it worthwhile to be in the market, a reasonable assumption is that their expected valuation for service is at least the service price, i.e. $\alpha v_h + (1-\alpha)v_l \geq p$. This corresponds to a lower bound on the low valuation of $v_l \geq \frac{p - \alpha v_h}{(1-\alpha)}$. However, this assumption is not required for the proceeding analysis.

[3] Equivalently, one could think of an overall walk-in population that includes both high and low types, but low types choose to stay home in the service period and therefore are not visible to the firm. This would simply scale the entire distribution of market size by the fraction of high-types in the walk-in population. Since low-type walk-ins would never visit the firm under any pricing and reservation policy, we do not explicitly include such a segment in our model.

Customers face a tradeoff because there may be a cost to making a reservation, but a customer who waits may be unable to get service. In order to evaluate the expected surplus from waiting until the service period, customers form expectations of the probability of getting service as a walk-in, $\hat{\phi}$. Following Tereyağoğlu and Veeraraghavan (2012) and Su and Zhang (2008), the advance customer expects to be served if and only if there are fewer customers requesting service than there are units of capacity. For each model, we characterize the subgame perfect Nash equilibria in pure strategies of the game between the firm and consumers. An equilibrium consists of the actions taken by the forward-looking firm and consumers, given their expectations of the actions chosen by the other agents. The beliefs of all agents are consistent with the equilbrium outcome, so that no player has a profitable deviation from his chosen strategy.

2.4 Free Reservations

As a baseline for comparison, we first analyze the case in which the firm offers free reservations. There is no decision for the firm in this case. With no fees, a reservation holder will use the reservation if and only if he ends up being a high-type, since it gives him net utility $v_h - p > 0$ in that case. For notational convenience, define $v := v_h - p$ as the net value gained by a high type who purchases service. If a reservation-holder ends up being a low type, his value for keeping the reservation is $v_l - p < 0$, so he prefers to become a no-show, getting value 0, rather than using the reservation and getting a negative surplus. Given that low types will become no-shows, the expected utility from reserving is αv. A customer who does not reserve will attempt to walk in if and only if he is a high type, so the expected utility from not reserving is $\hat{\phi}\alpha v$, where $\hat{\phi}$ is expected fill rate, i.e. the expected probability of getting a unit of service as a walk-in. Because the maximum possible fill rate is 1, all customers choose to reserve when reservations are free.

In the first period, the firm gives out A reservations to advance customers. In the second period, there are $\kappa - A$ seats remaining for walk-in customers. The expected number of walk-in sales is $\int_0^{\kappa-A} x dF(x) + (\kappa - A)\overline{F}(\kappa - A)$. Letting $S_D(z)$ denote the general sales function expression $E_D[\min\{D, z\}]$ for the minimum of a random variable D and some $z \geq 0$, the number of walk-in sales can be written more compactly as $S_X(\kappa - A)$. In addition to the walk-in sales, there are αA high-type reservation holders who visit the firm and pay p for service. The remaining $(1-\alpha)A$ reservation holders who have low realized valuations become no-shows, and do not make any payment to the firm. The firm earns expected revenue:

$$\Pi_0 = (\alpha A + S_X(\kappa - A))p \tag{2.1}$$

Total customer surplus is:

$$CS_0 = (\alpha A + S_X(\kappa - A))\, v \tag{2.2}$$

We will use customer and firm welfare levels under free reservations as a basis for comparison in the next section.

2.5 No-show Fee

We now analyze the firm's choice of a no-show fee when service price is fixed and all capacity is offered through reservations, i.e. $k = \kappa$. Note that because consumers are homogeneous in the reservation period, in a pure strategy equilibrium, all advance customers make the same decision about whether to request reservations.

Consumer Behavior

When the firm charges a strictly positive no-show fee η, a customer who reserves and ends up having a low valuation for service faces a different decision than in the free reservations case. He can keep the reservation and get net utility $v_l - p < 0$, or become a no-show, pay the fee, and get utility $-\eta$. Both options yield negative net utility; the best he can do is choose the action that leads to the less negative outcome. To simplify notation, define $c := p - v_l$ as the cost to a low-type customer of keeping the reservation. Low-type customers will keep their reservations if $\eta \geq c$; otherwise, they will become no-shows and pay the associated fee η. In effect, customers commit to pay the firm at least η when making a reservation. Note that a no-show fee strictly greater than the low type's cost c has exactly the same effect as a fee equal to this cost: if $\eta > c$, then all reservation-holding customers will visit the firm (regardless of realized valuation v) to avoid paying the no-show fee. Therefore, without loss of generality, we restrict $\eta \leq c$. A customer would never pay the firm a no-show fee higher than c, and a fee higher than c would have no incremental impact on customer no-show behavior.

For any no-show fee $\eta \in [0, c]$, an advance customer's expected utility from reserving is $\alpha v - (1-\alpha)\eta$. With probability α, the advance customer has a high valuation, visits the firm to purchase service, and gets a surplus of v. With probability $1-\alpha$, the advance customer has a low realized valuation for service, and he either pays η as a no-show, or, if $\eta = c$, he gets utility of $-c = -\eta$ by purchasing service at a price higher than his valuation. Meanwhile, expected utility from waiting to walk in during the second period is $\hat{\phi}\alpha v$, where $\hat{\phi}$ is the customer's expectation of the probability of obtaining a unit of capacity if he waits. The scalar α reflects the fact that if an advance customer waits, he will only attempt to walk in if his valuation for service is high.

Because the advance population is homogeneous in the first period, all advance customers make the same buy-or-wait decision in equilibrium. Thus, there are two possible equilibrium strategies for the advance population: everyone reserves, or everyone waits. We will refer to these two categories of optimal advance consumer behavior as "reservations equilibria" and "waiting equilibria," respectively.

If the advance population reserves, then there are $\kappa - A$ units of unreserved capacity in the service period, and an advance customer who deviates by waiting will expect to get a unit if and only if realized walk-in demand is below $\kappa - A$. Therefore, the expected fill rate for advance customers in a reservations equilibrium is $\hat{\phi}_r = F(\kappa - A)$. Meanwhile, if the advance population does not reserve, then all κ units of capacity are still available in the second period. There will be αA high-type advance consumers present and desiring service in the second period, in addition to walk-in demand, so the expected fill rate for an advance customer who waits is the probability that there are fewer than $\kappa - \alpha A$ walk-ins, or $\hat{\phi}_w = F(\kappa - \alpha A)$. As expected, the fill rate is lower when the advance population reserves, i.e. $\hat{\phi}_r < \hat{\phi}_w$.[4]

Lemma 1. *For any possible no-show fee η, there is an equilibrium customer response. There exists a threshold η_1 such that for any no-show fee $\eta \leq \eta_1$, the unique equilibrium is for advance customers to reserve.*

Proof. From the consumer indifference equation, all advance customers prefer to reserve if $\alpha v - (1-\alpha)\eta \geq \hat{\phi}_r \alpha v$. Defining $\eta_2 = \frac{\alpha(1-\hat{\phi}_r)v}{1-\alpha}$, advance customers reserving is an equilibrium if $\eta \leq \eta_2$. Similarly, advance customers prefer to wait if $\alpha v - (1-\alpha)\eta < \hat{\phi}_w \alpha v$. Defining $\eta_1 = \frac{\alpha(1-\hat{\phi}_w)v}{1-\alpha}$, advance customers waiting is an equilibrium if $\eta > \eta_1$. Note that $\hat{\phi}_w > \hat{\phi}_r$, therefore, $\eta_1 < \eta_2$. Every possible fee η satisfies at least one of the equilibrium conditions.

For $\eta \leq \eta_1$, the unique equilibrium is for all advance customers to reserve. For $\eta \in (\eta_1, \eta_2]$, both equilibrium conditions are satisfied, so either an all-wait or an all-reserve equilibrium is possible. For $\eta > \eta_2$, all advance customers waiting is the unique equilibrium customer response. □

Lemma 1 shows that there is a threshold no-show fee, η_1, up to which the unique equilibrium is for all customers to reserve. This threshold fee is a function of market parameters: the fraction of high valuation customers in the market; the amount of surplus that high types gain from purchasing service; the size of the advance population relative to capacity; and the distribution of walk-in demand.

Corollary 1. *If $\eta_1 \geq c$, then advance customers always reserve in equilibrium.*

An immediate corollary to Lemma 1 is that if the threshold fee η_1 in a given market is weakly higher than the loss to low-type consumers from purchasing service, then advance customers will reserve under any possible fee that the firm can charge.[5] Lemma 1 also shows that there may be a non-empty range of fees for which there are two different customer best responses, i.e. fees for which both all-reserve and all-wait are equilibria for the advance population.

Corollary 2. *For the range of fees for which the consumer equilibrium is not unique, the all-wait equilibrium makes advance customers better off than the all-reserve equilibrium.*

[4] In pure strategy equilibrium, all advance customers make the same buy-or-reserve decision, so the equilibrium fill rate will be either $\hat{\phi}_r$ or $\hat{\phi}_w$. However, the fill rate can be defined similarly off the equilibrium path: if a fraction γ of the advance population waits, then the expected fill rate will be $\hat{\phi}_\gamma = F(\kappa - \alpha A - (1-\alpha)\gamma A)$, and $\hat{\phi}_r < \hat{\phi}_\gamma < \hat{\phi}_w$ for $\gamma \in (0, 1)$.

[5] This is not dependent on the assumption $\eta \leq c$, since any fee strictly above c is exactly the same as $\eta = c$ from an advance customer's perspective.

Proof. From the definition of η_1, an advance customer's expected utility in an all-reserve equilibrium with a fee greater than η_1 is strictly less than $\alpha v - (1-\alpha)\eta_1 = \alpha v - (1-\alpha)\left(\frac{\alpha(1-\hat{\phi}_w)v}{1-\alpha}\right) = \hat{\phi}_w \alpha v$, which is the expected utility of an advance customer in an all-wait equilibrium. □

For fees in the range $(\eta_1, \eta_2]$, either a reservations equilibrium or a waiting equilibrium is possible, but advance customers are strictly better off under the waiting equilibrium. In the proceeding analysis, we will focus on the conditions under which all-reserve is the unique consumer equilibrium, and assume that advance consumers play the higher-utility equilibrium strategy in the case where both types of equilibria exist simultaneously.

In the next section, we finish the characterization of equilibrium by solving for the optimal no-show fee in the firm's profit maximization problem.

Firm's Profit and Consumers' Surplus

The firm's profit function depends on customers' reservation request and service request decisions. In the previous section, we showed that the set of possible no-show fees can be partitioned into those that will induce a reservations equilibrium and those that will induce a waiting equilibrium. Corollary 1 shows that the second set may be empty under some market conditions. To determine the optimal no-show fee, we first find the best fee in the range $[0, \eta_1]$ (i.e., the best fee such that advance customers reserve), which we denote by η_r^*. We then compare the firm's profit with a fee of η_r^* to the profit when all advance customers wait, if waiting is a possible equilibrium response, to determine the firm's equilibrium strategy. We use η^* to denote the equilibrium no-show fee across the full range of possible fees that make customers either wait or reserve.

In a reservations equilibrium, the firm earns revenue p from each customer who purchases service (either with or without a reservation), and η from any advance customers who make reservations but fail to show up. There are two cases for the firm's revenue function when advance customers reserve. In a reservations equilibrium where $\eta < c$, only high-type reservation holders show up to purchase service in the second period; low types become no-shows and pay the fee. In a reservations equilibrium where $\eta = c$, all advance customers show up and pay the service price. Expected profit as a function of the no-show fee is:

$$\Pi_r(\eta) = \begin{cases} (\alpha A + S_X(\kappa - A))\, p + (1-\alpha)A\eta & \eta < c \\ (A + S_X(\kappa - A))\, p & \eta = c \end{cases} \tag{2.3}$$

Over the full range of possible fees $\eta \in [0, c]$, customer surplus can be written as:

$$CS(\eta) = (\alpha A + S_X(\kappa - A))\, v - (1-\alpha)A\eta \tag{2.4}$$

Lemma 2. *If $v_l \geq 0$, then the best possible reservations equilibrium for the firm is to charge the highest fee that will make all customers reserve. That is, $\eta_r^* = \min\{\eta_1, c\}$.*

Proof. Firm profit is increasing in η for all possible no-show fees $\eta \in [0, c]$.

Case 1. $\eta < c$. *For no-show fees in this range, the firm earns the fee $\eta < c$ from low-valuation advance customers who reserved, and $\frac{d\Pi_r}{d\eta} = (1-\alpha)A > 0$.*

Case 2. $\eta = c$. *Conditional on the advance population being willing to reserve, the same $(1-\alpha)A$ low-valuation advance customers purchase service, paying p, instead of becoming no-shows and paying η. Any no-show fee that customers actually pay to the firm must be strictly less than c, and if $v_l \geq 0$, then $\eta < c \Rightarrow \eta < p$. Therefore, the firm's expected profit increases over the entire range of possible no-show fees, up to and including $\eta = c$.*

□

If $v_l < 0$, i.e. the low valuation is below the outside option value, then the maximum possible no-show fee $c = p - v_l$ is greater than p. Thus, there is a non-empty range of possible fees (p, c) that will result in no-shows making a payment of more than the service price, conditional on advance customers reserving when such a fee is charged. For all-reserve to be the unique equilibrium strategy at a fee higher than p, it must be the case that $\eta_1 > p$, or equivalently, $v_h > \frac{(1-\alpha F(\kappa-\alpha A))p}{\alpha \bar{F}(\kappa-\alpha A)}$. Qualitatively speaking, this condition is satisfied when the high valuation multiplied by the probability of wanting service is high relative to the service price, and the chances of getting service without a reservation are relatively low. This corresponds to a popular firm (small capacity relative to expected demand) and a large variance of the advance customer service valuation: the low-type value is very low (below the outside option value, so customers would not want the service even if it were free), while the high-type value is very high. Valuation is realized before customers consume the service in our model, so a high variance does not reflect a polarizing service experience (i.e. "love-it-or-hate-it"), but rather a large disparity in how much customers end up wanting to visit the firm *a priori*. For the remainder of this paper, we focus on markets in which v_l is at least as high as the outside option value. This captures a setting where advance customers do not always value the service at more than its price but would always want service if it were free, which best describes the service settings we are interested in analyzing.

In a waiting equilibrium (i.e. if no advance customers reserve), in the second period there are κ seats available and $\alpha A + X$ total customers interested in those seats: αA advance high types and a random number X high-valuation walk-in customers. The firm's expected profit is:

$$\Pi_w = (\alpha A + S_X(\kappa - \alpha A))\, p \tag{2.5}$$

Total customer surplus in a waiting equilibrium is:

$$CS_w = (\alpha A + S_X(\kappa - \alpha A))\, v \tag{2.6}$$

Theorem 1. *Three types of equilibria are possible.*

1. If $\eta_1 \geq c$, then $\eta^* = c$. All advance customers reserve and purchase service.

2. If $c > \eta_1 \geq \theta$, where $\theta = \frac{(S_X(\kappa-\alpha A)-S_X(\kappa-A))p}{(1-\alpha)A}$, then $\eta^* = \eta_1$. Advance customers reserve, but only high-type reservation holders purchase service.

3. If $\eta_1 < \theta$, then $\eta^* > \eta_1$, i.e., any fee strictly above η_1 is optimal and yields the same outcomes for the firm and consumers. Advance customers do not make reservations.

Proof. Lemma 2 shows that if advance customers reserve in equilibrium, the firm will choose the highest fee such that they are willing to reserve. In case (1), a waiting equilibrium is not possible, since advance customers reserve for any fee the firm charges. It follows immediately that $\eta^* = c$ is optimal in this case. In cases (2) and (3), the highest fee such that customers will reserve is some $\eta_1 < c$. The firm can induce customers to wait by charging $\eta > \eta_1$, which yields higher revenue if and only if $(\alpha A + S_X(\kappa - \alpha A))\, p > (\alpha A + S_X(\kappa - A))\, p + (1-\alpha)A\eta_1$, i.e. if $\eta_1 < \frac{(S_X(\kappa-\alpha A)-S_X(\kappa-A))p}{(1-\alpha)A} = \theta$. □

Equilibrium Analysis

Theorem 1 shows that there are three types of equilibria possible in the game between the firm and consumers. The first type of equilibrium occurs when service price p is low, the value of a reservation is high (as indicated by a high α and/or v_h), and capacity is low relative to demand (measured by A and the distribution of X). When the condition in Theorem 1 Part 1 is satisfied, reservations are so valuable to advance customers that they will reserve at any fee level; as a result, the firm chooses a fee that makes even low-types keep their reservations, and the equilibrium outcome has all advance customers reserving and showing up to purchase service. In the second type of equilibrium, advance customers aren't willing to reserve at the maximum fee level $\eta = c$, but the highest fee they will accept is large enough to compensate the firm for the capacity that will be wasted due to no-shows. Thus, the firm chooses to charge a fee that will make advance customers reserve, advance high-types purchase service, and advance low-types pay the no-show fee. In the third type of equilibrium, the maximum fee that could be collected from no-shows is too low to make it worthwhile to allocate capacity to advance customers who may not show up. This type of equilibrium occurs when service price p is high (resulting in a large loss from wasted capacity) and α is low (advance customers are likely to be no-shows).

In the first and third types of equilibria, the no-show fee influences the behavior of strategic advance consumers. In Case 3, the no-show fee is not directly generating revenue at all; its purpose is to make customers *not* reserve (which they always would do in a free reservations scenario) in cases where the wasted capacity from giving out reservations is too high to be justified by the fee revenue that could be collected. In Case 1, the no-show fee impacts customer behavior in a different way, by incentivizing low types to show up when they would not have done so without a fee. These customers do make a payment to the firm, so the no-show fee increases revenue, but the incremental revenue comes from customers consuming the service, not from the firm collecting on the no-show fee. In the second type of equilibrium, Case 2, advance customer actions are exactly the same as in a free reservations setting: everyone reserves; high types show up to purchase service; low types become no-shows. The difference is that low-types pay a fee to the firm, so in this type of equilibrium the no-show fee's purpose is revenue generation. Figure 2.2 illustrates conditions under which the different types of equilibria occur, as a function of model parameters.

When advance-segment consumers place a sufficiently high probability on having a high valuation, the firm can charge the maximum no-show fee and customers will still reserve.

Lemma 3. *For any advance market size A, there is a threshold $\underline{\alpha}_A < 1$ such that the Type 1 fee policy emerges in equilibrium. The larger is the advance market size relative to capacity, the lower is the threshold high-valuation probability $\underline{\alpha}_A$.*

Proof. The maximum fee such that customers will reserve is given by $\eta_1 = \frac{\alpha \overline{F}(\kappa - \alpha A)v}{1-\alpha}$. Type 1 equilibrium occurs when $\eta_1 \geq c$. Note that as $\alpha \to 1^-$, the numerator approaches the finite value $\overline{F}(\kappa - A)v$ while the denominator approaches 0. Therefore, $lim_{\alpha \to 1^-} \eta_1 = +\infty$. For any finite c there exists an $\alpha < 1$ such that $\eta_1 > c$. Additionally, we have η_1 increasing in both A and α. To show this for A we check the partial derivative: $\frac{\partial \eta_1}{\partial A} = \frac{\alpha^2 f(\kappa - \alpha A)v}{1-\alpha} > 0$. For α, we see that the numerator of the η_1 expression, $\alpha \overline{F}(\kappa - \alpha A)v$ has partial derivative with respect to $\frac{\partial}{\partial \alpha} = \overline{F}(\kappa - \alpha A) + \alpha A f(\kappa - \alpha A) > 0$. The denominator $1 - \alpha$ is decreasing in α; therefore, η_1 in increasing in α. The higher is α, the lower is the A such that the condition $\eta_1 \geq c$ is satisfied. □

Theorem 2. *A no-show fee policy results in higher customer surplus, compared to free reservations policy, if and only if $\eta^* > \eta_1$.*

Proof. From equation (2.2), total customer surplus in the free reservations case is $(\alpha A + S_X(\kappa - A))\, v$. In equilibria of types (1) and (2), where advance customers reserve, total customer surplus is given by $(\alpha A + S_X(\kappa - A))\, v - (1 - \alpha)A\eta$, which is the previous expression less some non-zero no-show fee payments (or utility losses from low-types purchasing service). Consumers are not better off under a positive no-show fee if they reserve in the resulting equilibrium. In a type (3) equilibrium, total customer surplus is $(\alpha A + S_X(\kappa - \alpha A))\, v$, which is higher than the free reservations utility because $S_X(\cdot)$ is an increasing function. Thus, consumers are better off under a positive no-show fee if and only if a type (3) equilibrium emerges, i.e. if $\eta^* > \eta_1$. The total increase in customer surplus is equal to $(S_X(\kappa - \alpha A) - S_X(\kappa - A))\, v$. □

Theorem 2 shows that in our base model, a positive no-show fee can increase total customer surplus, but only in cases where the fee discourages advance customers from reserving. Under free reservations, advance customers always reserve (they have no incentive not to), and low type reservation-holders always become no-shows. However, the deadweight loss that results from unused capacity lowers the total customer surplus across the advance and walk-in populations, relative to the free reservations case. In this model, customers are always better off on the aggregate when advance customers do not reserve, and a no-show fee policy is able to incentivize that behavior in markets where the optimal fee induces a type (3) equilibrium.

Corollary 3. *In markets where Case 3 of Theorem 1 emerges, i.e. where $\eta^* > \eta_1$, the no-show fee pricing policy leads to higher social welfare than a free-reservations policy.*

Corollary 3 follows directly from the characterization of total customer surplus in Theorem 2. The firm must be at least weakly better off with a no-show fee, since it chooses the optimal fee level from a set that includes zero (free reservations). Therefore, if customer surplus is higher under a given equilibrium, then total social welfare (the sum of firm profits and customer surplus) is also higher.

This section characterizes equilibrium no-show fee policies without opportunities for cancellation or reoffering. Depending on market conditions, the firm may (1) charge a high fee that makes all advance customers reserve, but eliminates no-shows; (2) charge a somewhat lower fee that does not eliminate no-shows, but recovers some revenue from no-show customers; or (3) charge such a high fee that advance customers don't reserve at all, thereby eliminating no-shows through a different mechanism. We show that the no-show fee policy can make customers better off, in addition to benefiting the firm, although only in the case where the fee discourages advance customers from reserving.

2.6 Conclusion

This paper analyzes the no-show fee decision of a monopolist service firm that offers reservations to strategic advance customers. We show that there are three different categories for the optimal no-show fee, based on the market outcome it produces, and in equilibrium the firm may use the no-show fee to *influence behavior* of strategic customers or to *generate revenue*. Furthermore, no-show fees can make both customers and the firm better off, by allowing the firm to implement a better allocation of capacity across different consumer populations. Our analysis helps to explain why no-show fees are becoming more and more common in service industries, and shows that customer backlash against no-show fee policies is unfounded in some cases. From a managerial perspective, firms may be able to overcome the negative perception of fees by communicating their societal benefit to consumers.

Figure 2.1: Sequence of Events

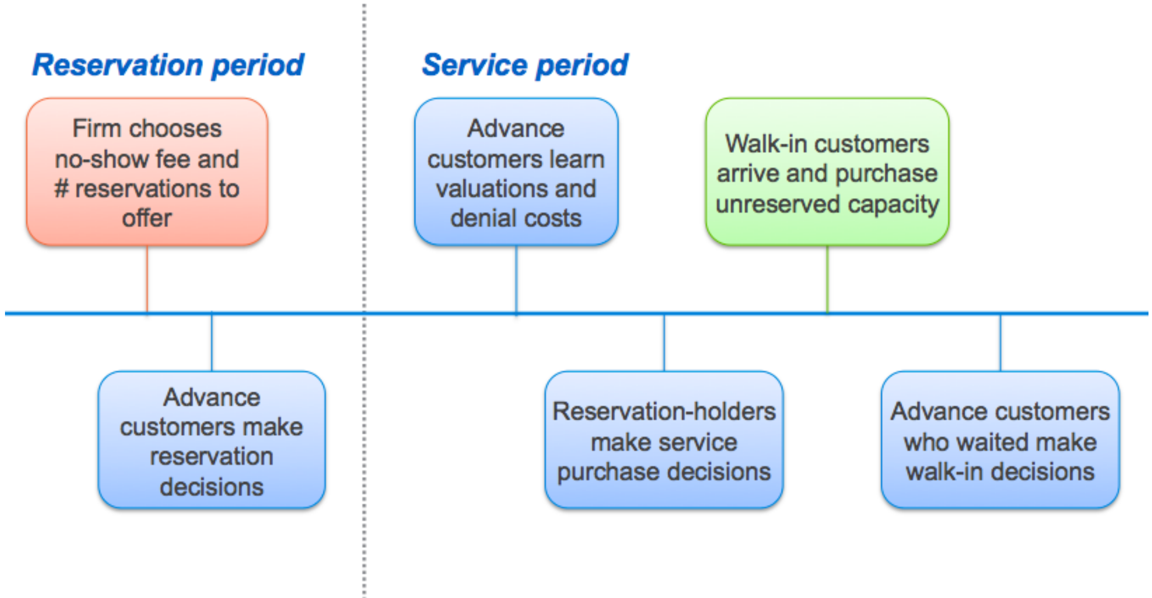

Figure 2.2: Types of Equilibria as a Function of Model Parameters

$v_h = 2,\ p = 1,\ v_l = 0.5,\ X \sim U[0, 1.5\kappa]$

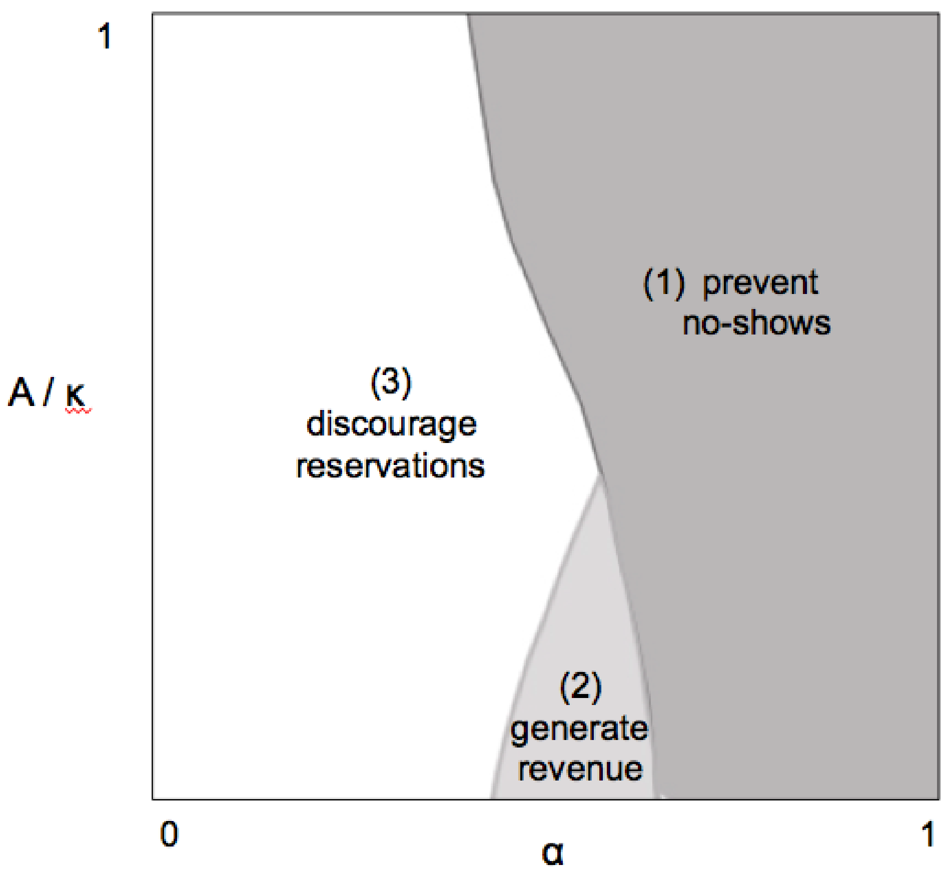

Chapter 3

Inventory Announcements Setting

3.1 Introduction

In recent years, online retailers and service providers have increasingly chosen to provide information about future price and inventory levels to their customers, using messages such as "Only 1 left in stock" or "2 units left at this price." Just as a nearly empty shelf in a brick-and-mortar store suggests that inventory is limited, these announcements convey a sense of scarcity. In the latter example, the inventory message is specific to the current price level, suggesting that more units may be available, but at a different or uncertain price.

How informative are the inventory announcements made by firms, and how do they affect customer decisions? In this paper, we analyze these questions in the context of the air travel industry. If customers are rational and forward-looking, then we would expect them to interpret messages from the firm strategically, and to incorporate any credible information into their purchase timing decisions. However, it is not clear how much information these messages contain. Their content is not directly verifiable, and if the firm prefers that customers buy early, it has an incentive to encourage purchases by creating a feeling of scarcity. On the other hand, repeated interactions with a firm could reveal "untruthful" or misleading announcement patterns to customers, who might subsequently punish firms that do not announce truthfully. Thus, long-term reputational concerns create a plausible incentive for firms to report accurate inventory information.

Evidence from the online frequent flyer community FlyerTalk supports a lack of agreement among consumers about the informational value of airline inventory announcements. Travelers questioning the veracity of an inventory message appear repeatedly in the FlyerTalk forums, and responses range from unequivocally dismissive ("It's a classic case of buy-or-die marketing...I would just ignore it") to essentially trusting ("It will generally reflect what is available at that moment at that price") (FlyerTalk 2015). Given that the information content of these messages is unclear, the response we would expect to see from rational customers is ambiguous. Furthermore, even if consumers believe a firm's announcements, the signaling effect can work in both directions. If customers know that an airline uses inventory announcements, but no announcement is shown for a particular itinerary, it may signal that there is no rush to buy, as the price is unlikely to increase in the near future. Similarly, if the announced quantity remaining is high enough, it may incentivize waiting rather than immediate purchase. In light of the uncertain customer response to announcements, it is not obvious how implementing an inventory messaging policy will affect firm revenues.

Several modeling papers have examined the credibility of cheap-talk announcements in both inventory and queueing settings (Allon and Bassamboo 2011; Allon, Bassamboo, and Gurvich 2012). However, there is scant empirical evidence of how firms use these messages in practice (i.e., how accurately they communicate the amount of inventory available at a given price), and of how consumers actually respond to them. In one recent paper, Yu et al. (2016) use data from a call center to measure the impact of delay announcements on customers' decisions about whether to abandon the queue. To the best of our knowledge, there has been no similar empirical investigation in an inventory setting, where the messages relate to the number of units remaining at the current price. Inventory announcements are qualitatively different than wait-time announcements in terms of the customer's ability to verify the message. A service customer who decides to wait in a queue will eventually learn the actual waiting time, but customers who are in the market for airline tickets or other goods are unable to observe true the inventory level at the current price. We contribute to the literature by analyzing the use of inventory announcements in the air travel industry.

The market for air travel is a particularly interesting setting for this research question, for several reasons. First, the airline industry as a whole operates under wafer-thin margins of less than 3%, with even lower margins for the U.S. legacy carriers (PWC 2015). Therefore, any decisions that influence demand patterns are particularly consequential. A thorough understanding of consumer response to inventory announcements could yield millions of dollars' worth of value for an airline by allowing for even finer control of revenue management practices than is currently standard in the industry. Second, although inventory messages are common enough to influence millions of purchase decisions, there is considerable variation in announcement practices across firms. For example, a customer searching for airline tickets on United.com may see an itinerary with the announcement "8 seats left at this price"; on Delta and American, messages are shown only when the announced number of seats is 3 or fewer. Meanwhile, the booking website for the now-defunct US Airways did not show inventory announcements for any flights. The disparity in policies both provides fertile ground for empirical research and suggests a possible lack of consensus about the optimal announcement strategy. Third, inventory announcements could be particularly valuable to customers due to the volatile nature of air-ticket prices. Because price changes for airline tickets are frequent and often large in magnitude, the problem of correctly timing the ticket purchase is an important one for consumers. Li et al. (2014) show that there is evidence of strategic consumer behavior in the market for airline tickets; if the messages provided by the firm are informative, then the strategic customer segment will take them into account when deciding when to buy. In such an uncertain environment, any information that enables better purchase timing decisions is highly valuable to customers, if they can interpret it correctly.

This research answers two central questions. First, how informative are airline inventory announcements about future price changes? We define 'informativeness' as the ability of inventory announcements to improve customer predictions of future price changes: both direction and magnitude. Second, how do announcements affect customer purchase decisions? Using multinomial logit regression to model the log-odds of directional price changes, we find that itineraries displaying announcements are much more likely to change price than are comparable itineraries without announcements. Both price increases and price decreases are more likely after an inventory announcement, and although the relative probability of increase increases in most cases, there are times in the booking horizon (close to departure) when an announcement can signal a higher likelihood of price decrease. Inventory announcements also predict price changes that are larger in magnitude, compared to price changes that occur when no announcement is shown. While the implications of announcement for price changes are qualitatively the same for the three airlines in our study, we find differences in the estimated sales impact across carriers. For United Airlines, our analysis shows a significant positive impact of announcement changes on immediate sales. This result is robust to the inclusion of various market- and itinerary-level controls, and persists when we control for potential endogeneity using instrumental variables. After correcting for bias from endogeneity, we also estimate a positive sales impact of announcements for American Airlines, while Delta's announcements have a null effect in our sample. The specific sales impact findings are explored in more detail in the results section.

The rest of the paper is organized as follows. Section 3.2 reviews related literature. Section 3.3 discusses our data collection, preparation, and validation processes. Section 4.1 develops a model for quantifying the informational value of announcements, and discusses results. Section 4.2 describes the empirical approach for measuring the sales impact of inventory announcements, and discusses results and implications for the airlines. Section 4.3 concludes the paper and highlights opportunities for future research.

3.2 Related Literature

We seek to identify the impact of inventory announcements on customer purchase timing; as such, our research relates to the substantial theoretical literature on operations management with strategic consumer behavior. Previous research has examined many different aspects of strategic customer behavior, including its implications for optimal pricing decisions (Besanko and Winston 1990, Jerath et al. 2010, Cachon and Feldman 2015); procurement strategies (Cachon and Swinney 2009, Swinney 2011); and capacity rationing (Liu and van Ryzin 2008). Papers in this area show that the presence of strategic customers can decrease firm revenues in some cases (Aviv and Pazgal 2008, Levin et al. 2009), but also benefits firms under certain conditions (Su 2007). However, any potential benefits are contingent upon the firm recognizing and accounting for the existence of strategic customers: a common theme in much of this work is that neglecting strategic customer behavior can decrease firm revenues substantially.

On the empirical side, most of the relevant work comes from economics and marketing.

In the economics literature, Hendel and Nevo (2006) and Hendel and Nevo (2013) analyze consumer stockpiling behavior for storable goods that are subject to price promotions. They show that the firm's optimal pricing policy includes temporary price reductions to discriminate between consumers who are more price-sensitive and those who are less price-sensitive. Erdem and Keane (1996) model dynamic brand choice when brand attributes are uncertain, and forward-looking consumers maximize expected utility over a planning horizon. Nair (2007) analyzes purchase timing decisions of rational customers in the market for console video games, when the firm uses intertemporal price discrimination to segment consumers. An underlying assumption in these papers is that customers behave strategically, in the sense that they consider future prices and availability when making the current-period purchase decision. Li et al. (2014) were the first to test empirically for evidence of strategic consumer behavior; like our paper, they consider the air travel setting. They estimate that between 5.2% and 19.2% of air-travel customers are strategic, and find that evidence of strategic behavior is robust to a variety of modeling assumptions. Osadchiy and Bendoly (2010) also test for the presence of strategic behavior, but in a lab experiment rather than with field data.

Within the marketing and psychology literatures, a number of papers analyze the impact of perceived limited availability on purchases or intent to purchase. A common theme in this area of work is that products are viewed more positively when they are in limited supply, a theory that is discussed in Lynn (1992). Aggarwal et al. (2011) measure the impact of two types of 'scarcity messages,' which indicate either a limited quantity available at a promotional price or limited time to take advantage of a promotion, on consumer intent to purchase. Using data from lab experiments, they find that limited-quantity scarcity messages have a larger impact on customer purchase intention than limited-time messages, although both message types have a positive and significant effect. Byun and Sternquist (2012) study the effects of perceived limited availability using survey data collected from fast fashion retail shoppers. In a key difference from our paper, there are no inventory announcements in the setting they consider; customers perceive availability to be limited because they know that the retailer's inventory renewal cycles are short. Therefore, the perception of scarcity does not influence buying decisions at an item level, but rather through an overall acceleration of purchases. Our research differs from these papers in our use of a large-scale real-world dataset and our focus on measuring the information content of messages claiming low inventory at a given price.

The emphasis of our paper differs from the aforementioned empirical literature in that our analysis focuses specifically on the impact of inventory announcements on purchases. The question of how much information announcements contain relates to the classical cheap talk literature that began with Crawford and Sobel (1982), and to more recent papers in the operations literature that look at the abliity of cheap talk messages to convey information in both retail and queueing settings (e.g., Allon and Bassamboo 2011; Allon, Bassamboo, and Gurvich 2012). These papers develop an analytical framework to understand the dynamics of a cheap-talk messaging game between a firm and customers: specifically, the conditions under which messages convey information (are "informative," in the language of our paper), and the equilibrium consumer response to the firm's announcements. We complement and expand this literature with an empirical study of these research questions, focusing on the market for airline tickets.

3.3 Data

Empirical analyses of strategic customer behavior are scarce in the existing literature, in large part due to the difficulty of obtaining the data to facilitate such a study. For our paper, we use a unique original dataset consisting of itinerary listing and seat map data for three airlines. This section provides a description of the data we collected, high-level summary statistics and variable definitions, and an overview of the data preparation process.

Data Description

To answer our research questions, we use data on inventory announcements and ticket sales for a large sample of domestic flights on U.S. airlines. With the assistance of a third-party data collection company, we constructed the original dataset for this research by extracting itinerary listing and seat map data from the booking websites of three major U.S. carriers that make inventory announcements: American Airlines, Delta Air Lines, and United Airlines. For each carrier in our study, we used data from the Bureau of Transportation Statistics' (BTS) Airline Origin and Destination Survey to identify near-monopoly markets, i.e., origin-destination pairs where at least 80% of itineraries between the two cities are operated by the relevant carrier.[1] Of the routes meeting this criterion, we selected markets with enough records in the BTS sample to imply an average of at least 20 ticket sales per day across all available flights. Because a single airline accounts for almost all traffic on these near-monopoly routes, the search results shown on the operating airline's website are a good representation of the choice set customers face when booking a trip in that market. Across the three airlines, we identified a total of 64 near-monopoly markets for inclusion in this study. A full list of these origin-destination pairs is provided in Table 3.1.

For each selected market, we collected itinerary and seat map data for roundtrip flights with daily departure dates from January 15, 2015 through April 15, 2015. For each departure date, we collected data for round-trip itineraries with return dates both three and seven days later, giving us 182 departure/return date combinations. Data collection took place daily, beginning two months in advance of each departure date and continuing until the travel date (see Table 3.2). In our analyses, we exclude observations captured less than 48 hours prior to a flight's departure, as strategic waiting is not a factor at that time.

Each observation in the itinerary listing dataset is a fare record that contains origin and destination airports; carrier; inventory class; ticket price; travel date and time; observation date; inventory announcement; and itinerary information such as number of stops, connecting

[1] On average, the markets in our study are even closer to true monopoly, with 90% of itineraries operated by the target airline.

city, and travel duration. This information is collected from the list of itineraries shown after searching for tickets between a given origin and destination on a given set of travel dates. We limit our data collection to itineraries with at most one connection.

Across all of the markets included in our analysis, we have between 1.6 and 3.3 million departure-leg itinerary listing observations for each airline.[2] For the three announcing airlines, the percent of itineraries showing inventory messages are 15.9% (Delta), 21.4% (American), and 32.0% (United). There is a strong time trend of more frequent announcement later in the booking horizon, as shown in Table 3.3. There is signficant price variation in our data, as shown in Table 3.4, and price volatility increases closer to the departure date.

Seat Map Data

We cannot directly observe ticket purchases for the flights in our study; instead, we use changes in seat availability over time as a proxy for sales. Within each itinerary option listed on a search results page, the airlines in our study include a link to view available seats on the aircraft serving the flight(s). From the seat maps, we collect the number of economy-cabin seats that are available and unavailable.[3] Occasionally, seat maps cannot be displayed due to an error loading the page or the inavailability of seat information from an operating partner.[4] However, more than 85% of the itineraries in our sample have seat map data available for all flights.

Previous research, such as Williams (2013), has shown that airline seat map changes are a useful proxy for bookings, although they overstate the load factor by an average of roughly 10% due to blocked seats (i.e., seats that are not filled by travelers, but which appear unavailable due to various booking restrictions). Because we are interested in availability changes rather than absolute levels of availability, overstating unavailable seats by a constant amount or factor is not particularly troublesome.[5] What is potentially more problematic is if there are seat map changes that result from opening and closing fare classes; traveler upgrades from economy to premium cabins; or other events besides actual ticket sales. Although we cannot perfectly control for false implied sales, we perform a number of robustness checks to test whether they are likely to bias our results. As a first step, we identify observations that are clear outliers or suggest an error in collection (roughly 2% of observations), and exclude them from our analysis. We also check whether the distribution of seat map changes shifts

[2] American: 1.6 million observations; Delta: 3.3 million; United: 1.9 million.

[3] For flights that have more than one type of economy cabin seat (e.g. regular economy and "Economy Plus"), we collected availability at the seat-category level, although our analysis in this paper does not distinguish between different types of open seats.

[4] In most cases, even flights that are not operated by the main airline have seat maps available.

[5] If load factors are indeed inflated by a constant factor, such as 10%, then seat map changes will overstate the number of actual sales, because the difference between actual tickets sold and seats shown as unavailable will be increasing in load factor. However, if we know that overstatement of unavailable seats takes this form, we can easily correct for the disparity through a simple transformation of our implied sales variable. We verify that our central results are qualitatively unchanged if load factors are overstated in the manner described in Williams (2013).

in suspicious ways at different points in the booking horizon (i.e., if there are discontinuities at points in time when fare classes might be more likely to open and close, or when airlines would be processing upgrades to premium cabins), and do not find any concerning patterns in implied sales. In particular, after cleaning the data and controlling for outliers, the mean of our sales variable is a smooth function of advance booking days, with the coefficient of variation of sales trending down as the time to departure approaches.

A second limitation of seat map data is that any given flight can be used in many different itineraries, and when a seat fills we are unable to observe which specific itinerary containing that flight was purchased by a customer. The network structure of airline operations means that there are hundreds of itineraries that contain a given leg, with many different destinations and travel date combinations. However, for our purposes, we only need seat map changes to have sufficient correlation with actual sales on the route in question. If announcements across different itineraries that contain the same leg are highly correlated–both for itineraries with the same origin and destination but different dates for either departure or return, and for itineraries that share a leg but connect to different origins or destinations–then even if observed seat map changes result from sales on a different itinerary than the one in our study, the announcement status of the trip that was purchased is consistent with the announcement status of the corresponding observation in our dataset. That is, we will still correctly link sales to either an 'announcement itinerary' or a 'no-announcement itinerary.' Our data allows us to test whether this condition is satisfied.

The two lengths of stay for which we collect data for each departure date make it easy to test the correlation of announcements across itineraries that differ only in their return date, e.g. the same flight from Houston to Cleveland on March 1, with either a March 4 or March 8 return date. Both itineraries use the same departing flight; observing the same announcements on both supports the hypothesis that even if sales on the March 1 flight came from itineraries with different return dates, those sales likely took place with the same announcement status as the one used in our dataset. The results are very encouraging: over 98% of the time, if a given departing flight shows an announcement for the trip returning 3 days later, it also shows an announcement for the trip returning 7 days later (and vice versa). Additional details of our empirical methodology and identification strategy for the sales impact of announcements are provided in Section 4.2.

Variable Definitions

Throughout this paper, we use itinerary-observation day as the unit of analysis. In section 4.1, we have two different dependent variables. The first is one-period directional price change, a categorical variable equal to 'increase,' 'decrease,' or 'no change,' according to the type of price change observed for itinerary i from day t to day $t+1$. The second is price change magnitude, the amount of the price change over the same time frame. In section 4.2, the outcome variable is our seat-map change proxy for sales, which we calculate using single-period changes in available seats. A more detailed discussion of various ways to define this variable, and the intuition and implications of each, is presented in that section.

Our independent variable of interest is announcement status, which takes on several different definitions. In analyzing the predictive power of announcements, we use a simple binary indicator *announce* for most of our analyses, which is equal to 1 if any announcement is shown and 0 otherwise. When we are interested in measuring the differences in information content across specific announcements, we use individual indicator variables for each level (1 seat, 2 seats, etc). For United Airlines, we perform some regressions with an indicator for an announcement of 3 or fewer seats remaining, to facilitate comparison with the other carriers. In Section 4.2, our variable of interest is *announceDecrease*, an indicator for a new or decreased inventory message. This variable is equal to 1 when there is a change from no announcement to any announcement level, or when there is a change from some announcement level x to announcement level $x - y$, $y \in \{1, ..., x - 1\}$. We use this variable instead of the binary indicator for whether any announcement is shown because the latter could lead us to underestimate the impact of announcement information. If the same message is shown for multiple days, periods beyond the initial announcement day will not show the full impact of announcement, because much of the effect will have been realized when the new information was first revealed. At the same time, if the effect of a new message does span multiple days, then the 1-day sales response to our announcement variable will be a conservative measure of the true effect.

The explanatory variables in our regression analyses include the natural logarithm of itinerary price; origin-destination pair, day-of-week, and length-of-stay indicators; load factor (measured as the percent of seats on the flight that are full); the number of itinerary options available for a given origin-destination pair and travel date, which is a measure of the availability of flight substitutes; market-level announcement intensity (measured as the percentage of itinerary options showing an announcement); and market price level (measured as median price, average price, or minimum price across all itinerary options). We also include *advweek*, the number of weeks in advance of the departure date that a given observation is recorded, as well as the square and cube of this variable. We include the squared and cubic terms because we believe there is a nonlinear relationship between sales and time to departure. Because we do not know the precise functional form of this relationship, we follow the standard practice of approximating the function with a polynomial. Our final specification uses a cubic function because the incremental explanatory power from adding higher-order polynomial terms to the regression is essentially zero. A common practice in constructing fare classes is to require 7-, 14-, or 21-day advance purchase, so we also include indicators corresponding to those points in the booking horizon, *7days*, *14days*, and *21days*. The cubic time trend and threshold day indicators allow us to flexibly estimate the underlying demand pattern for airline tickets, while the market and departure date indicators allow the regression intercept to vary with market- and date-specific factors. For reference, definitions of the independent and explanatory variables are given in Table 3.5.

Table 3.1: Origin-Destination Pairs Included in the Analysis

(with both directions of travel)

Airline	**Airports**	**Origin & Destination**	**Airline**	**Airports**	**Origin & Destination**
American	BTR-DFW	Baton Rouge, LA - Dallas, TX	**Delta, cont.**	ATL-MEM	Atlanta, GA - Memphis, TN
	COS-DFW	Colorado Springs, CO - Dallas, TX		ATL-SRQ	Atlanta, GA - Syracuse, NY
	CRP-DFW	Corpus Christi, TX - Dallas, TX		BOI-SLC	Boise, ID - Salt Lake City, UT
	DFW-DSM	Dallas, TX - Des Moines, IA		DTW-IND	Detroit, MI - Indianapolis, IN
	DFW-MFE	Dallas, TX - McAllen, TX		GRR-MSP	Grand Rapids, MI - Minneapolis, MN
	DFW-OMA	Dallas, TX - Omaha, NE		MKE-MSP	Milwaukee, WI - Minneapolis, MN
	DFW-TUS	Dallas, TX - Tucson, AZ		MSP-OMA	Minneapolis, MN - Omaha, NE
	DFW-XNA	Dallas, TX - Northwest Arkansas	**United**	ALB-ORD	Albany, NY - Chicago, IL
	EGE-MIA	Vail, CO - Miami, FL		AUS-IAH	Austin, TX - Houston, TX
	MCO-MIA	Orlando, FL - Miami, FL		BIL-DEN	Billings, MT - Denver, CO
Delta	ATL-BNA	Atlanta, GA - Nashville, TN		CLE-IAH	Cleveland, OH - Houston, TX
	ATL-CVG	Atlanta, GA - Cincinnati, OH		CLE-MKE	Cleveland, OH - Milwaukee, WI
	ATL-FNT	Atlanta, GA - Flint, MI		FAT-LAS	Fresno, CA - Las Vegas, NV
	ATL-GSO	Atlanta, GA - Greensboro, NC		HNL-IAH	Honolulu, HI - Houston, TX
	ATL-JAN	Atlanta, GA - Jackson, MS		IAH-OMA	Houston, TX - Omaha, NE
	ATL-LIT	Atlanta, GA - Little Rock, AR		IAH-SAT	Houston, TX - San Antonio, TX

Table 3.2: Travel Dates and Data Collection Dates

DateID	**Departure Date**	**Return Date**	**Data Collection Start Date**
1	Jan. 15, 2015	Jan. 18, 2015	Nov. 15, 2014
2	Jan. 15, 2015	Jan. 22, 2015	Nov. 15, 2014
3	Jan. 16, 2015	Jan. 19, 2015	Nov. 16, 2014
4	Jan. 16, 2015	Jan. 23, 2015	Nov. 16, 2014
...	...	...	...
181	Apr. 15, 2015	Apr. 18, 2015	Feb. 15, 2015
182	Apr. 15, 2015	Apr. 22, 2015	Feb. 15, 2015

Table 3.3: Average Announcement Frequency by Time to Departure Date

	Weeks to Departure									**Overall**
Airline	**9**	**8**	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**average**
American	14.5%	14.5%	16.3%	18.3%	19.6%	21.8%	22.5%	24.2%	27.4%	21.4%
Delta	9.8%	9.8%	10.7%	12.2%	13.6%	17.1%	17.6%	18.5%	26.1%	15.9%
United	23.9%	25.2%	27.1%	27.7%	29.0%	33.6%	33.2%	35.8%	44.3%	32.0%

Table 3.4: Frequency and Size of Price Changes by Market and Booking Time

		Price change probabilities			**Avg. change if fare decreases**		**Avg. change if fare increases**	
	Avg. Fare	Decrease	No change	Increase	$	%	$	%
American	426.38	13.8%	60.6%	25.6%	-67.86	-10.4	70.67	15.3%
Delta	601.91	11.4%	67.1%	21.5%	-57.02	-9.6%	83.35	14.5%
United	575.32	16.3%	55.9%	27.9%	-50.06	-9.2%	74.38	14.0%
By weeks to departure	**United Airlines***							
9	484.35	16.7%	65.5%	17.8%	-34.25	-7.3%	40.15	9.3%
8	482.41	16.4%	65.1%	18.5%	-37.60	-8.1%	40.78	9.5%
7	488.19	16.7%	62.3%	21.0%	-36.05	-7.7%	41.36	9.7%
6	499.67	17.2%	60.7%	22.1%	-35.95	-7.6%	44.72	10.2%
5	512.68	18.7%	57.6%	23.7%	-40.47	-8.4%	54.13	11.1%
4	524.30	19.0%	52.6%	28.4%	-46.08	-9.1%	60.13	13.6%
3	563.69	16.6%	48.0%	35.5%	-59.98	-11.4%	89.48	19.0%
2	680.76	13.0%	49.2%	37.8%	-67.47	-10.2%	101.48	17.0%
1	850.34	12.0%	56.0%	31.9%	-95.14	-11.6%	108.27	14.9%

*The price change trends by weeks-to-departure are qualitatively the same for American and Delta, and are omitted to avoid repetition.

Table 3.5: Variable Definitions

(1) Price Change Regressions	
price change category	forward-looking directional price change (from t to $t+1$)
price change	forward-looking amount of price change (from t to $t+1$)
announce	binary indicator for whether any announcement was shown*
booking time	cubic time trend in weeks in advance + 7/14/21 day indicators
(2) Sales Impact Regressions	
sales	average of seats filled from time $t-1$ to time t and seats filled from t to $t+1$
announceDecrease	binary indicator for a new announcement or reduction in announced quantity available at price
price	log of price
load factor	percent of total seats unavailable
market announcement	announcement intensity across other itinerary options, measured as % showing announcement
booking time	cubic time trend in weeks in advance + 7/14/21 day indicators

Chapter 4

Inventory Announcements Results

4.1 Announcement and Price Changes

The first question that we answer related to inventory announcements is the "informativeness" of these messages. Conceptually, we want to understand how much information is contained in airline inventory announcements, and the extent to which customers can use this information to make better purchase timing decisions. A comprehensive answer to this question is a necessary precursor to the next section, in which we evaluate whether the sales response we see after inventory announcements is consistent with strategic consumer behavior: the rational response to an announcement will depend on the information it contains.

Because carriers use sophisticated pricing algorithms that result in the constant opening, closing, and reopening of more than two dozen fare classes, measuring the informational value of inventory announcements is not as straightforward as it may seem on the surface. The inventory messages shown by airlines make a claim about the number of seats left *at the current price*, not about the total number of units remaining, and they say nothing about what the next price level will be. Carriers establish an initial set of 'fare buckets' or 'booking classes' for each itinerary, with some number of tickets assigned to each one, but these buckets and their capacities change over time in response to realized demand on the route. Furthermore, industry conventions such as fare class advance booking requirements introduce another level of nuance to the interpretation of inventory messages. Eight days before the departure of Flight 200 from Atlanta to Nashville, the airline may have three seats left in a certain booking class priced at $400. But the price of that fare bucket may increase tomorrow, regardless of how many tickets are sold today.

Given the complex pricing dynamics of the air travel industry, it is clear that we cannot merely evaluate whether an announcement is truthful. Instead, we seek to understand the marginal information provided by an inventory announcement, conditional on all relevant market conditions (e.g. time to departure; itinerary characteristics; availability of substitutes). We define 'informativeness' as the ability of inventory announcements to improve customer predictions of future price changes: both direction and magnitude. The incremental value of incorporating announcements into the information set is the difference in expected price paid by a strategic customer when inventory messages are used to predict future price changes, compared to the expected price paid without announcements to aid in purchase timing. Because we are interested in predictive power in this section, we do not seek to establish a causal relationship between announcements and price changes.

In the rest of this section, we develop a model to compare the expected direction and magnitude of price changes after an announcement to price changes when there is no announcement, and conclude with a simple model for quantifying the value of inventory messages to a rational consumer who times his ticket purchase optimally conditional on available information.

Directional Price Changes

When a customer is deciding whether to purchase an airline ticket or wait until the next period, his overarching consideration is whether the price of the ticket is likely to increase, decrease, or stay the same. Therefore, we begin by modeling the impact of an announcement on expected directional price changes from one observation period to the next. As background for this section, it is helpful to understand the prevailing price trends in the industry and for our markets in particular.

Price Change Trends

In general, airline ticket prices trend upward as the time remaining until departure decreases; however, the typical price path is non-monotonic. The prices of the flights in our sample change price an average of 11 times in the final month of the booking horizon, with price increases representing two-thirds of those changes. Across all observations for the near-monopoly markets in our data, prices increase from one period to the next approximately 22-28% of the time (depending on the airline), decrease 11-16% of the time, and remain unchanged 56-67% of the time. As expected, price change probabilities display a strong time trend, with increases being more likely closer to departure. Price decreases display a non-monotonic time trend, occuring most frequently 4-5 weeks before departure. At a high level, this pattern is consistent with airlines observing early market signals and adjusting prices in response to realized demand. Plotting daily price change probabilities for United (see Figure 4.1), we see that price increase probabilities are markedly higher on days that correspond to industry advance booking requirements (including 21 days, 14 days, and 7 days before departure).

Looking at the raw directional price changes with and without announcement, a clear pattern emerges. For example, on United's website, if any announcement is shown, there is a 33% chance that the price of that flight will be higher the next day, compared to a 20% chance of price increase if no announcement is shown. If United announces that 3 or fewer seats are left at the current price, the probability of a price increase tomorrow jumps to

41% across all markets in our sample,[1] and as high as 55% for some origin-destination pairs. Price change summary statistics by announcement status for all three airlines in our study are shown in Table 4.1.

The high-level summary statistics suggest a correlation between announcements and price increases, but do not tell us about the marginal information provided by inventory announcements, conditional on other factors that are correlated with prices (such as time to departure). For this, we use a multinomial logistic regression of price change category (increase, decrease, or no change) on our announcement variable and a vector of explanatory variables. Under the multinomial logit model, the logarithm of the probability (log-odds) of observing price change category k for itinerary i (relative to baseline category K) follows a linear model, i.e.,

$$\eta_{ik} = \frac{\pi_{ik}}{\pi_{iK}} = \alpha + \beta_k \mathbf{announce}_i + \mathbf{x}_i' \gamma_k$$

We use 'no change' as our baseline category, and estimate the log-odds of price increase and decrease as functions of announcements and a set of controls. In our primary specification, the vector of explanatory variables $\mathbf{x}$ includes fixed effects for origin-destination pair, departure date, and length-of-stay, as well as a cubic time trend for days until the travel date and indicators for the days commonly used in fare class advance booking requirements (21/14/7 days before departure).

Results

For all three announcing airlines, we obtain consistent results: positive and highly significant coefficients on announcement for both price increase and price decrease. That is, conditional on trip characteristics and time to departure, a given flight is more likely to increase in price tomorrow if an announcement is shown today. However, that flight is also more likely to *decrease* in price than if there had been no announcement. The coefficient magnitudes for the two categories indicate a larger effect for price increases: after announcement, both price increases and price decreases are nominally more likely, but a price increase also becomes more likely relative to a decrease. For example, on United flights, a price increase is on average 77% more likely after announcement compared to after no announcement, while a price decrease is 29% more likely. The relative odds of seeing a price increase vs. a price decrease are 37% higher after announcement, after controlling for market factors and time to departure. Full regression results are shown in Table 4.2.

The multinomial logit model's linear prediction function allows us to control for variables that are correlated with both price change direction and announcement, but does not reveal how the effect of announcements changes with the explanatory variable values. Of particular interest is the relationship between announcement informativeness and advance booking:

[1] Recall that United's announcements can be as high as "8 seats left at this price"; statistics for 3 or fewer are reported to facilitate comparison with American and Delta, which both use 3 as their maximum announcement number.

does an announcement two months before the travel date provide the same amount of information as an announcement close to departure, or does the marginal predictive power of inventory messages change over the booking horizon? To answer this question, we perform separate multinomial logistic regressions for each advance booking week from one to nine weeks before the departure date. Weekly results for the change in the odds of a price increase and a price decrease after announcement (a transformation of the MNL coefficient estimates) are plotted in Figure 4.2.[2] Performing separate regressions by week allows us to model a nonlinear relationship between announcements and days-to-departure. This estimation procedure is more flexible than including announcement-advance week interaction terms in a combined regression, because the coefficients of the explanatory variables in our regressions are allowed to vary with advance booking week as well.

We find a strong time trend in the implications of announcement for the price decrease probability. Eight weeks prior to travel, if an announcement is shown, then the price of that itinerary is 5% more likely to decrease in the next period than if there is no announcement. Meanwhile, in the final week before departure, the probability of a price decrease nearly doubles if an announcement is shown. The change in price increase probability after announcement also exhibits variability over the booking horizon, but does not display the same upward trend. The results in Figure 4.2 imply that announcements have different interpretations at different points in time. Early in the booking horizon, announcement signals that a price increase is much more likely relative to decrease, compared to the case where no announcement is shown. Close to the departure date, seeing an announcement can actually signal that a price *decrease* is relatively more likely than if no announcement was shown. Table 4.3 shows how much the likelihood of price increase vs. price decrease changes after announcement, by booking week.

For customers, our directional price change results suggest that announcements are indicative of higher short-term price volatility. While price increases become relatively more likely post-announcement over most of the booking horizon, the fact that both directions have positive coefficients means that the implications of an announcement will depend on customer preferences. For a risk-averse or risk-neutral customer who is certain she will travel, buying now becomes more attractive relative to waiting if an inventory message is shown. However, for a customer who has other options, or wants to travel only if he can find a good deal, an announcement could signal that it is worth waiting to see if the price drops–at significant risk of a price increase, of course. Furthermore, an announcement two months prior to travel is not the same as an announcement two weeks prior to travel. The relative risk ratios for directional price change show that announcements are more informative further from the travel date, in the sense that they provide a less noisy signal of future price change direction. One interpretation is that there are more factors influencing price changes close to the departure date, and the baseline probability of a price change is higher than it is early

[2]These results use the same set of explanatory variables as our main directional price change regressions; variations that include booking day indicators (both for all days, and for just for the 7/14/21 day milestones, as applicable), yield consistent results.

in the booking horizon. Therefore, an announcement provides less marginal information to consumers when the travel date is near. Another interpretation is that scarcity is not expected early in the booking horizon, and any announcement implies that scarcity exists.

In this section we have shown that announcements improve prediction of price change direction, to a degree that varies over time. To quantify the expected utility impact of this predictive power more precisely, it is necessary to understand the relationship of announcements not just to price change directions but also to price change magnitudes. We turn to this question in the next section.

Magnitude of Price Changes

In the previous section we showed how the expected direction of price changes is affected by inventory announcements; we now test whether customers can also infer something about the expected magnitude of price changes.

Model and Results

To measure the impact of announcements on expected price change magnitudes, we use ordinary least-squares regression to determine the best fit of the dollar-amount or percentage price changes upon a new or updated announcement and the explanatory variables discussed in the previous section. The regression model is:

$$\mathbf{priceChange}_i = \beta_0 + \beta_1 \mathbf{announceDecrease}_i + \mathbf{x}'\boldsymbol{\gamma} + \varepsilon_t \tag{4.1}$$

The explanatory variables that we include in our primary specifications are price; fixed effects for origin-destination pair, departure date, and length-of-stay; a cubic time trend in advance booking week; and indicators for 7, 14, and 21 days before departure. We also performed regressions including market-level announcement intensity and load factor, but found that they were not significant predictors of price change magnitude. Our objective in this section is to establish the main effect of announcements on price change magnitude, so we do not include interaction terms in these regressions.[3]

To check whether the coefficient estimates could be biased by multicollinearity among our explanatory variables, we calculate the variance inflation factor (VIF) of each variable after performing the regression. A VIF equal to 1 indicates no multicollinearity; as a rule of thumb, a variable with a VIF greater than 10 merits further investigation. In our primary regression model, the only variables with VIF greater than 10 (or, in fact, greater than 5) are the advance booking week terms, which have very high values (>100). However, this is not a concern in our setting, for two reasons. First, we expect to have high VIF values when we include powers of an explanatory variable, which are likely to have high correlation

[3] As an input to our calculations in Section 4.5, we perform the regression in specification (1) separately for each advance booking week. This procedure allows us to estimate the trend in announcement impact over the booking horizon and is more flexible than including *announce*/*advweek* interactions in a single regression.

with the variable itself (indeed, the correlations of advweek with its square and cube are 0.97 and 0.92, respectively). If we center advweek at 0 by subtracting its mean from each observation, the VIF values fall to reasonable levels[4] with no change in the point estimates of the other coefficients. Additionally, multicollinearity is only problematic for the variables that are collinear. Most important for our analysis, the announcement variable of interest has a very low VIF of 1.13, indicating that its coefficient estimate is not unreliable due to multicollinearity.

Across the seven specifications for which we report results in Table 4.4, we find that expected price changes are larger in magnitude after announcement, both across all observations and conditional on price increase. Specifications (1)-(5) use the natural logarithm of price as an explanatory variable; specifications (6) and (7) use absolute price. Our results are consistent for both definitions of the price level explanatory variable. In our United markets, the average price change is \$10.50-\$11.20 larger after any announcement compared to no announcement; the difference grows to \$12.97-\$16.37 if we consider just announcements of 3 or fewer seats left. The specific announcement level also matters at a more granular level: using indicators for each distinct announced inventory level as our announcement variables, we see that after a message of "1 seat left at this price," the expected price change increases by nearly \$25. Meanwhile, after a message of "8 seats left," the expected price change is just \$4 larger than with no announcement. The coefficient estimates for announcements of 1-5 seats have non-overlapping 95% confidence intervals; messages of 6-8 seats left have coefficient estimates that are not statistically different from one another.

We obtain qualitatively similar results if the price change magnitude variable is expressed as a percentage instead of a nominal amount. Our estimates suggest different trends in the contribution of the advance-booking-week cubic between the two variable definitions: the fitted function is flat for high values of advweek using absolute price changes, but continues to decline for the percentage price change definition. However, both price change variables follow similar patterns over time (see Figure 4.3), so the difference can be attributed to imperfect curve fitting of the cubic.

The same pattern holds if we condition on a price increase: lower announcement levels consistently predict larger (i.e. more positive) price changes in the next period. However, conditional on price decrease, an announcement of 1 seat left or 2 seats left actually predicts a more negative price change, i.e. a larger decrease. Meanwhile, higher announced quantities predict a decrease smaller in magnitude. One possible interpretation is that while all announcements levels predict higher probability of price change, the change is more likely to be a jump from one fare bucket to another after a very low quantity is announced–because it is more likely that 1 or 2 seats sell from one day to the next, compared to 7 or 8 seats. Other price changes (which are still more likely to occur after an announcement, based on our results) could reflect adjustments to the specific price level of the current fare bucket, which might be plausibly smaller in magnitude (in either direction) than movements from one bucket to the next. Although we cannot be certain that this is the explanation, our results

[4] VIF = 5.51, 1.61, and 6.02 for *advweek*, $advweek^2$ and $advweek^3$, respectively.

provide strong evidence that price change magnitudes are largest–in both directions–after an announcement of just 1 or 2 seats left.

For the regressions here and in the previous section, our results could overstate the marginal information content of announcements if there are unobserved factors that are positively correlated with both price changes and probability of announcement (conditional on the explanatory variables), and are known to consumers. Certain types of demand fluctuations, such as those caused by holidays or special events in a destination city, could fall into this category. However, it is not immediately obvious how such variation in demand would affect announcement probability. If the airline uses finer revenue management techniques (e.g. smaller fare class buckets) when it expects higher-than-average demand, then we might see both more announcements and more (or larger) price increases than usual. But there could also be fewer (or smaller) price movements when demand is known to be higher, e.g. if the airline sets a higher initial price or offers fewer discount fares. Additionally, customers are unlikely to have full visibility into such demand shocks and their impact on prices and announcements, so any overstatement of the predictive value of announcements is likely to be small.

Value of Announcement Information

To quantify the informational value of inventory announcements to a rational customer, we require a model of the customer decision process. For the purposes of illustration, we consider a simple case in which consumers wait for at most one period, and we treat announcement as a binary indicator (i.e., we ignore the differential information content of different announcement levels).

Suppose a customer arrives in the market and observes the current price p_t of his preferred itinerary. His objective is to maximize expected utility, which in this case will be equivalent to minimizing the expected present value of price paid for the ticket. We assume he must buy in either the first or second period, so the problem can be distilled to a single decision: buy or wait in the first period.

First suppose that inventory messages are never available, and that on average (conditional on market characteristics and time to departure), ticket prices in this market increase π_i percent of the time and decrease π_d percent of the time, with average sizes of η_i and η_d, respectively. The expected next-period price is $\hat{p}_{t+1} = p_t + \pi_i\eta_i - \pi_d\eta_d$, and the customer applies some discount factor δ to purchases made in the second period, which reflects the strength of his preference for delaying the purchase decision. Thus, the customer will buy if $p_t \leq \delta\hat{p}_{t+1}$ and wait otherwise.

Now suppose the carrier has a policy of making inventory announcements, and the customer is aware of the impact of inventory information (either announcement or lack thereof) on expected price changes. If the customer sees no announcement, then he will expect the price to increase in the next period with probability π_i^0 and decrease with probability π_d^0. If it increases, let η_i^0 denote the expected amount of the increase, given that there is no announcement. Similarly, let η_d^0 denote the expected amount of a decrease, so expected next period price is $\hat{p}^0_{t+1} = p_t + \pi^0_i \eta^0_i - \pi^0_d \eta^0_d$. Meanwhile, if an inventory announcement is shown, then the customer expects the next period price to increase with probability π^a_i (by expected amount η^a_i), and to decrease with probability π^a_d (by expected amount η^a_d), yielding expected next-period price $\hat{p}^a_{t+1} = p_t + \pi^a_i \eta^a_i - \pi^a_d \eta^a_d$.

Announcements add value for the customer when the decision made under the announcement policy is different than the decision made using the unconditional price change probabilities; that is, if an announcement is shown and p_t falls strictly between $\delta\hat{p}_{t+1}$ and $\delta\hat{p}^a_{t+1}$, or if no announcement is shown and p_t falls strictly between $\delta\hat{p}_{t+1}$ and $\delta\hat{p}^0_{t+1}$. In Table 4.5, we first calculate the unconditional (on announcement status) expected price change by advance booking week. We then combine our results on price change directions and price change magnitudes (performing separate regressions for each advance booking week) to estimate the expected price change after announcement and after no announcement.[5]

Without any further assumptions, the expected price change calculations indicate that even without discounting (i.e. if $\delta = 1$), the existence of inventory information has a positive impact on customer utility over most of the booking horizon. Wherever the expected price change direction differs between the unconditional case and announcement, or between the unconditional case and no announcement, the customer will make a different buy/wait decision as a result of the firm's policy of sending inventory messages. Note that because the optimal decision with no information is usually to buy (due to prices trending up over time), the value of the announcement policy often comes from the *absence* of a message. The estimated impacts of announcements on price-change expectations reported in Table 4.5 are small, but represent high-level averages across all flights at the relevant point in the booking horizon. The value of announcements can be calculated at a more granular level (e.g., the informational value of a message of "3 seats left" on a 7-day trip from Houston to Cleveland) by using only the data for flights meeting specific criteria. Furthermore, if customers have either a multiplicative discount factor or an additive waiting cost/benefit, this will lead to more cases where announcement impacts the customer's decision.

There are a number of simplifying assumptions underlying this example, including the assumption that the customer knows for certain which itinerary he wants to purchase. In reality, customers may wait because they have uncertainty about their travel plans, which makes the optimal itinerary uncertain. This model can be generalized to capture the case in which customers form price change expectations for multiple flights, and maximize expected utility across the different options. The underlying intuition is the same as in the single-itinerary setting, but with higher computational burden for the consumer. The calculation in this section provides a "back-of-the-envelope" method of evaluating announcement utility;

[5]The new increase and decrease probabilities after announcement, π^a_i and π^a_d, are obtained by multiplying the relative risk ratios (from the multinomial logit directional price change regression) by the unconditional directional price change probabilities. The expected price change amounts η^a_i and η^a_d are obtained by summing the unconditional expected price change amounts with the announcement coefficient estimates from weekly regressions of price change amount on announcement + controls. π^0_i, π^0_d, η^0_i, and η^0_d are backed out using the post-announcement parameters, announcement frequencies, and the unconditional price change probabilities and magnitudes.

in particular, we use the observed price change probabilities across the whole sample as the 'unconditional' probabilities π_i and π_d. In reality, the airline's policy decision to make announcements could affect the overall price change frequencies; a full counterfactual analysis would require a structural model.

4.2 Announcements and Customer Purchases

In this section, we shift our focus from the predictive power of announcements to their effect on customer purchase behavior, and resulting implications for the firm. We begin by discussing the specific definition of the sales variable used in this section, and the implications of various alternative definitions. Next, we develop hypotheses for how inventory announcements could affect short-term sales. Using the sales variable that we derive from observed seat map changes, we use multiple linear regression for a preliminary analysis of announcement impact. To conclude the section, we use an instrumental variables approach to control for potential endogeneity issues with the OLS model. We interpret the results of the estimated regression coefficients, and perform robustness checks for our main results.

Sales Variable Definition

To prepare our data for the analysis in this section, we first need to calculate the implied sales values by taking the difference between the number of available seats on each flight from one period to the next.[6] Implied sales should reflect as closely as possible the number of sales that took place while the announcement associated with a given observation was shown on the airline's website. We observe a snapshot of announcements and seat map availability once per day, typically in the late morning or early afternoon hours in the departure airport's time zone.[7] Depending on when the announcement for a given flight first appeared on the airline's website, we could observe the flight's seat map after some sales have already taken place with that announcement showing. Additionally, there may be sales that coincide with the current announcement after we observe the seat map and before the announcement is next updated. Therefore, when deciding which seat map changes to link with a given announcement observed at time t, there are two different measures to consider: the change in available seats from time $t-1$ to t, and the change in available seats from time t to $t+1$.

In preliminary analysis, we explored three potential sales variables: one defining sales as $availableSeats_{t-1} - availableSeats_t$, one defining sales as $availableSeats_t - availableSeats_{t+1}$, and one defining sales as the average of these two values. There are advantages and drawbacks to each of these definitions. For the backward-looking definition, the primary concern

[6] All results discussed in this paper define a period as one day, but we replicate our main analyses using a two-day period length, with consistent results.

[7] Our data daily collection began shortly before noon Eastern Standard Time. Collection for the travel dates we include in our study typically required only a few hours per day.

is reverse causality: it might be the case that sales over this time horizon cause an announcement change, rather than the other way around. However, using the forward-looking definition may systematically underestimate the impact of announcements by failing to include sales that occur shortly after an inventory message is first shown (before we take our snapshot of seat maps and announcement status). Indeed, if announcements shift demand earlier in time (resulting in an immediate increase followed by an offsetting decrease), then using the forward-looking sales definition could yield a negative estimated sales impact, even if the true impact is positive.

We find that announcements have the most explanatory power when we use a sales variable that defines period t sales as the average of seats filled from time $t-1$ to time t and seats filled from t to $t+1$. This corresponds to an assumption that we observe announcements and seat availability in the middle of the horizon between announcement changes, and that sales over this horizon are evenly distributed in the periods before and after observation. In our final analysis, we use an instrumental variables approach to address the potential simultaneity bias issue caused by the inclusion of a backward-looking sales measure. Since we expect any bias introduced by the forward-looking sales measure to be positive, after addressing the endogeneity issue we can interpret the announcement coefficient estimates as a lower bound on the true sales impact.

A limitation of our analysis is that we collect data once per day; however, an airline may change the announcement for a given flight multiple times in a day as sales take place. Especially for collection dates close to departure, we are likely to miss some announcement changes between one observation day and the next if multiple sales occur. As a result, we are inherently limited in our ability to detect a differential impact on sales between rapid changes in announcement status and a single change from one day to the next. However, announcement changes can only influence customer purchases if the customer observes them. If a strategic consumers who is watching an airline ticket checks its price roughly once per day, then our data is a good approximation of the information that the customer collects and uses to make his purchase decision.

Linear regression model for sales

For our preliminary analysis, we specify a linear regression model that includes our *announceDecrease* variable and the controls discussed in Section 3.3. The regressions in this section have the form:

$$\mathbf{sales}_i = \beta_0 + \beta_1 \mathbf{announceDecrease}_i + \mathbf{x}'\boldsymbol{\gamma} + \varepsilon_t \tag{4.2}$$

where the vector of covariates $\mathbf{x}$ consists of the logarithm of itinerary price; a cubic trend in weeks-to-departure; 7-, 14-, and 21-day indicators; and indicators for origin-destination pair and departure date.[8] Table 4.6 shows the results for regression equation (4.2), as well as variations that add market-level announcement intensity and load factor. Standard errors

[8] We performed regressions with additional explanatory variables described in Section 3.3, including load

for the *announceDecrease* coefficient estimates are reported in parentheses. We cluster the standard errors by *marketID* to allow for correlation in the error terms within markets.

The OLS estimates are inconsistent across airlines. In the United sample, the OLS coefficient on our announcement variable is positive and highly significant, indicating a positive relationship between announcement and same-period sales. For American, the OLS model fit is very poor ($R^2 \approx 0.01$), and the coefficients on announcement and our additional explanatory variables are not significant. For Delta, the estimated coefficient is significant and negative, although small in magnitude.

Although these regressions are useful for preliminary estimation, there are two potential sources of endogeneity that could be biasing the estimates of announcement impact. First, *omitted variable bias* will occur if we exclude from our regression variables that are correlated with both announcements and sales. The main concern in our setting is that there could be unobserved factors, such as special events in a destination city, that cause higher demand in that market and are also correlated with announcement changes. In this case, the error term in the OLS model will be correlated with at least one regressor, violating the strict exogeneity assumption $E[\epsilon|X] = 0$. While we cannot explicitly test for or measure bias from omitted variables that are not observed, we can comment on the expected direction and magnitude of the bias. For an omitted variable w, the amount of bias is given by $\frac{\rho_{Aw}\beta_w s_w}{s_A}$, where ρ_{Aw} is the correlation between the announcement variable (denoted here by A) and the omitted variable w, β_w is the partial effect of the omitted variable on sales, and s_A and s_w are the deviations of announcements and the omitted variable, respectively. The direction of bias could be positive if the excluded variables have a nontrivial impact on sales ($\beta_w > 0$) and a positive correlation with announcements ($\rho_{Aw} > 0$). However, if above-average demand causes multiple purchases in quick succession, then we could observe fewer announcement changes relative to purchases when sales are high. This would result in a negative omitted variable bias on the announcement coefficient. The second possible source of endogenity is *simultaneity bias,* or reverse causality. We are interested in measuring the partial effect of announcements on same-period sales, but sales may also cause more announcements. Thus, the positive relationship between announcements and sales may overstate the true effect of announcements. To address both of these issues and isolate the effect that we are interested in measuring, we use an instrumental variables (IV) two-stage least squares (2SLS) model (cf. Angrist and Krueger (1991)), as outlined in the next section. We return to the interpretation of the differing results across airlines after the IV analysis.

2SLS model

The instrumental variables approach can be used when we observe one or more variables that are correlated with an endogenous regressor but have no relationship with the outcome

factor, market-level price measures, and number of alternate flights available in the market. These additional variables are excluded from the results we report here because they did not increase the explanatory power of the regression model.

variable except through that regressor. That is, the new variables (the 'instruments') satisfy *relevance* and *exclusion restriction* assumptions: they are sufficiently correlated with the endogenous variable in question, but uncorrelated with the error term in the original OLS regression, making it valid to exclude them from that model. The intuition behind this approach is that even when a dependent variable is endogenous, some portion of its variation will be uncorrelated with the independent variable, and can be used to identify the endogenous variable's partial effect on the outcome. With a single endogenous variable and a single instrument, IV and 2SLS are equivalent, so we use the terms interchangeably throughout this section.

As an instrument for an itinerary's announcement status, we use the announcement status of the same itinerary, but with a travel date two weeks in the future, observed the same number of days before departure. For example, for an observation of flight UA 100 departing on March 1, captured 12 days before departure, the value of the instrumental variable will be equal to the announcement variable for flight UA 100 departing on March 15, captured 12 days before departure.[9] The relevance of this instrument derives from the carrier's use of the same high-level pricing strategy across the paired flights: for the same itinerary within a short time horizon, the airline is likely to use the same fare bucket structure, resulting in a similar pattern of announcements. Note that we do not need to assume that individual announcement statuses are the same across flights, or even that the fare buckets themselves are the same–all we need is a nontrivial level of correlation between the two variables. It is straightforward to check the relevance assumption by measuring the correlation between the endogenous announcement variable and the instrument, which we denote by z. In our samples the correlations are 0.09 (American), 0.12 (Delta), and 0.15 (United), and in all three cases a standard test of the first stage regression rejects the null hypothesis of weak instruments. The exclusion restriction cannot be tested directly, but there is no plausible mechanism for the instrument to influence sales on the itinerary in question. We choose a travel date far enough away that the instrument itinerary would rarely be a substitute for travelers, so we can be confident that both IV assumptions are satisfied.

The first stage of the 2SLS approach regresses the instrumental variable, and the vector of covariates, on announcement status, i.e.,

$$\mathbf{announceDecrease}_i = \gamma_1 z_i + \mathbf{x}'\boldsymbol{\gamma}_2 + u_i \tag{4.3}$$

In the second stage, the predicted values of announcement decrease are regressed on sales, along with the same set of covariates included in the first stage regression. Although we have a binary dependent variable in the first stage, we follow the recommendation of Angrist and Pischke (2009) and use the standard 2SLS procedure with linear fitted values,

[9] We use a two-week gap to maximize the number of itineraries we are able to match with an instrument. For longer gaps (such as three or four weeks), there will be more travel dates for which we do not have the data required to construct an instrument. Furthermore, while flight numbers are generally stable in the short term, the longer is the gap, the greater is the chance that some flight numbers will change, complicating the itinerary matching process. Aside from these concerns, our analysis is robust to instrument definitions based on three- or four-week gaps.

which is preferred over instrumenting with nonlinear fitted values. The regression equation is:

$$\mathbf{sales}_i = \beta_0 + \beta_1 \mathbf{announce\hat{D}ecrease}_i + \mathbf{x}'\boldsymbol{\beta}_2 + \epsilon_i \tag{4.4}$$

To obtain the correct standard errors, we use Stata's *ivregress* command to implement 2SLS directly, with standard errors clustered by market. Results of the IV model, using the set of regressors from OLS model (1), are presented in Table 4.7, with standard errors shown in parentheses.

For all three airlines, the 2SLS coefficient estimates move in the positive direction compared to OLS. For United, the positive effect size increases from 0.208 to 0.495 after correcting the endogeneity; for American, we obtain a positive coefficient of 0.313 that is significant at the $p = 0.05$ level. For Delta, the coefficient estimate is directionally more negative, but no longer significant due to the lower efficiency of 2SLS relative to OLS. Tests of endogeneity for United and American reject the null hypotheses that announcements are exogenous at the 5% signifiance level ($p = 0.043$ and $p = 0.021$, respectively). However, for Delta, the null hypothesis of exogeneity is not rejected ($p = 0.60$).

The direction of the shift in coefficients from OLS to 2SLS suggests that reverse causality is not driving the positive results; if it were, then the 2SLS estimates would be lower than those from OLS. Rather, the more rigorous analysis provides additional evidence than consumers do respond to announcements, at least on some carriers. For United (American), we can interpret the coefficients as predicting an average of about 0.5 (0.3) more sales after an announcement is shown, compared to the same market conditions with no announcement. For reference, the average of our sales measure across the entire sample is roughly 1.4 for United, and just over 1.0 for American. Given our findings about price change probabilities after announcement, the positive impact of announcements on immediate sales for these two carriers is broadly consistent with rational consumer behavior. Our results suggest that consumer response to announcements varies across carriers, inviting further exploration of the differential impacts of announcement policies. We observe carrier differences in both announcement threshold (eight for United, vs. three for Delta and American), and in frequency of announcement across airlines with the same threshold (21.4% of itineraries show announcements for American, vs. 15.9% for Delta). These disparities in how the airlines use inventory messaging may explain some of the differences we see in consumer response.

4.3 Conclusion

In this paper, we evaluate the information content of airline inventory announcements, and estimate their effect on customer purchases. We show that announcements contain significant information about the direction and magnitude of future price changes, which can increase a rational customer's expected utility by improving his ability to time the ticket purchase to minimize the paid. We present a rigorous empirical approach for evaluating the impact of announcements on customer purchases, which allows us to control for potential endogeneity of announcement changes and sales. We identify a novel instrument for inventory announcements, and use a 2SLS model to correct for endogeneity in our regression analysis. We find a positive immediate sales response to announcements for two out of our three carriers, consistent with rational consumer behavior given the information content of announcements.

There are a number of questions related to this topic that remain unexplored here. We have not yet addressed the central inventory announcement question faced by the firm: namely, what is the optimal announcement policy? Our results here show that inventory announcements can benefit customers by reducing the expected price paid for a flight, which implies that the firm earns less on some ticket sales as a result of announcements. On the other hand, if announcements encourage purchases by customers who would not have patronized the firm otherwise, or if they can also shift demand from lower to higher prices, then the revenue impact could be positive. The overall impact of announcements will depend on the fraction of the market that is strategic, and in non-monopoly markets, on the competitive implications of revealing this information. An analysis of the firm's announcement policy decision, using a structural model to allow for evaluation of counterfactual messaging policies, is a natural next step for this research. There are many promising avenues for future empirical work in this area, which have the potential to yield valuable insights for many different industries.

Figure 4.1: Price Change Probabilities by Days-to-Departure, United Airlines

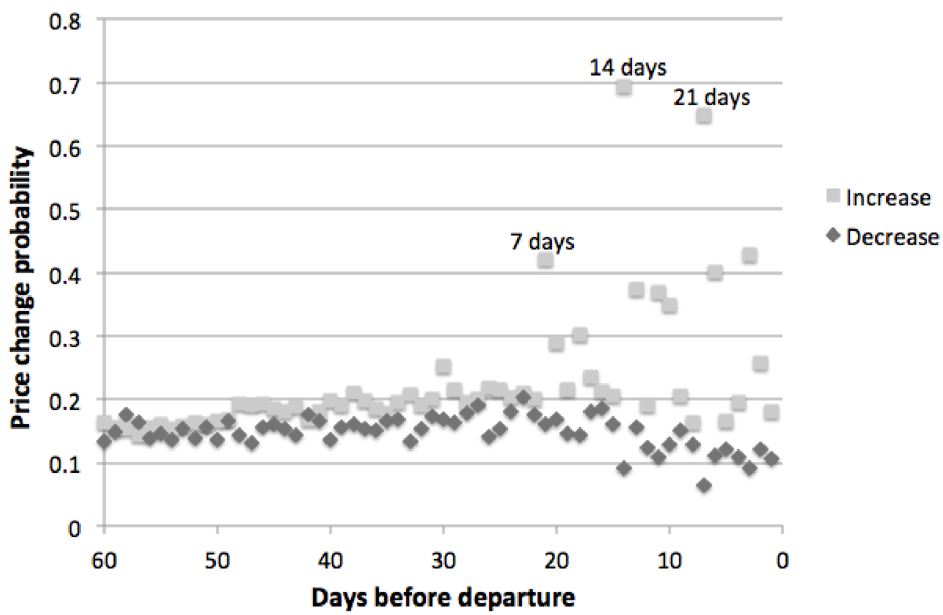

Figure 4.2: Change in Odds of Price Increase and Price Decrease, After Announcement

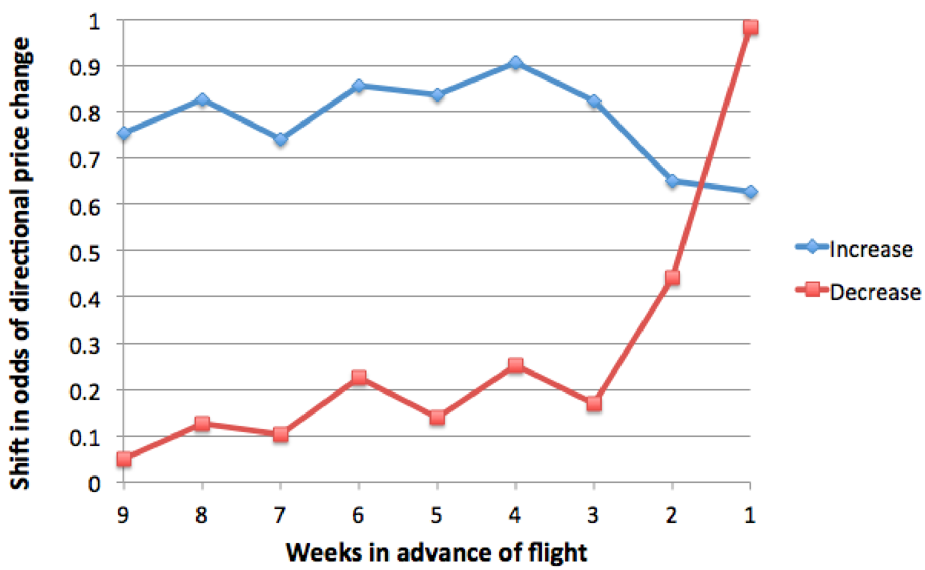

Figure 4.3: Average Price Change Amount and Percentage, by Advance Booking Week

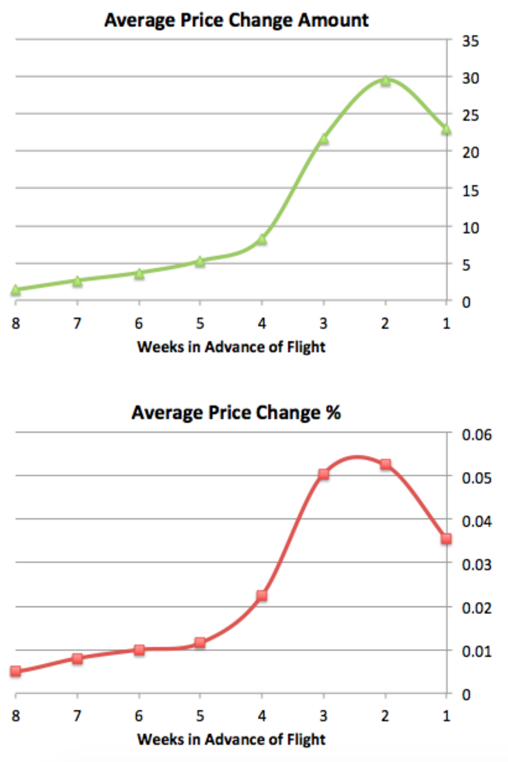

Figure 4.4: OLS Coefficient Estimates: Regression of Price Change Amount on Announcement-Level Indicators and Covariates Listed in Table 4.4, with 95% Confidence Intervals

Conditional on Price Increase

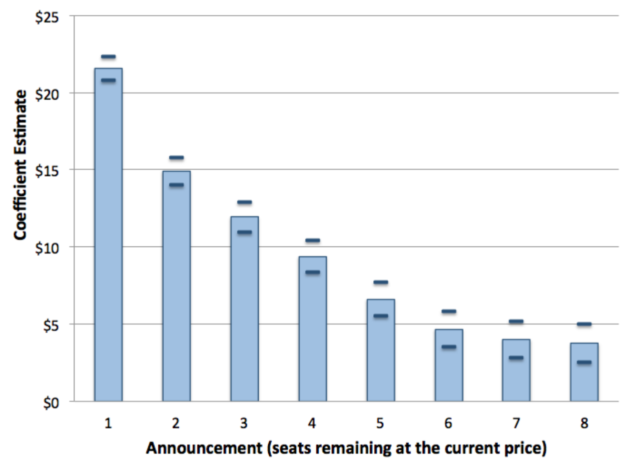

Conditional on Price Decrease

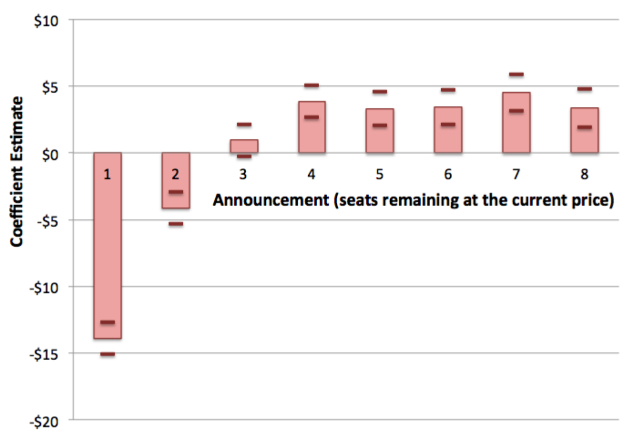

Table 4.1: Summary Statistics: Price Changes by Airline

1-period directional price change probabilities, by announcement status

	American		**Delta**		**United**		
	No announce	Announce	No announce	Announce	No announce	Announce	Announce ≤3
Increase	20.3%	40.0%	14.8%	40.3%	21.6%	32.9%	40.5%
Decrease	12.6%	20.6%	8.9%	14.3%	15.3%	16.6%	16.2%
No change	67.2%	39.3%	76.2%	45.4%	63.1%	50.4%	43.3%

Average 1-period price change magnitudes, by itinerary announcement status

	American		**Delta**		**United**		
	No announce	Announce	No announce	Announce	No announce	Announce	Announce ≤3
All	$5.20	$14.95	$5.91	$27.07	$5.86	$14.38	$19.96
Increases	$52.54	$73.66	$71.75	$88.96	$60.46	$70.29	$73.91
Decreases	-$43.86	-$73.54	-$53.19	-$62.00	-$46.88	-$52.77	-$61.56

Table 4.2: Multinomial Logit Results:

Regression of Directional Price Change on Announcement

	United	
	any announce	announce≤3
No change (base outcome)		
Price decrease		
announce coeff.	0.256***	0.314***
	(.005)	(.007)
announce RRR	1.293	1.369
Price increase		
announce coeff.	0.573***	0.855***
	(.004)	(.005)
announce RRR	1.774	2.360
Increase v. decrease		
RRR	1.372	1.724
Pseudo R-squared	0.050	0.055
n	1,719,100	1,719,100

RRR = Relative Risk Ratio

Additional regressors: cubic time trend; 7-, 14-, and 21-day indicators; market, departure date, and length-of-stay fixed effects

Regression results for American Airlines and Delta Air Lines are qualitatively similar, and are available from the authors.

$^{*}p < 0.05$ $^{**}p < 0.01$ $^{***}p < 0.001$

Table 4.3: Change in Relative Odds of Price Increase vs. Price Decrease, After Announcement

United

Weeks until departure	**Fractional change in odds of increase vs. decrease, after announcement**
1	-18%
2	14%
3	56%
4	52%
5	61%
6	51%
7	58%
8	62%
9	67%
10	55%

Table 4.4: OLS Regression Results:

Regression of Price Change Magnitude on Announcement

United

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
announcement	\$10.52***	\$12.97***	2.1%***	2.0%***	1: \$24.00***	\$11.22***	\$16.37***
					2: \$15.65***		
					3: \$10.67***		
					4: \$8.05***		
					5: \$6.38***		
					6: \$5.22***		
					7: \$4.79***		
					8: \$4.05***		
lnprice (1)-(5); price (6)-(7)	-48.63***	-23.71***	-11.0%***	-20.4%***	-49.43***	-0.08***	-0.08***
advweek	-23.53***	-34.18***	-4.1%***	-5.8%***	-23.71***	-37.05**	-37.31**
advweek^2	4.04***	5.53***	0.6%***	0.7%***	4.07***	5.88***	5.92***
advweek^3	-0.23***	-0.30***	-0.1%***	-0.1%***	-0.23***	-0.31***	-0.31***
7-day indicator	74.20***	57.32***	11.9%***	9.2%***	74.16***	73.83***	73.90***
14-day indicator	52.46***	40.81***	11.6%***	9.1%***	52.36***	52.87***	52.96***
21-day indicator	10.12***	5.10***	2.6%***	1.0%***	9.95***	10.52***	10.50***
price change variable	dollar amount	dollar amount	% change	% change	dollar amount	dollar amount	dollar amount
price changes	all	all	all	increase only	all	all	all
announcement variable	any announce	announce≤3	any announce	any announce	level indicators	any announce	announce≤3
Adjusted R-squared	0.091	0.181	0.113	0.206	0.095	0.096	0.098
n	1,719,100		1,719,100	431,490	1,719,100	1,719,100	

1-8 in the first row of specification (5) indicate the announcement-level indicators All regressions include market, departure date, and length-of-stay fixed effects

$^{*}p < 0.05$ $\quad ^{**}p < 0.01$ $\quad ^{***}p < 0.001$

Table 4.5: Expected Price Changes, by Advance Booking Week and Announcement Status

Weeks in advance of flight	**Expected price change (p_t to p_{t+1})**		
	Unconditional	**After announce**	**After no announce**
9	-2.32	**3.85**	-3.47
8	0.60	6.84	**-0.98**
7	0.63	8.19	**-1.33**
6	1.66	9.36	**-0.64**
5	2.19	1.86	2.63
4	3.04	14.61	**-0.66**
3	3.99	17.77	**-1.39**
2	13.25	37.08	3.58
1	22.07	45.14	11.78
0	19.96	33.75	11.56

Bold text indicates that the customer's optimal decision differs under the announcement policy (for the relevant message status, either announcement or no announcement), compared to the case where no inventory information is provided.

Table 4.6: OLS Regression Results:

Regression of Sales on Announcement Changes

	United			**American**			**Delta**		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
announceDecrease	0.208***	0.215***	0.175***	0.010	0.011	0.021	-0.092*	-0.092*	-0.088*
	(0.039)	(0.036)	(0.037)	(0.030)	(0.030)	(0.031)	(0.042)	(0.040)	(0.038)
lnprice	-0.362***	-0.372***	-0.450***	-0.484***	-0.483***	-0.459***	-0.605	-0.606***	-0.612***
market announcements	-	-0.117	-0.184***	-	-0.008	-0.005	-	-0.004	0.003
load factor	-	-	0.636***	-	-	-0.059**	-	-	0.597*
advweek	0.397***	0.396***	0.407***	0.081**	0.081***	0.082***	0.809***	0.798***	0.774***
advweek^2	-0.114***	-0.114***	-0.109***	-0.037***	-0.037***	-0.037***	-0.307***	-0.305***	-0.285***
advweek^3	0.007***	0.007***	0.007***	0.003***	0.003***	0.003***	0.025***	0.025***	0.023***
7-day indicator	0.411***	0.411***	0.397***	0.139***	0.139***	0.152***	0.566***	0.568***	0.631***
14-day indicator	0.307***	0.309***	0.292***	0.072***	0.072***	0.076***	0.426***	0.435***	0.424***
21-day indicator	0.082***	0.086***	0.076***	0.028*	0.028*	0.030*	0.209***	0.220***	0.205***
R-squared	0.085	0.085	0.091	0.014	0.014	0.014	0.136	0.136	0.147
n		508,840			709,684			1,095,031	

All regressions include market, departure date, and length-of-stay fixed effects.

$^{*}p < 0.05$ $^{**}p < 0.01$ $^{***}p < 0.001$

Table 4.7: 2SLS Regression Results

	United	**American**	**Delta**
announceDecrease	0.495***	0.313*	-0.264
	(0.184)	(0.149)	(0.354)
lnprice	-0.507***	-0.396***	-0.599***
R-squared	0.034	0.0146	0.135
n	508,840	709,684	1,095,031

Additional regressors: cubic time trend; 7-, 14-, and 21-day indicators; market, departure date, and length-of-stay fixed effects.

$^{*}p < 0.05$ $^{**}p < 0.01$ $^{***}p < 0.001$

Chapter 5

Conclusion

In this dissertation, I study two revenue management tools that are widely used by service-industry firms: reservation no-show fees and inventory announcements. I use both analytical and empirical techniques to develop managerial insights for firms that may consider using these mechanisms in their practice of revenue management.

In my analysis of no-show fees, I characterize the fee policy that emerges in equilibrium and show that despite consumers' aversion to fee-based reservation policies, both customers and firms can benefit when the firm uses a no-show fee pricing scheme. The no-show fee policy results in a more efficient allocation of scarce capacity, and serves as a partial antidote to strategic behavior.

In my analysis of inventory announcements, I show that inventory announcements in the airline industry contain significant information about price changes, and have an impact on customer purchases. I find that both price increases and price decreases are more likely after an inventory announcement, and although the relative probability of increase is higher after an announcement in most cases, there are times in the booking horizon (close to departure) when an announcement can signal a higher likelihood of price decrease. Additionally, I find a positive immediate sales response to announcements for two out of the three carriers in my study, consistent with rational consumer behavior given the information content of announcements.

Bibliography

Aggarwal, P., S.Y. Jun, and J.H. Huh (2011). “Scarcity messages”. *Journal of Advertising* 40.3, pp. 19–30.

Alexandrov, A. and M.A. Lariviere (2012). “Are reservations recommended?” *Manufacturing & Service Operations Management* 14.2, 218–230.

Allon, G. and A. Bassamboo (2011). “Buying from the babbling retailer? The impact of availability information on customer behavior”. *Management Science* 57.4, 713–726.

Allon, G., A. Bassamboo, and I. Gurvich (2012). ““We will be right with you”: Managing customer expectations with vague promises and cheap talk”. *Operations Research* 59.6, 1382–1394.

Angrist, J. and A.B. Krueger (1991). “Does compulsory school attendance affect cchooling and earnings?” *Quarterly Journal of Economics* 106.4, 979–1014.

Angrist, J. and J.-S. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton, NJ: Princeton University Press.

Aviv, Y. and A. Pazgal (2008). “Optimal pricing of seasonal products in the presence of forward-looking consumers”. *Manufacturing & Service Operations Management* 10.2, pp. 339–359.

Bertsimas, D. and R. Shioda (2003). “Restaurant Revenue Management”. *Operations Research* 51.3, 472–486.

Besanko, D. and W.L. Winston (1990). “Optimal price skimming by a monopolist facing rational consumers”. *Management Science* 36.5, pp. 555–567.

Byun, S.E. and B. Sternquist (2012). “Here today, gone tomorrow: Consumer reactions to perceived limited availability”. *Journal of Marketing Theory and Practice* 20.2, pp. 223–234.

Cachon, G. and P. Feldman (2011). “Pricing services subject to congestion: Charge per-use fees or sell subscriptions”. *Manufacturing & Service Operations Management* 13.2, pp. 244–260.

— (2015). “Price commitments with strategic consumers: Why it can be optimal to discount more frequently...than optimal”. *Manufacturing & Service Operations Management* 17.3, pp. 399–410.

— (2016). “Is advance selling desirable with competition?” *Marketing Science*. Forthcoming.

Cachon, G. and R. Swinney (2009). “Purchasing, pricing, and quick response in the presence of strategic consumers”. *Management Science* 55.3, pp. 497–511.

Cachon, G. and C. Terwiesch (2011). *Matching Supply with Demand: An Introduction to Operations Management*. 3rd ed. McGraw Hill.

Çil, E.B. and M.A. Lariviere (2013). "Saving seats for strategic consumers". *Operations Research* 61.6, pp. 1321–1332.

Crawford, V. and J. Sobel (1982). "Strategic information transmission". *Econometrica* 50.6, pp. 1431–1451.

DeGraba, P. (1995). "Buying frenzies and seller-induced excess-demand". *RANDJournal of Economics* 26.2, pp. 331–342.

Erdem, T. and M. Keane (1996). "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets". *Marketing Science* 15.1, pp. 1–20.

FlyerTalk (2015). `http://www.flyertalk.com/forum/`.

Gale, I.L. and T.J Holmes (1993). "Advance-purchase discounts and monopoly allocation of capacity". *The American Economic Review* 83.1, pp. 135–146.

Gallego, G. and O. Sahin (2010). "Revenue management with partially refundable fares". *Operations Research* 58.4, pp. 817–833.

Georgiadis, G. and C.S. Tang (2014). "Optimal reservation policies and market segmentation". *International Journal of Production Economics* 154, pp. 81–99.

Guillet, B.D. and I. Mohammed (2015). "Revenue management research in hospitality and tourism: A critical review of current literature and suggestions for future research". *International Journal of Contemporary Hospitality Management* 27.4, pp. 526–560.

Hendel, I. and A. Nevo (2006). "Measuring the implications of sales and consumer inventory behavior". *Econometrica* 74.6, pp. 1637–1673.

— (2013). "Intertemporal price discrimination in storable goods markets". *American Economic Review* 103.7, pp. 2722–51.

Jerath, K., S. Netessine, and S.K. Veeraraghavan (2010). "Revenue management with strategic customers: Last minute selling and opaque selling". *Management Science* 56.3, pp. 430–448.

Kimes, S.E. et al. (1998). "Restaurant revenue management: Applying yield management to the restaurant industry". *Cornell Hotel and Restaurant Administration Quarterly* 39.3, pp. 32–39.

Levin, Y., J. McGill, and M. Nediak (2009). "Dynamic pricing in the presence of strategic consumers and oligopolistic competition". *Management Science* 55.1, pp. 32–46.

Li, J., N. Granados, and S. Netessine (2014). "Are consumers strategic? Structural estimation from the air-travel industry". *Management Science* 60.9, pp. 2114–2137.

Littlewood, K. (1972). "Forecasting and control of passenger bookings". *Proceedings of the Twelfth Annual AGIFORS Symposium*. Nathanya, Israel.

Liu, Q. and G.J. van Ryzin (2008). "Strategic capacity rationing to induce early purchases". *Management Science* 54.6, pp. 1115–1131.

Lynn, M. (1992). "Scarcity's enhancement of desirability: The role of naive economic theories". *Basic and Applied Social Psychology* 13.1, pp. 67–78.

Nair, H. (2007). "Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games". *Quantitative Marketing and Economics* 5.3, pp. 239–292.

Oh, J. and X. Su (2012). "Pricing restaurant reservations: Dealing with no-shows". Working paper.

Osadchiy, N. and E. Bendoly (2010). "Are consumers really strategic? Implications from an experimental study.", Working paper, Emory University.

Osadchiy, N. and G. Vulcano (2010). "Selling with binding reservations in the presence of strategic consumers". *Management Science* 56.12, 2173–2190.

Png, I.P.L. (1989). "Reservations: Customer Insurance in the Marketing of Capacity". *Marketing Science* 8.3, 248–264.

PWC (2015). *2015 Aviation Trends*. `http://www.strategyand.pwc.com/perspectives/2015-aviation-trends`.

Shen, Z.-J. and X. Su (2007). "Customer behavior modeling in revenue management and auctions: A review and new research opportunities". *Production and Operations Management* 16.6, pp. 713–728.

Shugan, S.M. and J. Xie (2000). "Advance Pricing of Services and Other Implications of Separating Purchase and Consumption". *Journal of Service Research* 2.3, pp. 227–239.

Su, X. (2007). "Intertemporal pricing with strategic customer behavior". *Management Science* 53.5, pp. 726–741.

Su, X. and F. Zhang (2008). "Strategic customer behavior, commitment, and supply chain performance". *Management Science* 54.10, pp. 1759–1773.

Swinney, R. (2011). "Selling to strategic consumers when product value Is uncertain: The value of matching supply and demand". *Management Science* 57.10, pp. 1737–1751.

Szkaradnik, C. (2014). *Volver abolishes tickets, adjusts pricing.* `http://philly.eater.com/2014/9/19/6573595/volver-abolishes-tickets-adjusts-pricing`.

Talluri, K.T. and G.J. van Ryzin (2004). *The Theory and Practice of Revenue Management.* Boston, MA: Kluwer Academic Publishers.

Tereyağoğlu, N. and S. Veeraraghavan (2012). "Selling to conspicuous consumers: Pricing, production, and sourcing decisions". *Management Science* 58.12, pp. 2168–2189.

Wells, P. (2015). *Restaurants add reservation cancellation fees to the menu.* `http://nyti.ms/1dcOBYg`.

Williams, K.R. (2013). "Dynamic airline pricing and seat availability." Working paper, University of Minnesota.

Xie, J. and S.M. Shugan (2001). "Electronic tickets, smart cards, and Online Prepayments: When and how to advance sell". *Marketing Science* 20.3, pp. 219–243.

Yu, Q., G. Allon, and A. Bassamboo (2016). "How do delay announcements shape customer behavior? An Empirical Study". *Management Science,* Forthcoming.