

DISCUSSION PAPER SERIES

No. 7315

**A DYNAMIC MODEL OF PRICE
DISCRIMINATION AND INVENTORY
MANAGEMENT AT THE FULTON
FISH MARKET**

Kathryn Graddy and George Hall

INDUSTRIAL ORGANIZATION



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP7315.asp

A DYNAMIC MODEL OF PRICE DISCRIMINATION AND INVENTORY MANAGEMENT AT THE FULTON FISH MARKET

Kathryn Graddy, Brandeis University and CEPR
George Hall, Brandeis University

Discussion Paper No. 7315
June 2009

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Kathryn Graddy and George Hall

CEPR Discussion Paper No. 7315

June 2009

ABSTRACT

A Dynamic Model of Price Discrimination and Inventory Management at the Fulton Fish Market*

We estimate a dynamic profit-maximization model of a fish wholesaler who can observe consumer characteristics, set individual prices, and thus engage in third-degree price discrimination. Simulated prices and quantities from the model exhibit the key features observed in a set of high quality transaction-level data on fish sales collected at the Fulton fish market. The model's predictions are then compared to the case in which the dealer must post a single price to all customers. We find the cost to the dealer of posting a uniform price to be extremely small.

JEL Classification: C15, D21, D4, L1 and L81

Keywords: dynamic programming, fish, indirect inference, price discrimination and yield management

Kathryn Graddy
Department of Economics
Brandeis University
415 South Street
Waltham, MA 02454-9110
USA
Email: kgraddy@brandeis.edu

George Hall
Department of Economics
Brandeis University
415 South Street
Waltham, MA 02454-9110
USA
Email: ghall@brandeis.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=135475

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=151940

* We thank participants at the Brandeis IBS brown-bag seminar for helpful comments.

Submitted 21 May 2009

1 Introduction

Which type of selling method can achieve the highest revenue for the seller? Common examples of different types of selling methods are posted prices, privately negotiated prices, and auctions. This paper looks specifically at the revenue advantage of price discrimination resulting from privately negotiated prices in the context of a fish market.

We solve and estimate a dynamic profit-maximization model of a fish wholesaler. Stocks of fish arrive every morning and the fish must be sold within a relatively short period of time. Throughout the day, customers arrive sequentially and randomly, but when a customer shows up, the seller observes his type and thus knows his price elasticity. Therefore the seller can price discriminate across different types of buyers. The fish seller is solving two problems simultaneously: 1) how to optimally price his stock of fish which is falling in value over time and is replenished only once a day; and 2) how to optimally price discriminate across customers with differing price elasticities.

Our model is able to successfully match several key features of the Fulton fish market. In particular the model predicts, as we see in the data, that Asian customers pay about 5 cents per pound less for fish than white customers. More generally the model matches the mean and variance of prices with large differences in prices across different days as well as considerable intra-day price volatility. Overall, we conclude that the model is a reasonable approximation of the behavior of a Fulton fish market dealer.

We then impose the restriction that the fish seller must post a uniform price to all customers, and is therefore unable to price discriminate. In this case, the inelastic white customers pay lower prices, purchase larger quantities of fish, and make more frequent purchases. The more price-elastic Asian customers pay higher prices and purchase less fish. The revenue gain from the larger, more frequent, white purchases essentially offsets the loss in revenue from the smaller Asian purchases. Consequently, in this dataset and with this model, we find only a very small cost to posting a single price.

While our model is intended to provide an accurate representation of the behavior of a particular fish wholesaler, the problem of determining the optimal price of a stock of depreciating assets is a classic problem in economics and operations research. In operations research the study of dynamically pricing an inventory stock falls under the headings *revenue management* or *yield management*.¹ In the economics literature, work by Reagan (1982), Aguirregabiria (1999), Zettelmeyer, Scott Morton and Silva-Risso (2003), Chan, Hall, and Rust (2004), Sweeting (2008), and Copeland, Dunn, and Hall (2009) study the interaction

¹This literature which started with Whiten (1955) and Karlin and Carr (1962) is reviewed by Federgruen and Heching (1999) and Elmaghraby and Keskinocak (2003).

between inventory management and pricing. Our model differs from these others in that the timing of new procurements of inventories is fixed though the quantity is stochastic. While these other papers have focused on durable goods, such as automobiles and steel, or goods that expire at a pre-determined date, such as baseball tickets, hotel rooms and airline seats, our paper focuses on fish – a good that depreciates steadily but quickly.

This paper proceeds as follows. In the next section we discuss the importance of price discrimination and previous papers that have highlighted discriminatory settings. In section 3, we discuss the details of the Fulton fish market, the market on which our structural model is based. In section 4 we develop our model, and in section 5, we present our strategy for estimating the model. We report our findings in section 6 and conclude our analysis in section 7.

2 Price Discrimination versus a Single Posted Price

In an early empirical paper on price discrimination, Borenstein (1991) demonstrated price discrimination in the retail gasoline market, a market that at first glance appeared to be more competitive than monopolistic (and in this way similar to the Fulton fish market) by estimating price-cost margins. More recent work has focused on demonstrating price discrimination through structural models. For example, Hastings (2009) focuses on wholesale price discrimination in the gasoline market by using wholesale transaction prices, retail prices and quantities. Hastings specifically looks at the question of whether eliminating wholesale price discrimination would improve welfare by constructing and estimating a structural model. Villas-Boas (2009) also looks at the question of whether eliminating wholesale price discrimination would improve welfare, but does it in the context of the yogurt market and does this without having access to wholesale prices. The paper that is closest in methodology to this paper is Chan, Hall and Rust (2004), in which the authors compare the case of price discrimination in the steel market by a wholesaler to the case of uniform pricing by estimating a dynamic model of price discrimination and inventory adjustment.

Understanding the advantages of uniform pricing – often imposed by posted prices – over price discrimination is an important endeavor. With advances in information technology, more markets are going to posted prices or public auctions. For example, the secondary market in Treasury securities used to be a market with privately negotiated prices, but now much of the market is a posted-price market. As Rust and Hall (2003) argue, posting a fixed price eliminates the need for consumers to engage in costly search and thus may increase demand from consumers with high search costs. Yet, some markets, such as the

market for buying a car, essentially remain a market with negotiated prices. Negotiated prices allow third-degree price discrimination – where different groups of customers are charged a different price based on an observable characteristic. More elastic consumers are charged lower prices, and less elastic consumers are charged higher prices. This can lead to “unfair” outcomes, such as when Ayres and Siegelman (1995) showed that blacks and women paid higher prices for cars at dealerships.

Why do some markets remain markets with negotiated prices and others have posted prices? Wholesale markets, and markets with large demand and supply shocks, such as steel, automobiles, secondary security markets, and fish markets often have negotiated prices, whereas most items that are purchased on a regular basis have posted prices. Negotiated prices allow the ability to price discriminate. Everything else equal, if a firm has the ability to price discriminate, then the firm can raise its revenue. However, everything else is very rarely equal. The ability to price discriminate raises negotiation costs for the seller. It also raises search costs for the buyers and psychic costs to being treated unfairly. Through posted prices, firms can often raise quantity demanded. The extent to which a market remains privately negotiated or goes to posted prices largely rests on whether the often hard-to-measure costs of price discrimination outweigh the measurable revenue benefits. Even if a part of the market remains a privately negotiated market, markets that have partly migrated to the internet can lose their ability to price discriminate. Scott Morton, Zettelmeyer, and Silva-Risso (2003) find that the groups who pay the highest prices under negotiated prices benefit the most from buying cars off the internet.

3 The Fulton Fish Market

In this section, we provide details of the Fulton fish market. Through a contact at the Fulton fish market in New York City, one of the authors, Kathryn Graddy, collected a detailed set of transaction level data for a single fish dealer for 22 weeks from December 2, 1991 to May 8, 1992 (111 business days and 2,868 transactions) for one type of fish: whiting. For each transaction, she collected the order the transaction occurred, the price paid, the quantity sold, the quality of the fish, and several characteristics of the buyer (e.g. race, location and type of retail outlet). For each day in the sample, Graddy also collected the quantity of whiting received and discarded by this dealer and the weather. These data have been described and extensively analyzed in Graddy (1995), Angrist, Graddy, and Imbens (2000), Graddy (2006), Lee (2007), and Graddy and Kennedy (2009). Given that these data are high quality and that the fish seller’s problem is relatively straight forward and well understood, we wish to formulate a structural model that can

replicate the main features of these data. We have attempted to make our model an accurate representation of a Fulton fish market dealer focusing on the following details of this market.

At the time the data were collected in the early 1990s, the market was open for business from 3:00 to 9:00 in the morning on Monday and Thursday and from 4:00 to 9:00 on Tuesday, Wednesday, and Friday. There were 60 registered dealers in the market, but only 35 wholesalers actively operating in the market. Not all dealers carried all types of fish. Only six major dealers sold whiting.²

The whiting was supplied by fishermen's cooperatives, packing houses, and smaller fishing boats in New Jersey, Long Island, and Connecticut each day before the market opens. The price the dealer paid a supplier for a particular day's supply of whiting was determined at the end of the day by the prices the dealer received for the whiting on that particular day. Each supplier serviced multiple dealers. A supplier would refuse to deliver fish to a dealer from whom he continually received a poor price relative to other dealers. In this way, dealers competed with each other for fish supplied in the future by the amount they paid the suppliers for fish on a particular day. The dealer stated that he kept 5-15 cents on each pound of whiting sold (although this margin could not be independently verified).

A dealer often had a good idea of the quantity that would be available before the market's close on the previous day. Quantity supplied was primarily determined by weather conditions with wind and waves being the most important determinant of the quantity of fish caught. Days were classified as stormy when wave heights were greater than 4.5 feet and wind speed is greater than 18 knots. Wind speed and wave height are moving averages of the last three days' wind speed and wave height before the trading day, as measured off the coast of Long Island and reported in the *New York Times* boating forecast. Quantities fall and prices rise when there are storms, and quantities rise and prices fall in good weather. Holding day of the week constant, the average quantity sold on a clear day was 2,371 pounds more than on a stormy day. Conversely, the average price was 32 cents less per pound than on a clear day.

The quantity of whiting delivered for a particular day was received in its entirety before the market opened. If one dealer did not receive what he perceived to be enough whiting, another dealer would sell him whiting before the market opened. The price was set not at the time of sale, but was determined toward the end of the trading day.

Anyone could purchase fish at the Fulton fish market, but small quantities were not sold. For whiting, the minimum quantity sold was one box, approximately 60 pounds (except as favors to regular customers).

²Data on whiting were collected for three reasons. More transactions take place in whiting than in any other fish. Whiting do not vary substantially in size and quality. Finally, and perhaps most importantly, the dealer from whom the data were collected suggested that the whiting salesman would be the only salesman amenable to having an observer.

Most of the customers were repeat buyers, but buyer-seller relationships varied. Some buyers regularly purchased from one seller, and other buyers purchased from various sellers. Most buyers followed distinct purchasing patterns. Monday, Thursday and Friday were big days, while Tuesday and Wednesday were relatively quiet. Whiting stayed sufficiently fresh to sell for at most four days after it was received. However, it was usually sold either the day that it was received or the day after. About half of the time, the total quantity received exceeded total sales. As would be expected, there was a small amount of oversupply as some inventory loss is inevitable in the selling process, especially given the perishable nature of fish. Over 111 days in late 1991 and early 1992, the total quantity received exceeded the total quantity sold by 11,237 pounds, which amounts to about 1.6 percent of total sales. The bulk of this difference consisted of fish that was not sufficiently fresh to sell and was literally thrown out.

The Fulton fish market had no posted prices and each dealer was free to charge a different price to each customer. If a customer wished to buy a particular quantity of fish, he would ask a seller for the price. The seller would quote a price and the customer would usually either buy the fish or walk away. Prices were generally quoted in five-cent increments. Sellers were discreet when naming a price. A particular price was for a particular customer.

Prices varied tremendously. For example, the average price in a whiting transaction on Friday, May 1, 1992, was \$.33/lb., and the average price on Friday, May 8, was \$1.75/lb., an increase of almost six-fold. Figure 1 is a high-low price chart, where the bottom of the vertical line shows the low price for that day and the top of the vertical line shows the high price for that day. The high intraday volatility was primarily due to changes in price throughout the day, but some of the volatility could be attributed to differences in qualities of fish. There was often a large decline in price after 7:00 a.m, but the pattern before 7:00 a.m. is not easily predictable. There was a higher intraday volatility on Friday than on other days.

Figure 2 indicates total daily sales, which varied considerably. Daily supply shocks, primarily caused by weather, were largely responsible for the high volatility in day-to-day prices. The simple correlation between total daily sales and average daily price is $-.32$.

About 60 percent of the purchasers at the market were Asian. Graddy observed that sellers were quoting lower prices to Asian customers for the same box of fish. The price difference would usually be about 5 cents, though it could go to as large as 10 cents. Regressions that carefully control for other variables, including time of sale and quality, indicated that whites pay on average 6.3 cents per pound more on each transaction for the same type and quality of fish than do Asian buyers (Graddy, 1995). While quality controls could naturally be imperfect, these regressions were consistent with observations.




Figure 1: Day-to-Day Prices for Whiting (in \$ per pound)

These price differences were able to persist due to the different characteristics of the buyers. There was strong anecdotal evidence that Asian buyers were more elastic than white buyers. Asian buyers would resell their whiting as whole fish in retail shops, fish sandwiches in fry shops, and would make them into fish balls. Most of these establishments were located in very poor neighborhoods. Resellers had little scope to raise the price of whiting for their ultimate customers. Consequently, they would bargain very hard at the market.

In contrast, many white buyers (though not all) had more scope to pass on prices to customers. For example, for the local fish dealer in Princeton, New Jersey, whiting was a very small part of his purchase; he would spend no time shopping around for price on whiting. If prices were especially high at the market, he could explain to his customers that fish were expensive at Fulton Street and his customers would often be willing to absorb costs by paying higher prices. Finally since prices were quoted discretely and white buyers rarely socialized with Asian buyers it is unlikely that either white buyers or Asian buyers knew that each group was being charged different prices.




Figure 2: Day-to-Day Volumes of Whiting (in pounds)

Given the small number of dealers (six at the time) who carried whiting and the long tenure of many of the dealers, the structure of the Fulton fish market for whiting, at least in the early 1990s, appeared conducive to tacit collusion, rather than perfect competition. Dealers traded with each other before the market opened. They would price fish on a daily and sometimes hourly basis and would receive feedback from the buyers on other dealers' prices. Dealers in such a situation could easily become very good at tacit communication. Moreover, entry into this market was limited because suppliers provided the whiting on credit to established dealers - and it is not clear they would be willing to provide fish on credit to a new entrant. The presence of organized crime in the market may also have discouraged entry. Hence there is little reason to expect the extraordinary profits earned from price discrimination to be competed away.

4 The Model

In this section we present our dynamic, profit maximization model of a fish dealer. The model is designed to be consistent with the structure of the Fulton fish market and to capture the basic features of the price

and quantity data presented above. In this spirit, there are shocks to both supply and demand. The daily supply of fish is stochastic and depends on the weather with storms reducing the average size of the catch. On the demand side, the arrival of customers each day is uncertain and their individual demand curves vary as well. Fish can be stored but depreciates quickly. As customers arrive sequentially, the fish dealer chooses a price for each individual customer. If he sets too high a price, he risks having fish go bad; if he sets too low a price, he foregoes the potential to sell the fish to subsequent customers at a higher price. The dealer can price discriminate since he observes each customer's type before quoting a price. The optimal pricing strategy is determined by recursively solving his dynamic optimization problem.

Formally, our fictional fish seller is open for business five days a week, Monday through Friday, for five hours each day. At the start of each workday, he receives a catch, C , which varies by the day of the week and the weather. The weather can be in one of two states, calm or stormy, and follows a Markov chain. The day's catch is drawn from a truncated log-normal distribution with a mean $\mu_c(d, w)$ and a standard deviation $\sigma_c(d, w)$ with $d = 1, 2, \dots, 5$ denoting the day of the week (Monday-Friday) and w denoting the day's weather, as being either calm or stormy.

Each five-hour workday is divided into 60 five minute periods indexed by t . At the start of each period t , the fish dealer holds inventory I_t . Customers arrive sequentially. There are two types of customers, Asians and whites, indexed by n , who arrive with probability λ^n , with

$$\lambda^0 + \lambda^{asian} + \lambda^{white} = 1 \quad (1)$$

where λ^0 is the probability that no customer shows up. Only a single customer may show up during a given period. In period t , if a customer arrives, the fish seller observes the customer's type. He knows each customer of type n has a stochastic constant elasticity demand curve for fish,

$$\log(q_t) = \mu_q^n + \alpha(d) - \gamma^n \log(p_t) + \varepsilon_t^n \quad (2)$$

where ε_t^n is distributed i.i.d $N(0, \sigma_{\varepsilon^n})$. As in Graddy (1995) and Angrist, Graddy and Imbens (2000), we do not allow for cross-price elasticities. Whiting is a good deal cheaper than any other fish, and within the observed range of prices, there did not exist good substitutes. The term $\alpha(d)$ allows the constant term to vary by day of the week but stays fixed with the different types of customers. We normalize $\alpha(5)$ (i.e. Friday) to zero. Sales, q_t , are constrained to

$$60 \leq q_t \leq I_t, \quad (3)$$

where I_t is the stock of fish available for sale at the start of period t . This constraint rules out back-orders and implements the observed minimum purchase size of 60 pounds (one box). If at the quoted price, the customer demands fewer than 60 pounds, no sale occurs. Thus the one-box minimum pins down the customer's reservation price above which he will walk away. If the customer demands more than the current level of inventory, we set $q_t = I_t$, and the dealer is stocked out until the following morning.

Fish depreciates at a constant rate throughout the day, so the evolution of the dealer's stock follows:

$$I_{t+1} = (1 - \tilde{\delta})I_t - q_t \quad \text{for } t = 1, 2, \dots, 59 \quad (4)$$

where $\tilde{\delta}$ is the intra-day depreciation rate. Between weekdays, the dealer's stock is augmented by the catch, so the law of motion becomes

$$I_1 = (1 - \hat{\delta})I_{60} - q_{60} + C(d', w') \quad \text{for } d' = 2, 3, 4, 5 \quad (5)$$

with $\hat{\delta}$ denoting the overnight depreciation rate. Over the weekend, this law of motion becomes

$$I_1 = (1 - \bar{\delta})I_{60} - q_{60} + C(1, w') \quad (6)$$

with $\bar{\delta}$ denoting the over-the-weekend depreciation rate.

We assume the fish dealer's single-period compensation function, π , is a non-linear function of the total quantity of fish sold and total revenue:

$$\pi(p_t, q_t) = \phi_0 q_t + \phi_1 (p_t q_t)^{\phi_2}. \quad (7)$$

The first term is a fixed commission for each fish sold and the second term provides the dealer an incentive to maximize total revenue. This compensation function is consistent with the observation that the price paid by the dealer to the fisherman (or to another dealer) is determined at the end of the day, after observing total revenue. More generally, this function is consistent with a contract between two parties in which one party (the supplier) wishes to induce effort and share risk with the other (the dealer). This specification allows for revenue maximization (i.e. $\phi_0 = 0, \phi_2 = 1$) and straight commission (i.e. $\phi_1 = 0$) as special cases.

Price Discrimination Case

We first consider the case in which the fish dealer observes the customer's type and is able to set a private price. Specifically, the fish dealer knows his inventory on hand, I , the type of customer he faces,

n , the weather w , the day of the week, d , and the period, t , when he chooses a price, p . We abstract from haggling, assuming the dealer makes a single take-it-or-leave-it price quote. After quoting a price, he learns ε and thus q . Let $V(I, n, w, d, t)$ denote the value to the fish dealer with state I, n, w, d, t . The fish dealer's value function is then:

$$V(I, n, w, d, t) = \max_p \left\{ E \pi(p, q(p)) + E \beta V((1 - \bar{\delta})(I - q(p)), n', w, d, t + 1) \right\} \text{ for } t = 1, 2, \dots, 59 \quad (8)$$

where $\beta \in (0, 1)$ is the intra-day discount factor, and $\pi(p, q)$ is given by (7). The dealer's optimization problem is subject to the demand function (2) and the constraint on sales (3).

Note that other dealers' prices are excluded from the value function. This specification is consistent with the situation in which all dealers are maximizing the same value function. Hence, dealers can be tacitly colluding. This is not an unrealistic assumption as dealers have been making identical pricing decisions, period after period, day after day, and year after year for nearly a hundred years in some cases.³

At the end of Mondays through Thursdays, $\hat{\delta}$, fraction of the inventory spoils, a new catch is drawn, and a new day starts all over again. In the last period of the day (period 60), the fish dealer's Bellman equation is:

$$V(I, n, w, d, 60) = \max_p \left\{ E \pi(p, q(p)) + E \hat{\beta} V((1 - \hat{\delta})(I - q) + C(d + 1, w'), n', w', d + 1, 1) \right\} \quad (9)$$

where $\hat{\beta}$ is the inter-day discount factor. If $d = 5$ (i.e. Friday), the continuation value in equation (9) becomes $E \hat{\beta} V((1 - \bar{\delta})(I - q) + C(1, w'), n', w', 1, 1)$.

The solution to the dealer's problem yields a decision rule for pricing:

$$\tilde{p} = \tilde{p}(I, n, w, d, t). \quad (10)$$

In general an analytical solution for the constraint on sales, \tilde{p} , is not available; but, ignoring the constraint (3) and setting $\phi_0 = 0$, $\phi_2 = 1$ and $\sigma_\varepsilon^n = 0$, the optimal pricing rule is

$$\tilde{p} = \left[\frac{\gamma^n}{\gamma^n - 1} \right] \beta \frac{\partial V}{\partial I'}(I', n', w, d, t + 1). \quad (11)$$

This is the standard monopoly pricing rule, but in this case the firm's marginal cost is the shadow value to the firm of a unit of inventory next period. In other words, the cost to the dealer of selling the marginal box of fish is the foregone opportunity of selling that box next period.

³We do not have data on other dealers' prices. No other dealers would allow an observer to record prices, although this was attempted during the time of data collection.

No-Price Discrimination Counter-Factual

We modify our model to obtain predictions under the assumption that our fictional fish dealer is unable to condition his price on the customer type, n . This case corresponds to a market structure in which dealers receive all requests for price quotes through an anonymous channel such as the internet. The fish dealer can still condition his choice of p on I, w, d, t ; but, only after announcing p does he learn n and ε and thus q .

Let $V(I, w, d, t)$ denote the value to the fish dealer on day d with weather w at time t of having I inventory on hand. The value functions can be written as

$$V(I, w, d, t) = \max_p \left\{ E \pi(p, q(p)) + E \beta V((1 - \tilde{\delta})(I - q(p)), w, d, t + 1) \right\} \quad \text{for } t = 1, 2, \dots, 59 \quad (12)$$

and

$$V(I, w, d, 60) = \max_p \left\{ E \pi(p, q(p)) + E \hat{\beta} V((1 - \hat{\delta})(I - q(p)) + C(d + 1, w'), w', d + 1, 1) \right\} \quad (13)$$

subject to (2) and (3). As in the price discrimination case, if $d = 5$, the continuation value in equation (13) becomes $E \hat{\beta} V((1 - \hat{\delta})(I - q(p)) + C(1, w'), w', 1, 1)$.

The solution to the dealer's problem in the case of no price discrimination yields a decision rule for pricing:

$$\tilde{p} = \tilde{p}(I, w, d, t). \quad (14)$$

5 Estimating the Model

In order to solve the model, we must first select parameter values. For a subset of the parameters we choose values to match observed characteristics of the market. For the remaining parameters we select values that generate simulations from our model as close as possible to the observed time paths of prices and sales. We are more precise about what we mean by "as close as possible" below.

We set the intra-day discount factor β to 0.999. Hence a dollar of revenue at the end of the workday is worth only 94.2 cents (i.e. 0.999^{60}) relative to the start of the workday. Since there are 288 five minute intervals in 24 hours, the inter-day discount factor is set to $\hat{\beta} = \beta^{(288-60)}$ or 0.796. We found that having the fictional dealer significantly discount the future allowed the model to better match the data. We suspect this is due to the model's simplicity with regard to quality. While Graddy (1995) considers five degrees of quality, in our model there are only two: fresh (salable) and spoiled (unsalable). Ideally, the model

would take into account the vintage of each fish and track its quality decline over time; but doing so would make the model untractable. We fix the value for β since estimating a discount factor so close to one is computationally challenging.

We assume fish depreciates at a constant rate throughout the day. Since whiting can be sold for up to four days, we set the daily depreciation rate, δ , to $\frac{1}{4}$. Since we divide the workday into five-minute intervals, we set the intra-workday depreciation rate $\tilde{\delta} = 1/288 \times \delta$, and set the inter-day and over-the-weekend depreciation rates to $\hat{\delta} = 228/288 \times \delta$, and $\bar{\delta} = 67/24 \times \delta$, respectively.

As stated above, each morning's catch is drawn from a log-normal distribution with a mean $\mu_c(d, w)$ and a standard deviation $\sigma_c(d, w)$. We set the ten values of $\mu_c(d, w)$ and the ten values of $\sigma_c(d, w)$ to the observed conditional means and standard deviations of quantity received in Graddy (1995). These values are reported in the appendix. Matching observed weather patterns we set the Markov transition matrix for w to:

$$\chi(w, w') = \begin{bmatrix} 0.85 & 0.15 \\ 0.41 & 0.59 \end{bmatrix} \text{ for } w, w' = \{\text{calm, stormy}\} \quad (15)$$

so that if it is calm today there is an 85 percent chance that tomorrow will be calm and a 15 percent chance it will be stormy. Likewise, if today is stormy, there is a 41 percent chance of calm weather tomorrow and a 59 percent chance of stormy weather. Consequently, today's weather not only affect today's catch (and thus today's supply) but also provides information about expected future catches and supply.

Since we do not have direct evidence on the values of the remaining parameters, we estimate them using indirect inference. This approach selects the values for the remaining structural parameters, Λ , that minimize the weighted distance between a set of observed moments and those generated by numerical simulations of the structural model.

To capture the dynamics of the price and sales data, we choose coefficients from six least-squares regressions as moments we wish to match. In the first set of four regressions, we summarize the within-day variation in prices and quantities by type of customer. We partition the data by the race of the customer. For each individual transaction, we compute the log difference between the price paid and quantity sold from the average transaction price and quantity on that day for that customer type. For each customer type i , we regress the de-meaned transaction prices, \hat{p}_i , on the de-meaned transaction quantities, \hat{q}_i , and the order in which the transaction occurred, x . In a similar manner, we then regress \hat{q}_i , on \hat{p}_i , and x .

Formally, we estimate

$$\hat{p}^{asian} = \eta_1 \hat{q}^{asian} + \eta_2 x + v_p^{asian} \quad (16)$$

$$\hat{p}^{white} = \eta_3 \hat{q}^{white} + \eta_4 x + v_p^{white} \quad (17)$$

$$\hat{q}^{asian} = \eta_5 \hat{p}^{asian} + \eta_6 x + v_q^{asian} \quad (18)$$

$$\hat{q}^{white} = \eta_7 \hat{p}^{white} + \eta_8 x + v_q^{white} \quad (19)$$

where the v 's are i.i.d. errors. We make no attempt to impose any economic interpretation to the coefficients. The coefficients η_1 , η_3 , η_5 , and η_7 pick up the observed intra-day, customer-type-specific correlations between prices and quantities, while η_2 , η_4 , η_6 , and η_8 capture the variation between transactions early and late in the day.

To capture the inter-day variation in prices and sales, we regress the log of total quantity purchased, q_t^i and log of average price paid p_t^i each day t by customer type i on a constant, day-of-week dummies, a dummy variable equal to one for Asians (ASIAN), and a stormy-dummy interacted with ethnicity (STORM*AS, STORM*WH). Specifically, we estimate

$$p_t^i = \eta_9 + \eta_{10} MON_t + \eta_{11} TUES_t + \eta_{12} WED_t + \eta_{13} THR_t + \eta_{14} ASIAN + \eta_{15} STORM * AS + \eta_{16} STORM * WH + v_t^p \quad (20)$$

$$q_t^i = \eta_{17} + \eta_{18} MON_t + \eta_{19} TUES_t + \eta_{20} WED_t + \eta_{21} THR_t + \eta_{22} ASIAN + \eta_{23} STORM * AS + \eta_{24} STORM * WH + v_t^q. \quad (21)$$

In addition to the 24 regression coefficients we augment the vector of moments with the standard deviations of the six regression error terms, the fraction of transactions conducted with Asians, and the total number of transactions for a total of 32 moments.

In the language of indirect inference, equations (16)-(21) and the additional moments comprise our auxiliary model. We report these moments in table 2. Let Λ denote the vector of the structural parameters we wish to estimate: $\Lambda = \{\mu_q^{asian}, \gamma^{asian}, \mu_q^{white}, \gamma^{white}, \alpha(1), \alpha(2), \alpha(3), \alpha(4), \sigma_{\epsilon^{asian}}, \sigma_{\epsilon^{white}}, \phi_0, \phi_1, \phi_2, \lambda^{asian}, \lambda^{white}\}$.

Our strategy to obtain point estimates for Λ is:

1. Use the Fulton fish market data to compute estimates of the 32 moments described above. Call this vector $\hat{\theta}_T$.
2. After choosing a set of parameters, Λ , solve the structural model for the value function and the pricing rule.

3. Simulate the structural model with price discrimination for S weeks to create a synthetic dataset $y_S(\Lambda)$. To simulate the model, draw four sets of random shocks. Each day draw the weather, w , then, conditional on the weather and the day of the week, draw the catch C . For each period draw the type of customers that shows up, n , (with a possibility of no customer showing up), and then conditional on the type of customer draw the disturbance in the demand function ϵ^n . The paths of prices, quantity sold, and inventories are determined by the solution of the structural model. The set of shocks is held fixed across every choice of Λ .
4. Using the simulated dataset, $y_S(\Lambda)$, run the regression described above to compute the structural model's predictions for the moments in the auxiliary model. Call these predictions, $\hat{\theta}_S(\Lambda)$.
5. Measure the distance between the vector of observed moments and the vector of simulated moments via the criterion:

$$(\hat{\theta}_T - \hat{\theta}_S(\Lambda))' \Omega_T (\hat{\theta}_T - \hat{\theta}_S(\Lambda)) \quad (22)$$

where Ω_T is a 32×32 weighting matrix.

6. Using a hill-climbing algorithm, repeat steps 2 - 5 to find the set of parameters $\tilde{\Lambda}_T$ that minimizes (22).⁴

We construct Ω_T so that the regression coefficients are weighted by the inverse of their variance-covariance matrices. The remaining moments are weighted by least squares. We “over-weight” the fraction of transactions conducted with Asians and the total number of transactions to ensure the model closely matches these two moments. We simulate the model for $S=1320$ weeks or 60 times the number of weeks in our dataset.

To numerically solve the model, we discretize the inventory grid into 45 points from 0 to 40,000 pounds. The grid points are more densely spaced near zero where the value function has more curvature. For each of the $45 \times 2 \times 2 \times 5 \times 60 = 54,000$ (I, n, w, d, t) quintuplets, we maximize the right hand side of equations (8) and (9) recursively over the sales price. Points off the inventory grid are approximated using linear interpolation, and all integration is done by quadrature.

⁴Although our dynamic model is too complex for us to provide analytical results on identification, visual inspection of concentrated slices of (22) and the gradients of $\hat{\theta}_S(\Lambda)$ give us confidence that the parameters are well identified.

μ_q^{asian}	γ^{asian}	μ_q^{white}	γ^{white}	$\alpha(1)$	$\alpha(2)$	$\alpha(3)$	$\alpha(4)$	$\sigma_{\varepsilon^{asian}}$	$\sigma_{\varepsilon^{white}}$
4.53	1.67	4.32	1.49	0.021	-0.181	-0.245	0.191	0.838	0.821
0.07	0.11	0.02	0.10	0.039	0.091	0.120	0.071	0.052	0.055

ϕ_0	ϕ_1	ϕ_2	λ^{asian}	λ^{white}
0.018	0.165	0.837	0.313	0.351
0.050	0.516	0.052	0.015	0.018

Table 1: Estimates of the Structural Parameters

Note: The first row of the table reports point estimates. The second row of the table reports estimated standard errors.

6 Empirical Results

In table 1 we report point estimates for the structural parameters together with estimated standard errors. The parameter values are economically sensible. The parameters in the top row determine the characteristics of the demand curves. Consistent with the findings of Graddy (1995), we find that Asians are more price elastic than whites. In particular the price elasticity for whites is estimated to be 1.49 while this elasticity for Asians is 1.67. Graddy (1995) employs a reduced-form instrumental-variables strategy that estimated the price elasticities to be 1.33 and 1.67, respectively. Unlike the IV approach which identifies the elasticities from exogenous shifts in supply, our simulation estimator pins down these elasticities primarily from the mark-up over the imputed marginal cost reported in (11). Since the basic economic assumptions underlying both approaches are the same, it is reassuring that even with the additional structural assumptions in the current approach, both deliver an identical price elasticity for Asians and price elasticities for whites that are well within two standard errors of each other.

Counterintuitively, despite the fact that whites have a more elastic demand than Asians, the estimated elasticities along with the constant terms imply their reservation prices are lower. From the constant elasticity demand curve (2) and minimum purchase size constraint (3), if we assume $\alpha(d)$ and ε are zero, the reservation price for whites is \$1.16 while the reservation price for Asians is 13 cents higher at \$1.29.

The standard deviations of the demand disturbances, $\sigma_{\varepsilon^{asian}}$ and $\sigma_{\varepsilon^{white}}$, are essentially identical. The large negative numbers for $\alpha(2)$ and $\alpha(3)$ are consistent with the observation that Tuesdays and Wednesdays are quiet relative to Mondays, Thursdays, and Fridays.

The positive point estimates of ϕ_0 , ϕ_1 , and ϕ_2 suggest that the dealer is not simply maximizing total revenue, though neither ϕ_0 nor ϕ_1 are significantly different from zero. Since the point estimate of ϕ_2

is less than one, the dealer's compensation is a concave function of revenue. These parameters imply the distribution of our fictional dealer's earnings has a mean of 8.1 cents per pound with a standard deviation of 2.2 cents. This matches the dealer's statement that he keeps 5-15 cents of each pound sold.

To evaluate the model's ability to match the basic characteristics of the data, we report three comparisons between the model's implications for price and quantity sold and the observed data. For the first set of comparisons, table 2 tabulates two sets of estimated regression coefficients: one using the observed Fulton fish market data and one computed from 1000 simulations of 22 weeks of the model at the estimated parameter values.⁵ As described in the previous section, the structural parameters reported in table 1 were chosen so that the model's predictions for these moments are as close as possible to the observed moments.

The model matches several key facts in the data. Consider the within day price regressions reported in the top panel of table 2. For both Asians and whites, the model matches that higher prices are associated with lower quantities and that prices decline on average throughout the day. For both races, the model's coefficients are well within two standard deviations of the coefficients estimated from the data.⁶

For the within day quantity regressions, reported in the middle panel, the model matches the negative relationship between quantity and price. For both races, the coefficients for the model are within a single standard deviation of the coefficients for the data. The model also predicts, consistent with the data, that quantity sold declines throughout the day, though not to the magnitude seen in the data.

In the inter day regressions reported in the bottom panel, the model matches the fact that during stormy weather, both Asians and whites pay higher prices and purchase less fish. From the estimated coefficients on the Asian indicator variable, we see the model also replicates the fact that Asians pay less and purchase more on average than whites. The model does less well matching the coefficients on the day-of-week indicator variables, but given their large standard errors, our estimation procedure puts relatively little weight on matching these coefficients.

The model predicts that 52 percent of the sales are conducted with Asians. The point estimates of λ^{asian} and λ^{white} reported in table 1 imply that on average, 18.8 Asian and 21.0 white customers show up per day. However since Asians have a higher reservation price than whites, they walk away less frequently than whites. On an average day, only 5.1 Asians, compared to 8.5 whites, walk away after receiving a

⁵That is, we simulate the model 1000 times creating 1000 synthetic 22-week data sets. We run the six regressions on each of these data sets, and report the mean and standard deviation for each statistic.

⁶Even though the lengths of the simulated data sets and the Fulton data set are the same, the model generates many more observations than reported under the observed column. In the Fulton data set, for 1527 transactions the race is either not observed or is another race beside Asian or white (e.g. African-American). In the model, the race of all customers is observed and all customers are either Asian or white. The model does match the observed total number of transactions.

Within Day Price Regressions								
variable	Asian				White			
	observed		simulated		observed		simulated	
	coeff	s.e.	coeff	s.d.	coeff	s.e.	coeff	s.d.
quantity	-0.055	0.007	-0.056	0.006	-0.048	0.011	-0.063	0.008
trans num	-0.045	0.007	-0.038	0.009	-0.049	0.011	-0.051	0.011
root MSE	0.158		0.133	0.006	0.125		0.146	0.007
R-squared	0.09		0.12	0.02	0.09		0.13	0.03
observations	992		1502	49	349		1373	55

Within Day Quantity Regressions								
variable	Asian				White			
	observed		simulated		observed		simulated	
	coeff	s.e.	coeff	s.d.	coeff	s.e.	coeff	s.d.
price	-1.160	0.141	-1.271	0.131	-1.052	0.244	-0.962	0.107
trans num	-0.173	0.033	-0.014	0.023	-0.096	0.054	-0.012	0.021
root MSE	0.724		0.633	0.014	0.581		0.571	0.015
R-squared	0.08		0.07	0.01	0.05		0.06	0.01
observations	992		1502	49	349		1373	55

Inter-Day Regressions								
variable	Price				Quantity			
	observed		simulated		observed		simulated	
	coeff	s.e.	coeff	s.d.	coeff	s.e.	coeff	s.d.
constant	-0.239	0.063	-0.245	0.064	7.362	0.142	7.938	0.127
monday	-0.101	0.079	0.108	0.092	-0.098	0.177	-0.195	0.180
tuesday	-0.043	0.077	0.226	0.090	-0.681	0.173	-0.733	0.187
wednesday	0.003	0.079	0.142	0.088	-0.484	0.176	-0.620	0.177
thursday	0.061	0.077	0.186	0.084	-0.006	0.172	-0.128	0.166
asian	-0.092	0.059	-0.108	0.012	0.627	0.131	0.419	0.068
stormy-white	0.359	0.078	0.166	0.071	-0.253	0.176	-0.380	0.162
stormy-asian	0.339	0.078	0.172	0.071	-0.659	0.174	-0.354	0.177
root MSE	0.369		0.299	0.019	0.826		0.687	0.044
R-squared	0.18		0.17	0.05	0.23		0.27	0.06
observations	221		220	–	221		220	–

Total number of transactions: observed = 2868; simulated = 2875 (std = 88)

Fraction of transactions conducted with Asian customers: observed 0.52; simulated 0.52 (std = 0.01)

Table 2: Regression Results Using Observed and Simulated Data

This table reports two sets of the six regressions described in section 5. For each regression, the first set reports the results using the Fulton fish market data. Along with each estimated coefficient we report the associate standard error. The second set reports the mean point estimate and associated standard deviation from 1000 simulations of 22 weeks from the model.

price offer. Thus on an typical day Asians make 13.7 purchases while whites make 12.5 purchases.

As a second measure of the model's ability to match the basic characteristics of the data, we compare the model's implications for prices and sales to the observed means and standard deviations first reported in Graddy (1995).⁷ In table 3, we see that the model does a good job matching the means and standard deviations of prices for Monday through Thursday. In particular, the model matches the average price difference paid by whites and Asians. For quantities, the model slightly underestimates the average purchase size. It also underestimates the volatility of the purchase quantity. The model does replicate that Asians tend to buy in larger quantities than whites.

As discussed in section 3, Fridays are a day of high volatility. While the standard deviation of prices is about 0.3 on the four other days, it is 0.45 on Friday. The model underestimates both the mean and the volatility of prices on Fridays. It get the mean quantity sold about right though it slightly underestimates the standard deviation of sales. This counterfactual prediction occurs despite the model taking into account the "specialness" of Fridays in three ways. First, through the α parameters the model allows the level of demand to differ on Fridays. Second, the model takes into account that a large fraction of the fish will depreciate over the weekend. Third, the model is parameterized so that Friday catches have both a high mean and a high standard deviation. It may be that demand is more inelastic or that demand is more variable on Fridays than on the other days. The model holds the elasticity of demand and the volatility of demand constant across days.

The model appears to dramatically overestimate the amount discarded. The reported discards in the data are the quantities recorded or observed thrown out. If discards are measured as the difference between fish received and fish sold, the average daily discard is about 175 pounds. The model still slightly overestimates the amount discarded, but its' estimates are reasonable.

As a third and final comparison between the model's predictions and the actual data, we plot in figure 3 a high-low price chart for a single 22 week simulation of the model. Comparing this graph to the high-low price chart in figure 1 we see that the model is successful in matching the wide variation in prices. Just as we see in the actual data, there are days in which the low price exceeds another day's high price. Both the actual and simulated price series exhibit 10-15 day cycles, though these cycles are much more pronounced in the actual data.

To illustrate the intuition underlying these results, we plot in figure 4 the dealer's derived pricing rule,

⁷See table 1 on page 79.

variable	Monday		Tuesday		Wednesday		Thursday		Friday	
	observed	simulated	observed	simulated	observed	simulated	observed	simulated	observed	simulated
Price/transaction (dollar/pound)	0.86 (0.34)	0.86 (0.29)	0.87 (0.37)	0.89 (0.33)	0.88 (0.31)	0.85 (0.28)	0.90 (0.34)	0.91 (0.30)	0.89 (0.45)	0.76 (0.26)
Price/trans (Asian) (dollar/pound)	0.86 (0.36)	0.85 (0.29)	0.80 (0.35)	0.84 (0.28)	0.84 (0.31)	0.80 (0.24)	0.86 (0.34)	0.90 (0.30)	0.83 (0.45)	0.73 (0.26)
Price/trans (white) (dollar/pound)	0.85 (0.30)	0.86 (0.29)	0.90 (0.33)	0.95 (0.37)	0.92 (0.30)	0.92 (0.31)	0.92 (0.33)	0.92 (0.31)	0.87 (0.40)	0.79 (0.27)
Average quantity per sale (pounds)	225 (358)	226 (240)	226 (323)	198 (205)	220 (249)	190 (186)	267 (340)	237 (251)	244 (318)	262 (288)
Average quantity per sale (Asian)	258 (298)	252 (275)	268 (303)	223 (234)	277 (231)	212 (212)	286 (323)	266 (288)	339 (350)	302 (326)
Average quantity per sale (white)	238 (346)	201 (196)	175 (247)	167 (154)	190 (296)	160 (140)	232 (297)	209 (204)	179 (267)	222 (236)
Average quantity discarded (pounds)	30 (91)	309 (303)	65 (176)	262 (242)	14 (64)	253 (188)	53 (166)	155 (129)	126 (334)	243 (275)

Table 3: Estimated Moments from Observed and Simulated Data

Note: The first row reports the means. The second row reports the standard deviations in parentheses.




Figure 3: Simulated Day-to-Day Prices

equation (10), for white customers on calm Thursdays.⁸ This graph shows the optimal price as a function of inventories and within-day time periods. Time goes from 1 (start of day) to 60 (end of day). It is clear that holding time fixed, prices are a decreasing function of the level of inventories. As inventories approach zero, the marginal value to the seller of an additional box of fish rises sharply; thus so does its' price. Similarly, holding inventories fixed, prices are a decreasing function of time of day. Since inventories are steadily decreasing throughout the day, depending on the initial level of inventories and the sequence of demand shocks, the model can generate upward or downward within-day price paths, though as reported in the top panel of table 2, on average prices fall throughout the day.

While we can statistically reject with a high degree of confidence that our structural model is not the true data generating process, overall we find it to be a realistic representation of a Fulton fish market seller. The model's design incorporates the characteristics of the fish seller's problem that Graddy observed first-hand. The economic intuition derived from the model's solution is consistent with the observed price and quantity data as well as the anecdotal evidence. Finally the model successfully imitates many of the key features of the observed transaction data. We now employ this model to perform a counterfactual exercise:

⁸In this case $n = \text{white}$ and $d = 4$. There are 19 other corresponding graphs for each day of the week, type of weather, and customer type. We only plot only one to conserve space since the other 19 are qualitatively similar.




Figure 4: Dealer's Optimal Pricing Rule for White Customers on Calm Thursdays

assuming the firm can no longer price discriminate.

Using the parameter values reported in table 1, we solve the model under the no-price-discrimination assumptions, equations (12) and (13). In table 4 we report the means and standard deviations of prices and quantities with and without price discrimination. Interestingly, if the fictional dealer is unable to charge different prices to Asians and whites, it has only a modest effect on the average price paid. Although they now face the same prices, Asians on average pay a higher price than whites. This is due to whites' lower reservation price. In the no-price-discrimination case whites walk away less frequently than they do under the price-discrimination case (7.6 versus 8.5 times per day) while Asians walk away from a slightly larger number of price offers in the no price discrimination case (5.6 versus 5.1 per day). This endogenous censoring of the price data implies that whites on average pay a slightly lower price than Asians.

In table 4 we see that since Asians pay higher prices in the no price discrimination case they purchase a little less. Since their price elasticity is greater than one, total revenue from Asian customers falls. White customers, on the other hand, now pay lower prices, so their average purchase size rises. Furthermore they make more frequent purchases. Hence, revenue from whites goes up, completely offsetting the lost

variable	observed	price discrim	no-price-discrim
Price/transaction (\$/pound)	0.865 (0.360)	0.850 (0.298)	0.841 (0.291)
Price/trans (Asian) (\$/pound)	0.842 (0.353)	0.823 (0.282)	0.846 (0.292)
Price/trans (white) (\$/pound)	0.892 (0.321)	0.881 (0.311)	0.836 (0.290)
Average quantity per sale (pounds)	249 (303)	226 (342)	224 (238)
Average quantity per sale (Asian)	288 (308)	253 (274)	247 (268)
Average quantity per sale (white)	203 (291)	196 (196)	201 (201)
Daily Revenue from Asians	\$2,684 (1,652)	\$2,499 (1,100)	\$2,414 (1,093)
Daily Revenue from Whites	1,552 (1,095)	1,906 (939)	2,015 (903)
Total Daily Revenue	3,993 (2,207)	4,405 (1,689)	4,429 (1,651)

Table 4: Means and Standard Deviations With and Without Price Discrimination

revenue from Asian customers. Hence there is no loss in revenue from setting a posted price. In this case revenue actually rises slightly.

Of course, in the price discrimination case, the dealer is able to charge both types of customers identical prices, so why is revenue higher in the constrained case? Recall that the dealer's compensation function, (7), is a non-linear function of total revenue. Under the no price discrimination case, average dealer earnings falls by 1/10 of a penny per pound. Hence the dealer makes only slightly more if he can price discriminate.

7 Interpretation and Conclusion

We conclude that our model is a realistic approximation to the observed fish seller's problem; however, like all models, it is a stark simplification of reality. It is possible that these simplifications non-trivially bias downward our computed cost of eliminating price discrimination. In the model, both prices and quantities are continuous variables, while at the Fulton fish market, prices tended to be quoted in 5 cent increments and fish were generally sold in 60 pound boxes. This endogenous discreteness may create kinks and non-convexities that create wedges between the prices it is optimal to charge Asians and whites. Since there is no physical reason for prices to be quoted in such coarse increments, this norm may provide a simple mechanism to limit bargaining and competition thus providing a greater opportunity to price discriminate.

More importantly, the model assumes that the price elasticity for both races is constant across the entire demand curve. While the replication of Graddy's (1995) estimates of these elasticities provides evidence that our point estimates are robust, it may be that these elasticities are constant only along a portion of the demand curves. It may be that as quantity demanded gets closer to the minimum purchase size (i.e. 60 pounds), demand become more inelastic. If this is the case, our analysis underestimates the reservation price for whites, underestimating the value to the seller of price discriminating.

While, the result that price discrimination yields very little additional revenue may be due to shortcomings of our model, it is also possible that sellers continue to price discriminate despite the lack of evidence that it results in increased revenue. Many of the dealers at the market have been operating for years. Indeed, the grandfather of the dealer from which the data were collected started his business in the 1920s. The way of transacting at the market is well-established and negotiated prices are the norm, based on history, despite any analytical evidence that these actually provide the highest revenue. Other fish markets, including the largest market, Skidjee in Japan, conduct business through an auction process. It is

possible that negotiated prices resulting in price discrimination may not be the optimal way to sell fish, but that this method simply exists because of inertia.

8 Appendix: Means and Standard Deviations for Daily Fish Received

In this appendix we report the means $\mu_c(d, w)$ and standard deviation $\sigma_c(d, w)$ of the catch conditional on day of the week and weather. These parameters, together with the Markov transition matrix, (15), yield unconditional daily means and standard deviations of quantity received consistent with those reported in table 1 on page 79 of Graddy (1995).

	Monday	Tuesday	Wednesday	Thursday	Friday
Calm	9058 (7110)	5197 (5034)	4823 (3568)	7783 (4386)	8123 (4463)
Stormy	6596 (3911)	2464 (2340)	5106 (3681)	4288 (3652)	7264 (7456)

Table 5: Means and Standard Deviations of the Catch Conditional on Day or Week and Weather

In pounds of Whiting. Standard deviations are in parentheses.

References

- [1] Aguirregabiria V., 1999. The Dynamics of Markups and Inventories in Retailing Firms. *Review of Economic Studies* 66, 275-308.
- [2] Angrist, J., Graddy K., Imbens G., 2000. The Interpretation of Instrumental Variables Estimators in Simultaneous Equations Models with an Application to the Demand for Fish. *Review of Economic Studies* 67, 499-527.
- [3] Ayres, A., Siegelman, P. 1995. Race and Gender Discrimination in Bargaining for a New Car. *The American Economic Review*, Vol. 85, No. 3, pp. 304-321
- [4] Borenstein, S., 1991. Selling Costs and Switching Costs: Explaining Retail Gasoline Margins. *The RAND Journal of Economics*, 22-3, 354-369.
- [5] Chan, M.H., Hall, G., Rust, J. 2004. Price Discrimination in the Steel Market. Univeristy of Maryland Working Paper.
- [6] Copeland A., Dunn W., Hall G., 2009. Inventories and the Automobile Market. Bureau of Economic Analysis Working Paper.
- [7] Elmaghraby W., Keskinocak P., 2003. Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions. *Management Science* 49, 1287-1309.
- [8] Federgruen A., Heching A., 1999. Combined Pricing and Inventory Control Under Uncertainty. *Operations Research* 47, 454-475.
- [9] Graddy, K., 1995. Testing for Imperfect Competition at the Fulton Fish Market. *The RAND Journal of Economics* 26-1, 75-92.
- [10] Graddy, K., 2006. The Fulton Fish Market. *Journal of Economic Perspectives* 20-2, 207-220.
- [11] Graddy, K., Kennedy, P., forthcoming. When are Supply and Demand Determined Recursively Rather than Simultaneously? *Eastern Economic Journal*.
- [12] Hastings, J., 2009. Wholesale Price Discrimination and Regulation: Implications for Retail Gasoline Prices. Yale University Working Paper.
- [13] Karlin S., Carr R., 1962. Prices and Optimal Inventory Policy. in *Studies in Applied Probability and Management Science*, K. Arrow, S. Karlin, H. Scarf (eds.), Stanford, CA: Stanford University Press.
- [14] Lee, S., 2007. Endogeneity in Quantile Regression Models: A Control Function Approach. *Journal of Econometrics* 141: 1131-1158.
- [15] Reagan P., 1982. Inventory and Price Behavior. *Review of Economic Studies* 49, 137-142.
- [16] Rust, J., Hall G., 2003. Middlemen versus Market Makers: A Theory of Competitive Exchange. *Journal of Political Economy*. 111-2, pp. 353-403.

- [17] Scott Morton, F., Zettelmeyer, F., and Silva Risso, J. 2003. Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities? *Quantitative Marketing and Economics*, 1, pp. 65-92.
- [18] Sweeting, A., 2008. Equilibrium Price Dynamics in Perishable Goods Markets: The Case of Secondary Markets for Major League Baseball Tickets. NBER Working Paper 14505.
- [19] Villas-Boas, S.B., 2009. An Empirical Investigation of the Welfare Effects of Banning Wholesale Price Discrimination. *The RAND Journal of Economics* 40-1, 20-45.
- [20] Whiten T., 1955. Inventory Control and Price Theory. *Management Science* 2, 61-80.
- [21] Zettelmeyer, F., Scott Morton F., Silva-Risso J., 2003. Inventory Fluctuations and Price Discrimination: The Determinants of Price Variation in Car Retailing. Yale School of Management Working Paper.