Price Discrimination in the Steel Market

Hiu Man Chan, *Charles River Associates*
George Hall, *Yale University*
John Rust, *University of Maryland* [*]

June 1, 2004

Abstract

This paper estimates a dynamic model of price discrimination and inventory investment under incomplete information. The model is motivated from an empirical analysis of operations of daily observations on inventories, sales, and purchases of over 2,300 individual products by a U.S. steel wholesaler. The model assumes the wholesaler has a distribution of beliefs about each retail customer's reservation values and posts individual take-it-or-leave-it offers to maximize discounted profits while simultaneously accounting for the firm's optimal inventory decisions. This model is compared to the case in which the the firm must post a uniform price to all customers. We simulate the estimated model and find that the simulated data exhibit the key features of inventory investment and pricing behavior we observe in the data.

Keywords: commodities, inventories, speculation, dynamic programming
JEL classification: D21, E22

---

1 Introduction

This paper estimates a dynamic structural model of price discrimination and inventory investment under incomplete information. The model is motivated by a dataset on the operations of a single steel wholesaler – or in industry lingo – a steel service center. This firm does only minimal production processing earning profits primarily by purchasing steel at wholesale prices from producers and other intermediaries and reselling to retail customers at a mark-up. In very simple terms, the object of the firm is to: 1) buy low and sell high; and 2) bargain with retail customers to obtain the highest markups possible. Hence each day the wholesaler must solve two problem simultaneously: decide how much inventory to hold, and set an retail price for each customer who wishes to purchase steel.

Even though steel is a relatively homogenous good, we document the presence of considerable price discrimination. Our model is designed to match the following pricing facts.

1. There is second degree price discrimination: different customers pay different prices for the same types of steel on the same day.

   (a) There are quantity discounts.

   (b) Distance matters: markups are higher for customers that are closer to the firm.

2. There is considerable day-to-day and within-day variation in retail prices.

3. The underlying wholesale price explains about 3/4 of the variation in the retail price.

4. Retail prices seem only weakly related to current inventory holdings.

But we also want it to match following inventory facts.

5. Purchases are made infrequently and vary considerably in size.

6. Purchases are more volatile than sales.

7. There is no stable inventory/sales relationship and stockouts and near stockouts occur regularly.

These facts are deduced from data on every transaction undertaken by the middleman. We develop a model of optimal speculation and price setting in commodities. We characterize both the inter-temporal procurement and pricing decisions and show that the model's implications are consistent with the stylized facts. We numerically solve the model and embed it in a simulated minimum distance (SMD) estimator to

2

estimate the unknown parameters of the model and test the hypothesis that the firm's behavior is governed by the optimal trading and pricing rules.

Our model builds on the work of Chan's (2002) Phd disseration. In our model, the wholesale price of steel $\{p_t^b\}$ is assumed to evolve according to an exogenously specified first-order Markov process. At the start of each period the firm decides how much new inventory q_t^b to buy at the wholesale price. There is a fixed transaction cost K to placing any order, so the firm will only place sufficiently large orders for which the incremental change in expected profits exceeds K. Each period, at most one customer approaches the firm requesting to purchase q^d steel for immediate delivery. Knowing the current wholesale price, the quantity of inventory on hand and the quantity demanded, the firm quotes a take-it-or-leave-it price to the customer. If the quoted price is below the customer's privately known reservation price, then the sale is consummated; otherwise, no trade occurs. The firm enters the next period with any unsold inventory from the previous period and learns the new wholesale price. The firm incurs the procurement costs at the time of purchase and revenues are collected as the stock is sold. The firm discounts the present value of the stream of profits. No backlogging of demand is allowed. Once the steel has been purchased, all procurement costs are sunk. We contrast this model with one that is identical except that the firm must quote a retail price that is independent of the quantity demanded.

We find that our initial specification is rejected by goodness of fit tests and that our model is only able to explain roughly 1/3 of the variation in retail prices. Never the less, the model does match several of the key facts in the data; it captures the existence of second-degree price discrimination and the presence of quantity discounts in particular. The optimal trading rules derived from the model are also consistent with the inventory facts reported above. Further while the optimal pricing rule is the solution to a relatively complicated dynamic model, along one dimension it can be well approximated by a rule of thumb. Our model outperforms the firm in terms of expected discounted profits but it is not clear whether the firm is not behaving optimally, or whether the profits from our trading rule are higher by luck or the fact that our law of motion for wholesale prices is misspecified. We may also need to include other variables as a basis for price discrimination beside quantity demanded. We are planning to include customer location and an unobserved variable characterizing the level of "naivete" of a customer.

Our interest in measuring the value of price discrimination was sparked by the opaqueness of prices in the steel industry. Unlike many other commodities and assets such as pork bellies and Treasury bills, no centralized market for steel exists, transaction price are not published, and almost everyone in the industry jealously guards their transaction prices. The steel market is best described as a "telephone market" in

which individual transaction prices are private information, a result of individual search, matching, and bargaining between buyers of steel and the middlemen who sell it.[1] In Rust and Hall (2003) we developed a theory of intermediation in which the microstructure of trade in a commodity or asset is endogenously determined. Depending on the relative costs of conducting a transaction there are equilibria consistent with all trade occurring via a *market maker* on a centralized exchange at posted prices, or all trade occurring via decentralized transactions with *middlemen,* at privately quoted prices, or trade segmenting between middlemen and market makers. If search is sufficiently costly and the costs of carrying out a centralized trade sufficiently small, our model suggests that it may be highly profitable for a market maker to enter a market and intermediate trade at posted prices. The secondary market for U.S. Treasury securities has gone through such a transformation and is now one of the most transparent financial markets. In the process, the first firms to publish such prices, such as Cantor Fitzgerald, rose in market dominance (Zuckerman, Davis, and McGee, 2001). So why has not such a transformation occurred in the steel market? Asked another way, why doesn't the firm post its prices on the web and become a market maker?

Recent attempts by potential market makers to create a "transparent" market in steel have so far failed.[2] The firm's primary concern about posting prices is losing its ability to price discriminate. While we cannot measure how much additional business the firm might receive from posting its prices, we can estimate the value of price discriminating that it could potentially lose. By our calculations, the firm makes roughly 85% of its profits from marking up the price of steel over the wholesale price and 15% of its profit from price speculation. What we want to know is how much of the profit due to markups is due to the ability to price discriminate.

While our model is intended to provide an accurate representation of the behavior of a particular middleman in the steel market, this model could also apply to other commodity and asset markets. Modified versions of our model that allow for discrete rather than continuous commodities could be used to describe pricing decisions of many types of finished consumer goods, such as automobiles. Although we are not aware of an existing model that is appropriate for this case, or any empirical models that have attempted to

---

[1] About 50% of all steel sales pass directly from producers to consumers, bypassing the intermediation by middlemen. However most of these sales are to large customers such as automobile producers or shipbuilders whose demand for steel is sufficiently large and regular that these companies find it worthwhile to have their own in-house steel purchasing departments that negotiate directly with steel producers. Most buyers of steel tend to make more irregular, less predictable purchases of different types of steel. These buyers find it cheaper to purchase their steel from middlemen. Also most steel producers only sell steel in large bulk quantities, so that customers who wish to make relatively small purchases go to middleman whose job is to purchase large quantities from producers and sell smaller quantities from their stockpiles at a markup.

[2] See Rust and Hall (2003) pages 396-398 for a discussion of the failure of Enron, e-STEEL, and MetalSite to garner a significant share of transactions in the steel industry.

test optimal trading behavior of a individual middleman in a commodity market, the joint production/price decision we model is a classic issue in the operations research literature going back to Whiten (1955) and Karlin and Carr (1962). We will not provide a comprehensive review of the operations research literature, but instead refer interested readers to the citations in Gallego and Ryzin (1994) and Federgruen and Heching (1999). Most of the papers in this literature focus on industries in which goods must be sold by a fixed deadline (e.g. fashion apparel, airline seats), thus these papers analyze models with only a single or small number of periods, or an exogenous number of (usually just one) order decisions, or deterministic demand, or a constant per-unit procurement (production) cost.

Thomas (1974), Polatoglu and Sahin (2000) and series of papers by Chen and Simchi-Levi (2002, 2003a, 2003b) study multi-period models with endogenous price setting, stochastic demand, and the opportunity to continually reorder with a fixed order cost. In these papers the per-unit procurement cost is constant. The main hurdle for studying models with a fixed cost to ordering and endogenous pricing is that the optimal ordering policy is no longer guaranteed to be (S, s). Both of these papers state sufficient conditions such that (S, s) holds and solve for the optimal policy under these conditions.[3] Since the per-unit procurement cost is a constant, inventories never exceed S. In Thomas' model the optimal policy is for inventory levels greater than s, the firm should place no orders and set the price as a decreasing function of current inventories. When inventories are below s, the firm should order up to S and and charge the price associated with S. Polatoglu and Sahin also find the optimal order policy in their model to be (S, s), but their computed pricing function is non-monotonically decreasing in the inventory. To us, these non-monotonicities appear to be due to computational/approximation errors, but Polatoglu and Sahin conclude "the optimal pricing trajectories ... can be complicated enough to make the evaluation and administration of the optimal pricing strategy impractical."

One recent related paper in the economic literature, Zettlemeyer, Scott Morton, and Silva-Risso (2003), studies the retail market for new automobiles. They model individual automobile dealers allowing for an endogenous retail pricing policy but assume an exogenous reorder policy; their model implies that the

---

[3]Thomas (1974) also constructs a counter-example. Like Thomas, we can find examples such that the optimal order policy in our model has multiple (S, s) bands. Multiple (S, s) bands can occur since choosing an optimal retail price implies taking the maximum across K-concave functions. Since the maximum of two K-concave functions is not guaranteed to be K-concave, the standard proof by induction of Scarf (1960) can not be implemented. One way to generate multiple (S, s) bands is to assume the firm can only choose prices from a coarse discrete grid (e.g. one high price, one low price). In these cases, there can be a set of (S, s) bands associated with low inventories and high prices, and a set of (S, s) bands associated with high inventories and low prices. We have found that if the price grid is sufficiently fine, the optimal policy is a single set of (S, s) bands. In the model presented below the firm can choose from a continuum of retail prices between a lower bound and upper bound.

retail price should be decreasing in the level of dealer inventories. Using price and inventory data for a set of California dealerships, they find the average retail price at a dealership with ample inventory is about $230 per car less than the retail price offered at a dealership with low inventories.

We extend this joint production/pricing literature by 1) adding a stochastic serially correlated per-unit procurement price (hence allowing inventories to exceed S), 2) embedding a bargaining problem into the endogenous price decision, 3) computationally solving the infinite horizon model, and 4) estimating the structural parameters of the model and comparing the model's implications directly to data.

2 The data

Through our contact with the general manager of a large U.S. steel wholesaler, we acquired a high frequency micro database on transactions in the steel market. This firm has provided us with daily data on all of its 2300+ individual steel products. Our empirical results are based on all transactions made between July 1, 1997 to June 2, 2003 (1500 business days) for two of its highest volume products. For each transaction we observe the quantity (number of units and/or weight in pounds) of steel bought or sold, the sales price, the shipping costs, and the identity of the buyer or seller. We also matched data from the wholesaler's records to data from *Reference USA*, a database of detailed information on over 12 million US businesses. For each buyer, we know the firm's location (i.e. street address, city, state, zip code, and MSA), line of business (primary SIC code and up to four secondary SIC codes), number of employees, annual sales volume, Yellow Page ad size, and credit rating.

Although this is an exceptionally clean and rich dataset, we only observe prices on the days the firm actually made transactions: the firm does not record any price information on days that it does not transact (either as a buyer or seller of steel). This shortcoming of our dataset is much more important for steel purchases than steel sales, since the firm purchases new steel inventory in the wholesale market much less frequently than it sells steel to its retail customers. Indeed, even for its highest volume products, it makes purchases only about once every two weeks. It is possible to get weekly and monthly survey data on prices for certain classes of steel products through trade publications such as *Purchasing Magazine* and *American Metal Market*. However, since there are no public exchange markets for steel products, transaction in the steel market are carried out in private negotiations. Hence these price surveys rely on participants in the steel market to report truthfully the prices they paid or received. The firm often faces considerably different prices than those reported in the survey data.

Figure 1: Times series plot of the inventory for product 4.

Figure 2: Purchase prices (solid line) and retail prices (dashed line) for product 4. For the purchase price series, the size of the marker is proportional to the size of the purchase.

We illustrate our data by plotting the time series of inventories and prices of one of the firm's primary products in figures 1 and 2. This product, which we call product 4, is one of the highest volume products sold by this firm. It is also a benchmark product within the industry since the prices of several other steel products are often computed as a function of this product's price.

In our plot of wholesale transaction prices in figure 2 (the lower curve with the large circles), we used straight line interpolations between observed purchase prices at successive purchase dates. The black circle at each purchase date is proportional to the size of the firm's purchase in pounds. Figure 1 plots the evolution of inventories over the sample period. Purchases of steel are easily recognizable as the discontinuous upward jumps in the inventory trajectories. As is evident from the saw-tooth pattern of the inventory holdings, the firm purchases the product much less frequently than it sells it. Although we observe 1500 business days in the sample, the firm made purchases on only 212 days, with an average size of 135,000 pounds and made sales on only 899 days, with an average size of 31,800 pounds. Purchases are also considerably more volatile than sales. The unconditional standard deviation of purchases is 3.3 times larger than the unconditional standard deviation of sales. If we condition on a purchase or sales occurring, the standard deviation of purchases is 7 times large than the standard deviation of sales. Clearly this firm is not "production smoothing."

Further, the firm does not target an inventory-to-sales ratio. The patterns of the dots in figure 2 suggests that the firm is more likely to purchase large quantities of steel when either the wholesale prices is low or the level of inventories are low (although other economic factors seem to be influencing the firm's purchase decisions as well.) As can be seen in figure 1 inventories over the period range from zero to over 3 million pounds. For example, in 2000 the firm had no inventory of product 4 for eight days in January and twelve days in October. While in the interest of space we do not plot a sales series, these inventory levels correspond to a days-supply of inventory varying from zero to over 275 days. While the general manager states that he likes to have about 60 days of supply on hand, he quickly follows up that he is willing to hold more or less inventory depending on his opportunities to buy at low prices. As discussed in more detail in Hall and Rust (2000a, 2003), the firm's inventory behavior is driven more by opportunistic buying than by targeting an inventory-to-sales ratio.

Figure 2 also plots the interpolated sequence of daily retail sales prices. Retail sales occur on about two out every three business days, so the amount of interpolation in the retail price series is modest. The wholesale and retail prices move in a roughly parallel way, although there is considerable day-to-day variation in retail prices. Retail prices are quoted net of transportation costs, but still much of the the high

frequency variation is due to observable factors. As a first attempt to explain this high frequency variation, we perform a reduced-form analysis.

Table 1 presents the results from a linear regression of the net retail price (in dollars per ton) on quantity of demand, transportation cost, material cost, most recent purchase cost, inventory level, and monthly dummies, using data for two products: one we label product 2 and the other, product 4, the one discussed above. The two variables that relate to the opportunity cost of an unit of inventory, material cost and most recent purchase cost, explain a large part of the variation in retail price. Material cost is a firm-calculated moving average accounting cost on which the commissions earned by sales agents are based. Hence, while this accounting cost differs from the true shadow value of steel, it is extremely relevant in the determination of the retail price. In fact, material cost alone explains as much as 77% and 71% of the variation in retail price of product 2 and product 4 respectively. The most recent purchase cost, on the other hand, proxies the replacement cost of steel and by itself accounts for 65% and 62% of the variation. For product 2, the coefficients on these two variables sum to .985, while for product 4, the sum is .85; hence, on average, a one-cent rise in the wholesale price leads to almost a one-cent rise in the retail price.

In addition to cost, quantity of demand and transportation cost also contributes significantly to the price variance, signaling the presence of price discrimination based on quantity and distance. For product 2, a customer can get a 0.12% discount off the average net price for each additional ton of steel bought. For product 4, the discount is lower, at 0.05% per ton. In both cases, each additional dollar of transportation cost translates into about 60 cents of reduction in net price. The firm is therefore absorbing part, but not all, of the increase in transportation cost incurred by customers located further away. The firm recognizes its comparative advantage over customers located nearby who can save on transportation costs and charges a premium on the net price. Yet, in gross price, more distant customers still end up paying more in total.

Somewhat surprising is the positive coefficient on current inventory for both products. Like Zettle-meyer, Scott Morton, and Silva-Risso (2003) and the OR papers in the procurement/pricing literature, the structural model we present below predicts that as the level of inventory increases, the marginal value of inventory and hence the optimal retail price decreases. One potential explanation is that firm's use of material cost to compute commissions does not accurately reflect the true shadow value of an additional unit of inventory; hence the sales force does not have the incentive (despite it being profit maximizing) to reduce prices when inventories are high. However, the general manager clearly understands that retail prices should fall when inventories are high, so we suspect the reason for this positive coefficient is that there are multiple regimes in the price data. For example from July, 1997 to December, 1998 can be con-

	Product 2		Product 4	
variable	point estimate	standard error	point estimate	standard error
Constant	0.5953	0.0695	1.1807	0.0656
Quantity of Demand	-0.0047	0.0009	-0.0019	0.0007
Transportation Cost	-0.6032	0.1620	-0.5996	0.1040
Material Cost	0.7506	0.0282	0.5925	0.0262
Most Recent Purchase Cost	0.2354	0.0241	0.2436	0.0219
Inventory Level	0.0233	0.0038	0.0214	0.0018
January	0.0064	0.0295	-0.1128	0.0366
February	0.0541	0.0313	-0.0949	0.0340
March	0.0611	0.0338	-0.1619	0.0350
April	0.0992	0.0343	-0.1086	0.0337
May	0.0604	0.0349	-0.0902	0.0338
June	-0.0177	0.0300	-0.1575	0.0323
July	0.0606	0.0275	-0.1047	0.0321
August	0.0038	0.0280	-0.1804	0.0308
September	-0.0164	0.0318	-0.1906	0.0393
October	-0.0779	0.0282	-0.1696	0.0283
November	-0.0332	0.0300	-0.1899	0.0351
No of Observations	1492		1484	
R^2	0.8070		0.7770	

Table 1: OLS Estimation Results on Net Price

sidered a "high price regime". However even if we run this regression on various subsamples in the data the inventory coefficient is rarely negative. Of course, both inventories and retail prices are endogenous variables, so instrument variable analysis may be appropriate in this case; but is not clear which potential instruments would be correlated with inventories but not prices.

From the regression analysis and discussions with firm executives, we conclude that the steel market is characterized by *second degree price discrimination,* i.e. middlemen charge different per unit prices depending on the quantity sold, or more generally, they charge different prices to different customers based on the individual characteristics of each customer. Price discrimination is possible due to the incompleteness of the steel market: as discussed above, there are no exchanges where steel is traded, and thus no publicly posted prices for steel products.

3   A Dynamic Model of Steel Pricing and Inventory Control

This section describes a dynamic model of optimal pricing and inventory investment decisions that is designed to explain the key facts presented in the previous section. We derive some of the model's general implications for pricing decisions, but a rigorous empirical test of the model (using a numerically solved version of the model) is deferred to section 4.

Although the middleman trades in over 2,000 steel products, we focus on trading and pricing decisions of two specific products, each treated as separate "profit centers" (i.e. the middlemen seeks to maximize expected profits for each product without "cross subsidization" across products). Customers arrive probabilistically, and the quantity q^d each wishes to purchase varies from customer to customer. Our hypothesis is that the middleman uses a pricing policy that maximizes expected discounted profits. We show that the optimal pricing rule involves second degree price discrimination and describe how this optimal pricing rule depends on q^d and other observable characteristics of the customer.

We model the middleman's pricing decision as a "stage game" within an infinitely repeated trading game in which inventory holdings are periodically replenished via bulk purchases of steel at wholesale market prices that evolve randomly over time. Our model predicts both customer-specific pricing decisions (i.e. price discrimination) and decisions on the timing and size of new purchases of steel. First degree price discrimination, i.e. charging a price that fully extracts all possible surplus from each customer, is not possible due to incomplete information on the customer's *reservation price.* Since customers are searching for the best price, it is reasonable to suppose they are using an optimal search strategy, and thus they purchase steel when they receive a price quote that is below their reservation price. While the middleman may observe some characteristics of each potential customer (e.g. their location and how much steel they want to buy), there is a significant amount of information that the middleman cannot observe. As a result, we model the middleman as having beliefs about the possible reservation values of each potential customer, which we model as a conditional distribution $F(r|q^d,l,z,p^b,x)$ that depends on the customer's location l, the quantity of steel the customer wishes to buy q^d, other characteristics z, some of which we may not observe as econometricians, and variables that are useful in forecasting steel prices, x. The middleman's beliefs about a customer's reservation values may also be affected by overall market conditions, which are captured by the variables (p^b,x).

Suppose that a customer arrives and the middleman observes the customer's characteristics (q^d,l,z).

We assume that the middleman will quote a "take it or leave it" sales price p^s given by

$$p^s = \underset{p}{argmax} \left[1 - F(p|q^d, l, z, p^b, x)\right] [p - c] \tag{1}$$

where c denotes the per unit *shadow price* of the q^d units of steel to be sold. Taking the first order condition to this optimization problem, it is straightforward to show that the optimal sales price can be written as

$$p^s = c + \frac{1 - F(p^s|q^d, l, z, p^b, x)}{f(p^s|q^d, l, z, p^b, x)} \tag{2}$$

where $f(r|q^d, l, z, p^b, x) = F'(r|q^d, l, z, p^b, x)$ is the conditional density function of customer reservation values. Equation (2) states that the optimal take it or leave price p^s is a markup over the shadow price c, where the markup term is an inverse hazard function involving the middleman's conditional distribution of beliefs about the customer's reservation value.

It is natural to ask why we don't model realized transaction prices as a result of a more involved dynamic bargaining process between the middleman and his customers. One reason we don't do this is empirical: we typically don't observe multiple rounds of bargaining between the middleman and his customers. In a typical transaction the middleman quotes a price, and the customer either agrees and takes the quote, or decides to continue looking. Sometimes the middleman might have a chance to make a second lower quote if the customer does not accept the initial quote (and is still on the phone), but our impression that this happens infrequently. The steel firm does not record its unsuccessful price quotes to customers. However we know from personal observation and from discussions with the firm's manager that the majority of transactions result from a single price quote to the customer.

A second reason why we don't model a dynamic bargaining "subgame" between the middleman and his customers, is that under reasonable conditions we can show that bargaining leads to suboptimal outcomes relative to presenting the customer with a take it or leave it price quote. There are two main reasons for this. The first reason is to economize on time. From a practical standpoint, bargaining is a time consuming process, and with over 2,000 different steel products and hundreds of different customers calling to purchase various amounts of these products each day, the middleman (even with the assistance of team of salesmen who are also in charge of price setting decisions) faces a tradeoff between the number of transactions per salesman and the incremental profits from allowing extended bargaining. Second, by appealing to Roger Myerson's (1981) seminal work on "optimal auction design", appendix 1 shows that under certain assumptions the middleman's optimal bargaining strategy is to precommit not to bargain and quote a fixed take it or leave it price. Appendix 1 shows that the optimal retail price in equation (2) is identical to the optimal reserve price in Myerson's optimal "auction" in the special case where there is only one buyer.

12

Although a formal proof of the conclusion that a take it or leave it price quote is optimal depends on the assumption that the potential buyer knows the middleman's shadow price c (so that the "auction" is a one-sided game of incomplete information, as in Myerson's analysis), we feel that the use of a take it or leave it price quotes should be nearly optimal even in the case where the potential customer does not know c. In any event, given that this assumption is also empirically "realistic" in the case of the steel middleman, we will adopt this assumption hereafter without further comment.[4]

3.1 Timing of Information and Decisions

We now describe our model in somewhat more detail, starting with a description of the sequence of events that take place each business day:

1. At the start of day t the middleman observes the value of the *state variables:* the inventory level q_t, the wholesale purchase price p_t^b at which the firm can buy steel, and a vector x_t of other variables that are useful for forecasting future steel prices and sales.

2. The middleman makes a *purchase decision* choosing the size of additional inventory q_t^b to be purchased for immediate delivery. Purchases must be non-negative and are subject to a maximum holding constraint: $0 \leq q_t^b \leq \overline{q} - q_t$ where $\overline{q}$ is the firm's maximum storage capacity.

3. With probability $\lambda(p_t^b, x_t)$ a single customer arrives and asks for a price quote. The customer has a demand to buy a quantity $q_t^d > 0$ (which is revealed to the middleman) and has a reservation price r_t (which is private information that is not revealed to the middleman). Besides q_t^d the middleman observes the location l_t of the customer, and possibly a vector z_t of other characteristics. Based on this information the middleman makes a single take it or leave it sales price quote p_t^s. If the price quote is less than the customer's reservation price, a sale is made, with a quantity sold of $q_t^s = \min[q_t^b + q_t, q_t^d]$. Otherwise no sale is made, and $q_t^s = 0$. If no buyer arrives for a price quote, $q_t^d = 0$.

4. Sales on day t determine the level of inventories on hand at the beginning of business day $t + 1$ via the standard inventory identity:

$$q_{t+1} = q_t + q_t^b - q_t^s. \tag{3}$$

---

[4]The assumption that the middleman can precommit to a take it or leave it price may not be a good approximation for other situations such as automobiles, where transactions occur less frequently, and the surplus per transaction tends to be higher. It is much more common to observe multiple rounds of offers and counter offers between customers and auto dealers.

5. New values of the state variables (p^b_{t+1}, x_{t+1}) are drawn from the transition density $g(p^b_{t+1}, x_{t+1} | p^b_t, x_t)$.

We assume that the firm incurs a fixed cost $K \geq 0$ associated with placing new orders, which implies that the *order cost function* $c^b(q^b, p^b)$ is given by

$$c^b(q^b, p^b) = \begin{cases} p^b q^b + K & \text{if} \quad q^b > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

The middleman also incurs costs for storing inventory. The per period cost of holding q_t units of inventory is $c^h(p^b_t, q_t, x)$. If the firm cannot fulfill an accepted sale (i.e. if $q^d > q^b_t + q_t$ and the $p^s_t < r_t$), the firm stocks out and incurs a fixed goodwill cost γ. The firm discounts future profits at rate $\beta = \exp^{r/365}$, where r is per annum interest rate.

3.2 The Value Functions

Let $V(p^b, q, x)$ denote the expected present discounted value of a firm *after* it makes its optimal inventory ordering decision, given that the current wholesale price is p^b (the price that the middleman can buy at), the start of day inventory is q, and where x represents other information useful for forecasting future prices and demand for steel.

$$V(p^b, q, x) = \max_{0 \leq q^b \leq \overline{q} - q} \left[W(p^b, q + q^b, x) - pq^b - I\{q^b > 0\}K \right] \tag{5}$$

The W function represents the expected discounted value of the firm *before* it has had a chance to replenish its inventories by buying more steel.

$$\begin{aligned} W(p^b, q, x) &= [1 - \lambda(p^b, x)]W_{na}(p^b, q, x) + \\ &\quad \lambda(p^b, x) \int_{q^d} \int_l \int_z W_a(p^b, q, q^d, x, l, z)\mu(dl, dz | p, x, q^d)m(dq^d | p^b, x), \end{aligned} \tag{6}$$

where W_a is the value of the firm conditional on the arrival of a customer and W_{na} is the value of the firm when no customer arrives. When a customer arrives, the firm must make a decision about what price p^s it should quote the customer. The W_a function represents the firm's value when it makes the best possible price quote

$$W_a(p^b, q, q^d, x, l, z) = \max_{p^s} \left[ES(p^b, p^s, q, q^d, x, l, z) - c^h(p^b, q, x) - \gamma I\{q^d > q\} + \beta EV_a(p^b, p^s, q, q^d, l, z, x) \right] \tag{7}$$

where q^d is the amount of steel a customer demands, ES is the firm's expected sales revenue from quoting the price p^s, c^h is the cost of holding steel inventories, and $\mu(dz, dl | p, x, q^d)$ is the conditional distribution

14

of the customer's other characteristics (l, z), where l denotes the customer's location and z denotes a vector of other characteristics that the middleman might observe.

The firm's expected sales revenue, conditional on an arrival and quoting the price p^s, is

$$ES(p^s, p^b, q, q^d, l, z, x) = p^s \min[q^d, q][1 - F(p^s | q^d, l, z, p^b, x)]. \tag{8}$$

This is the familiar "price times quantity sold" times the probability of making the sale.

The function EV_a is the expected discounted value of future profits conditional on the arrival of a customer and conditional on quoting a price p^s. It is given by

$$EV_a(p^b, p^s, q, q^d, l, z, x) = F(p^s | q^d, l, z, p^b, x) EV_{na}(p^b, q, x) +$$
$$[1 - F(p^s | q^d, l, z, p^b, x)] \int_{p'} \int_{x'} V(p', \max[0, q - q^d], x') g(dp', dx' | p, x). \tag{9}$$

Thus, with probability $F(p^s | q^d, l, z, p^b, x)$ the customer will "walk" after hearing the quoted price (this is the probability that the customer's unknown reservation price is less than the price the middleman quotes), and with the complementary probability the customer will accept the quoted price. The function EV_{na} represents the firm's expected value when it loses the sale, and EV_a is the firm's expected value when it makes the sale. When no customer arrives, the value of the firm is given by

$$W_{na}(p^b, q, x) = -c^h(p^b, q, x) + \beta EV_{na}(p^b, q, x), \tag{10}$$

where

$$EV_{na}(p^b, q, x) = \int_{p'} \int_{x'} V(p', q, x') g(dp', dx' | p^b, x). \tag{11}$$

3.3 Characterizing the Optimal Pricing and Trading Strategy

The solution to this model consists of a pair of decision rules, one for the optimal price-setting decision p^s, and another for the optimal purchase quantity q^b

$$p^s = p^s(q^d, l, z, p^b, q, x)$$
$$q^b = q^b(p^b, q, x) \tag{12}$$

Under fairly general conditions the optimal ordering strategy q^b can be shown to be a *generalized* (S, s) *rule.* That is, there exist two functions $S(p^b, x)$ and $s(p^b, x)$ such that $q^b(p^b, q, x)$ can be represented as

$$q^b(p^b, q, x) = \begin{cases} 0 & \text{if } q \geq s(p^b, q, x) \\ S(p^b, q, x) - q & \text{otherwise} \end{cases} \tag{13}$$

15

Figure 3: Example of $S(p)$ and $s(p)$ bands

The function $S(p^b,q,x)$ is the desired or *target inventory level* and $s(p^b,q,x)$ is the *ordering threshold*.[5] It is not difficult to show that $S(p^b,x) \geq s(p^b,x)$ and that $S(p^b,x) = s(p^b,x)$ when $K = 0$. Furthermore, we will show below that under additional assumptions both S and s are decreasing functions of p^b.

The fact that the optimal ordering strategy is an (S,s) rule implies that V is given by

$$V(p^b,q,x) = \begin{cases} W(p^b,q,x) & \text{if } q \geq s(p^b,x) \\ W(p^b,S(p^b,x),x) - p^b(S(p^b,x)-q) - K & \text{if } q < s(p^b,x) \end{cases} \tag{14}$$

This equation says that when $q < s(p^b,x)$, it is optimal for the middleman to order more inventory, and in this region, the value function V is a linear function of q with slope equal to p^b. That is, the marginal value of an additional unit of inventory at the purchase stage is

$$\frac{\partial}{\partial q} V(p^b,q,x) = p^b \quad \text{if } q < s(p^b,x). \tag{15}$$

Also at the target inventory level $q = S(p^b,x)$ we have

$$\frac{\partial}{\partial q} V(p^b,q,x) = p^b \quad \text{if } q = S(p^b,x). \tag{16}$$

---

[5]In previous work (see Hall and Rust, 2000b) we have provided a sufficient condition on sales price — the *no expected loss condition* — that implies the optimality of the generalized (S,s) rule. However this work assumed a 'reduced-form" markup rule which differs from the endogenously determined markup rules considered in this paper. As discussed in the introduction there are counterexamples to the optimality of the generalized (S,s) rule when prices are endogenously chosen to maximize profits even though the our no-expected loss condition is satisfied. We have verified via numerical solutions that the optimal ordering strategy is of the (S,s) for the specifications estimated below.

While the value function has a non-convex kink at $q = s(p^b, x)$ (and is thus not globally concave), V can be shown to be show to be strictly concave for $q > S(p^b, x)$. This implies that

$$\frac{\partial}{\partial q} V(p^b, q, x) < p^b \quad \text{if } q > S(p^b, x). \tag{17}$$

We summarize this discussion as

Proposition 1: *Assume that the optimal ordering strategy is a generalized (S, s) rule. At the purchasing stage, the marginal value of an additional unit of inventory is equal to p^b when $q < s(p^b, x)$, is greater than p^b when $s(p^b, x) < q < S(p^b, x)$ and is less than p^b when $q > S(p^b, x)$.*

From equation (7), we see that the firm's optimal pricing rule is given by

$$p^s(q^d, l, z, q, p^b, x) = \underset{p}{argmax}[ES(p, p^b, q, q^d, l, z, x) + \beta EV_a(p^b, p, q, q^d, l, z, x)] \tag{18}$$

where EV_a is given in equation (9). Rewrite EV_a as

$$EV_a(p^b, p^s, q, q^d, l, z, x) = \int_{x'} V(p', q, x') g(dp', dx' | p^b, x) -$$
$$[1 - F(p^s | q^d, l, z, p^b, x)] \int_{x'} [V(p', q, x') - V(p', q - \min[q, q^d], x')] g(dp', dx' | p^b, x). \tag{19}$$

Once a customer arrives and demands q^d, the value of an additional unit of inventory changes depends on q^d. Define the *shadow cost* of selling q^d units from inventory as

$$c^s(q^d, q, p^b, x) \equiv \beta \int_{x'} \frac{V(p', q, x') - V(p', q - \min[q, q^d], x')}{\min[q, q^d]} g(dp', dx' | p^b, x) \tag{20}$$

For small q^d, the difference quotient $[V(p, q, x) - V(p, q - \min[q, q^d], x)] / \min[q, q^d] \simeq \frac{\partial}{\partial q} V(p, q, x)$ is the marginal opportunity cost (in terms of lost discounted profits) of selling a unit of inventory. Thus c^s is the expected discounted value of this opportunity cost, and thus it provides the expected discounted cost of the loss in future value from selling amount q^d now.

We can rewrite the firm's optimal retail pricing problem in equation (18) as

$$p^s(q^d, q, l, z, q, p^b, x) = \underset{p}{argmax} \left[\min[q, q^d][1 - F(p | q^d, l, z, p^b, x)][p - c^s(q^d, q, p^b, x)] \right]. \tag{21}$$

Hence the optimal pricing rule involves setting a sales price p^s that is a markup over the shadow value of inventories $c^s(q^d, q, p^b, x)$

$$p^s = c^s(q^d, q, p^b, x) + \frac{1 - F(p^s | q^d, l, z, p^b, x)}{f(p^s | q^d, l, z, p^b, x)} \tag{22}$$

This is the same pricing rule that we claimed in equation (2).

17

The shadow price $c^s(q^d, q, p^b, x)$ given by (20) measures the loss in expected future discounted profits from selling q^d units today. Clearly selling q^d units today leads to an opportunity cost of forgoing the sale of those units in a future period. The firm can always order more inventory, so that one measure of the shadow price would be the discounted expected future wholesale price, since this is the per unit expected discounted cost of replacing the q^d units sold today with new units purchased in the wholesale market. However the formula for c actually depends on the level of inventory q since if the middleman is "overstocked", or if current wholesale prices are too high, he may not wish to replace the q^d units sold today via purchases in the wholesale market tomorrow. This is an expectation of a secant approximation to the partial derivative of V, $\partial/\partial q V(p^b, q, x)$. From Proposition 1, we know that if $q < s(p', x')$, then this secant approximation equals next period's wholesale price p'. However if $q > S(p', x')$ and $q - \min[q, q^d] > S(p', x')$, then Proposition 1 implies that the secant approximation will be less than the wholesale price p'. We conclude that

Proposition 2: *The shadow price of inventory is generally not equal to the expected future wholesale price:*

$$c^s(q_t^d, q_t, p_t^b, x_t) \neq E\left\{p_{t+1}^b | p_t^b, x_t\right\}. \tag{23}$$

Thus, the conditional expectation of next period's wholesale price is an upper bound on the shadow price of inventory. If the middleman always immediately replaced the inventory resulting from a sale to a customer in the next period, the shadow price would equal the discounted expected wholesale price, however our theory of optimal trading shows it is generally not optimal to do this. When the firm is "overstocked" (i.e. when $q > S(p^b, x)$), it will be optimal to wait to restock inventory in some subsequent period after enough inventory has been decumulated, or until wholesale prices have fallen low enough to push the firm below its ordering threshold, $q < s(p^b, x)$. The opportunity cost of holding inventory, the delay in restocking, and the potential of capital losses due to adverse movements in the wholesale price, are the primary reasons why the shadow price is generally less than the discounted conditional expectation of next period's wholesale price.

We conclude this subsection with a final proposition characterizing how different variables affect the firm's pricing decision.

Proposition 3: *Optimal sales prices $p^s(q^d, l, z, q, p^b, x)$ are a function of customer characteristics (q^d, l, z), current inventory q, and overall conditions affecting the steel market (p^b, x), including the current wholesale price p^b. Customer characteristics (l, z) affect sales prices only via the markup term in equation (22).*

Inventory holdings affect sales prices only via the shadow price of inventory, $c^s(q^d, q, p^b, x)$. The remaining variables, (q^d, p^b, x) affect the sales price through both the shadow price and the markup terms.

4 Findings

We solve and estimate a special case of the model presented above. First we assume there are no x and z variables. We assume the wholesale price evolves according to a truncated lognormal $AR(1)$ process:

$$log(p^b_{t+1}) = \mu_p + \phi_p log(p^b_t) + w^p_t \tag{24}$$

where w^p_t is an IID $N(0, \sigma^2_p)$ sequence. The holding cost function takes the form

$$c^h(q_t) = \phi_1 q_t + \phi_2 q^2_t$$

The arrival probability is a constant, λ.

Conditional on the arrival of a customer, the fixed demand of the customer q^d_t is a random variable whose distribution depends on the SSC's purchase cost. Let $m(q^d_t | p^b_t)$ denote the conditional density function of the fixed demand. Since the firm's spot purchase cost correlates positively with the general market price, the fixed demand is expected to correlate negatively with the wholesale price. The distribution of the customer's reservation price depends on the wholesale price because the wholesale price reflects the general market trend, and therefore the price at which the buyer may purchase steel from an alternative supplier. The customer's reservation price is likely to be higher when the wholesale is higher. On the other hand, the customer's marginal utility from steel is probably diminishing with quantity. Hence, the customer's reservation price per unit of steel, an average figure, is assumed to be a decreasing function of quantity demanded.

When a customer arrives for a price quote, the fixed quantity of demand q^d_t is a lognormal distribution conditional on the wholesale price p^b_t, with locational parameter

$$\mu_{q^d}(p^b_t) = \bar{\mu}_{q^d} - \eta_{p^d, p^b} \ln(p^b_t), \tag{25}$$

and standard deviation parameter σ_{q^d}. Conditional on the purchase cost and the fixed quantity of demand, the reservation price of the customer p^r_t is a truncated lognormal distribution with location parameter

$$\mu_{p^r}(p^b_t, q^d_t) = \bar{\mu}_{p^r} + \eta_{p^r p^b} \ln(p^b_t) - \eta_{p^r q^d} \ln(q^d_t) \tag{26}$$

and standard deviation σ_{p^r}.

Both our theory suggests and discussions with the general manager confirm that the endogenous sampling of the wholesale price is not innocuous. The firm makes purchases infrequently and at irregularly spaced intervals. The decision whether to purchase or not and how much to purchase depends critically on the expected path of the wholesale price. The firm is more likely to make purchases when prices are low (and expected to rise) than when prices are high (and expected to fall). Hence simple interpolation of prices between successive purchase dates provides a misleading picture of underlying price dynamics. Since steel prices are primarily determined via private bilateral negotiations, we cannot go to public datasets to overcome this sampling problem.

Since most of the variation in the retail price is explained by movements in the wholesale price, we must make accurate inferences about the evolution of the firm's best purchase price. Hence our primary econometric challenge is to infer the law of motion of the spot buy price $\{p_t^b\}$ that is observed only at a subset of times $\{t_1, \ldots, t_n\}$. To estimate the model we employ the simulated minimum distance (SMD) estimator presented in Hall and Rust (2003). This SMD estimator maintains the assumption that the firm's steel purchasing and pricing decisions are given by the decision rules of the dynamic programming model described in the previous section. This SMD estimator is a simulated moments estimator (SME) (Lee and Ingram, 1991 and Duffie and Singleton, 1993), applied to a situation where the data are endogenously sampled.

Let $\{\xi_t\}$ denote the censored price process (i.e. with $p_t^b = 0$ when $q_t^o = 0$), and let θ denote the $L \times 1$ vector of parameters to be estimated. The SMD estimator is based on finding a parameter value that best fits a $J \times 1$ vector of moments of the observed process:

$$h_T \equiv \frac{1}{T} \sum_{t=1}^{T} h(\xi_t), \tag{27}$$

where $J \geq L$ and h is a known (smooth) function of ξ_t that determines the moments we wish to match. We then simulate S realizations of the $\{p_t^b\}$ process for any candidate set of parameter values θ, and censor the price process in exactly the same way as the observed data using our the decision rule (12) as the sampling rule. For each of the simulations of the model we construct the vector of J moments; we then take an average of the moments across the S independent simulations, $h_{S,T}(\theta)$. The simulated minimum distance estimator $\hat{\theta}_T$ is defined by:

$$\hat{\theta}_T = \underset{\theta \in \Theta}{argmin} (h_{S,T}(\theta) - h_T)' W_T (h_{S,T}(\theta) - h_T), \tag{28}$$

where W_T is a $J \times J$ positive definite weighting matrix.

20

We have considerable freedom in our choice of moments to match h to use in the criterion. We match the means, variances and histograms (four of the five quintile bins) of the p, p^r, q^o, q^s, and q processes as well as the number of days a sale or purchase is made for a total of 42 moment conditions. We set the number of simulations to $S=30$.

parameter	Product 2		Product 4	
	point estimate	standard error	point estimate	standard error
r	0.0769	0.005	0.0521	0.007
λ	0.940	0.025	0.905	0.009
μ_{q^d}	4.86	0.036	5.19	0.010
σ_{q^d}	0.822	0.016	0.888	0.011
μ_{p^r}	0.0658	0.0031	0.0749	0.003
σ_{p^r}	0.0350	0.0022	0.0264	0.0005
η_{p^r,q^d}	0.0204	0.0005	0.0174	0.0007
η_{p^r,p^b}	1.067	0.002	1.027	0.0012
η_{q^d,p^b}	-0.200	0.011	-0.047	0.0018
K	14.57	1.51	8.05	1.78
γ	1.82	253.6	4.47	188.4
ϕ_1	-0.000128	4.37×10^{-5}	-0.000211	7.16×10^{-5}
ϕ_2	1.644×10^{-6}	6.2×10^{-8}	6.12×10^{-7}	2.67×10^{-8}
ρ_p	0.980	2.2×10^{-8}	0.979	5.2×10^{-5}
μ_p	0.0584	0.0001	0.0615	9.5×10^{-5}
σ_p	0.0285	0.0007	0.0254	0.0003
$\chi^2(26)$	1914		1347	

Table 2: Estimation Results using data for product 2 and product 4.

We estimate the model for the two products independently and report the point estimates and aympototic standard errors in table 4. The interest rate r is estimated to be quite sensible, 7.6 and 5.2 percent per annum. While the general manager said his borrowing costs were lower than this, the firm may be more impatient than the bank. The quantity data are in hundred-weight (i.e. in 100's of pounds) so the price parameters are in dollars-per-hundredweight (or cents per pound). The fixed cost to placing an order, K, is quite small for both products $14.50 and $8 per order. This is also the case for the lost sale parameter, γ which is under $5 for both products, however given the large standard errors the criterion is very flat right around the point estimate.

The point estimates of the AR(1) wholesale price process ϕ_p, μ_p, and σ_p described by equation XX imply the uncensored price process for product 4 has a mean of 18.82 cents per pound and a standard

deviation of 2.20; for product 2, the mean of the uncensored price process is 18.83 and the standard deviation is 2.48.

The arrival probabilities, λ of 94% and 90.5% imply that nine out of every ten days a customer calls for a price quote even though the customers buy on six out of very ten days.

Although we estimated the parameters for each of these products independently, it is reassuring that all the point estimates are reasonably similar across the two products. However, several of the standard errors are implausibly small. For example, the standard errors for the wholesale price process parameter ϕ_p are. As we conjectured in Hall and Rust (2003), we expect these small standard errors may be to small discontinuities in the estimation criterion.

The SMD estimator provides a formal criterion of the validity of the model. Since the number of moment conditions exceeds the number of parameters estimated ($J > L$) the model is over-identified. Following Hansen (1982), we use the objective function to test the over-identifying restrictions:

$$\frac{T}{(1+1/S)^2}(h_{S,T}(\hat{\theta}) - h_T)' \left[\hat{\Omega}(h)\right]^{-1} (h_{S,T}(\hat{\theta}) - h_T) \to \chi^2(J-L) \tag{29}$$

We report this criterion in the bottom row of table 4. By this criterion, both models are decisively rejected. We discuss below in more detail along which margins the models fail to match the data.

moment	model simulation	data
mean(buy price)	18.95	17.01
var(buy price)	6.13	5.64
mean(sell price)	19.57	19.06
var(sell price)	4.97	5.96
mean(purchase size)	1193	1349
var(purchases)	154	636
mean(sale size)	277	318
var(sales)	7	13
mean(inventory)	8151	7704
mean(markup)—small sale	1.19	2.11
mean(markup)—medium sale	0.95	1.80
mean(markup)—large sale	0.75	1.61
number of purchases	217	212
number of sales	888	899

Table 3: Product 4: Comparison of Selected Moments from 30 Simulations and Data

In table 3 we report a subset of the moments we used in the SMD estimator along with the average moments from 30 simulations of the model. Estimated model matches closely some of the In the model,

the firm make a sale on average 888 out of 1500 days. Given the point estimate of the arrival probability, λ of .90, that two-thirds of the customers who arrive actually purchases steel. It is interesting to note that the mean of the censored buy price series is *higher* than the uncensored mean (18.95 versus 18.82); the variance of the censored and uncensored series are about the same. this is due to the fact than when prices are high and inventories are the low the firm make a large number of small purchases.

While, in the data the sales price is more volatile than the purchase price, the opposite is the case for the model. The model implies declining markups as a function of the quantity purchased, but the estimated magnitude of the markup is about 80 cents per pound too low.

The model implies the procurements are more volatile than sales, but underestimates the variance of purchases. The reason for the underestimation is that the firm made several very large purchases (e.g. they bought an entire shipload of steel). The model implies that instead of making one large purchase, the firm should make a series of medium-sized purchases.

We plot the estimated value function in figure 4 and derivative of the value function with respect to inventory holdings, q in figure 5. The value of the firm is monotonically increasing the quantity of inventories at every price. When inventories are high (i.e. when the firm is going to do more selling than buying), the value of the firm is increasing in the wholesale price, but when inventories are low (i.e. the firm is going to do more buying than selling), the value of firm is decreasing in the wholesale price of steel. From figure 5 one can see that as stated by condition (15) in the previous section, the derivative of the value function with respect to q is equal to p^b when inventories and prices are low (i.e. below $s(p)$). While the kink in the value function as $s(p)$ is not apparent in value function itself given the resolution of the graph, the kink is apparent in the derivative of the value function.

Since the retail price is endogenous, the optimal inventory policy is no guaranteed to be on the (S, s) form. However, for the estimated parameter values (as well as for all economically interesting parameter values), an (S, s) inventory policy is optimal. In figure 6 we plot the procurement policy rule for product 4 as a function of the firm's beginning of period inventory holding and the current wholesale price. The firm places no orders when the current wholesale price is high and/or current inventories are high. Both the $S(p)$ and $s(p)$ bands can be seen the order function. The $S(p)$ band is the quantity ordered when $q = 0$. The order threshold, $s(p)$ is the set of q's such that the firm is indifferent between ordering, so it the set of points such the decision rule becomes greater than zero. Both $s(p)$ and $S(p)$ are decreasing functions of p, capturing the basic feature of successful price speculation, namely, that the firm should buy large quantities when prices are low and sell this inventory when prices are high. If we resolve the model assuming the

Figure 4: Value Function for Product 4

Figure 5: Derivative of the Value Function with respect to quantity on hand, q

firm must charge an uniform price to all customers (i.e. it can not condition its pricing decision on q^d), the order policy is almost identical to order policy allowing for price discrimination.

In figures 7-11 we illustrate the firm's endogenous pricing decision rule. In figure 9 we plot the expected markup (knowing p^b and q, but before learning q^d) as a function of the wholesale price and current inventory. Above the $s(p)$ band the markup is a decreasing function of both the wholesale price and current inventories. Below $s(p)$ the markup does not depend on the current level of inventory (since the firm always orders up to $S(p)$) and is increasing in the wholesale price. This difference in the relationship between the wholesale price and the markup can be seen in figure 10. When $q = 0$ the firm is below $s(p)$; in the case the markup is greater when the wholesale price equals 20 cents per pound. More importantly, figure 10 illustrates the presence of quantity discounts. For all (p^b, q) pairs, the markup falls dramatically for small orders, and then flattens out, but is always monotonically decreasing. Further note that when inventories are well above $S(p)$ (i.e the curve for q=2,000,000 lbs and p=20), the cost of holding cost inventories can be so large that it is optimal for the firm to charge a negative markup, i.e. the optimal retail price can be below the current purchase price.

In figure 8 we plot slices of the pricing rules for four different spot prices holding quantity demanded constant. As can be seen in this figure, when current inventories q are below $s(p)$, the retail price (and thus the markup) is a constant reflecting the firm's policy to order to $S(p)$ whenever $q < s(p)$. The markup
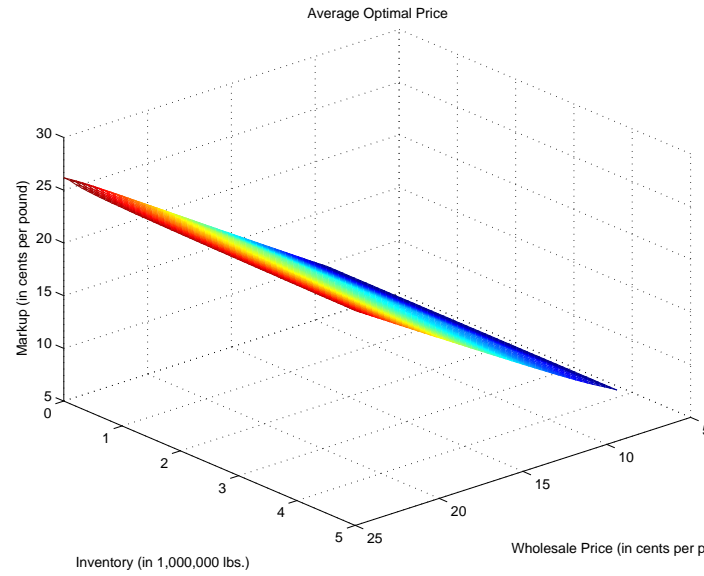
24

Figure 6: Optimal Ordering Rule for Product 4.

Figure 7: Optimal Expected Retail Prices for Product 4: $E\{p^s(p^b,q,q^d)|p^b,q\}$

Figure 8: Optimal Retail Price for Product 4 Given $q^d = 50,000$ lbs.
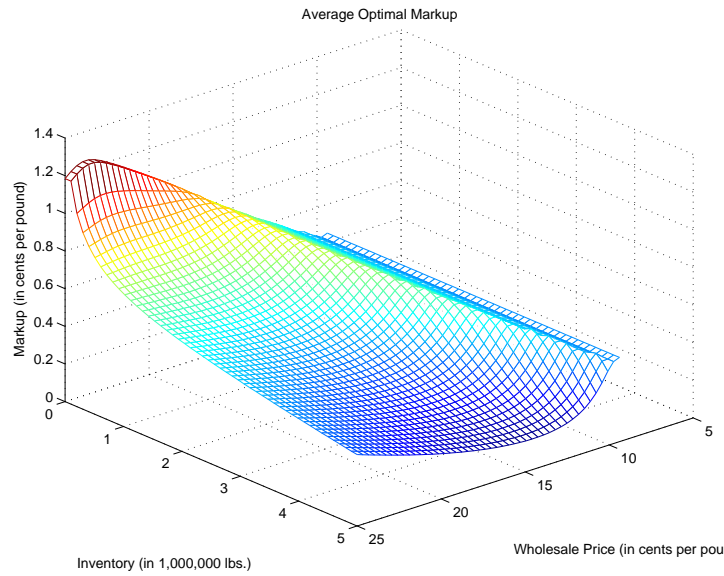

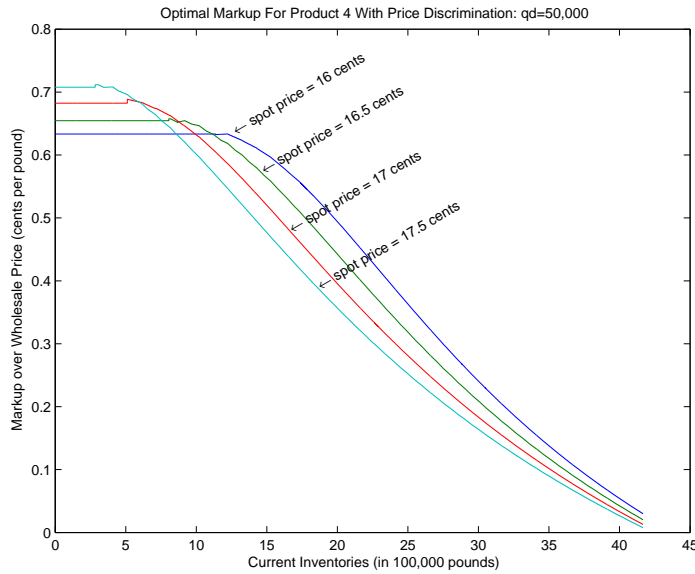Figure 9: Expected Markups over Wholesale Prices for Product 4: $E\{p^s(p^b,q,q^d)|p^b,q\} - p^b$

Figure 10: Optimal markup for Product 4 as a Function of Inventory.

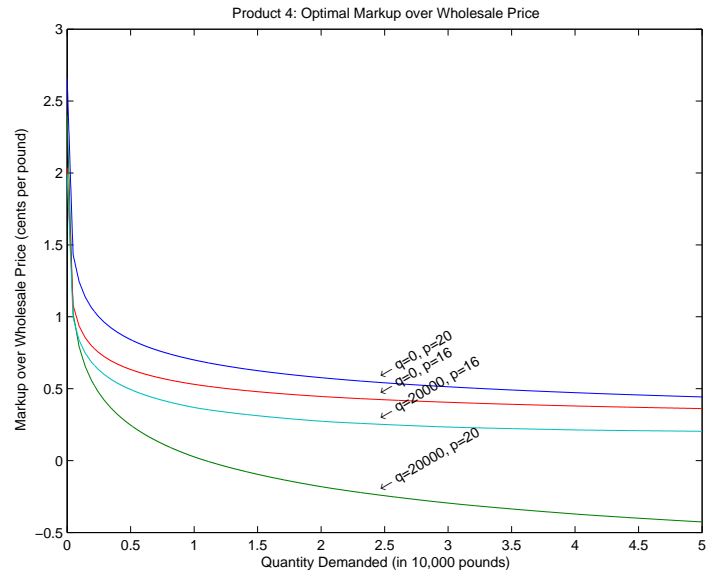Figure 11: Optimal markup for Product 4 as a Function of Quantity Demanded.

jumps at $q = s(p)$, by roughly $\frac{K}{S(p)-s(p)}$ and for $q > s(p)$, the markup is a decreasing function of quantity on hand. This pricing rule is consistent with our intuition concerning the shadow price of an additional unit of inventory for the firm. Below $s(p)$ the marginal value of an additional unit of inventory is a constant. For $s(p) < q < S(p)$ the marginal value of an additional unit exceeds the shadow value for $q < s(p)$, since the firm would like to postpone incurring the fixed cost K. However as q increases the opportunity cost of holding the inventory dominates this desire to postpone paying K; consequently in this region, the shadow value of inventories steadily decreasing at an increasing rate. In language of the operations research literature, this is a generalization of the *base stock list price policy*: if the inventory level is below $s(p)$, it is increased to the base stock level, $S(p)$ and the "list price" is charged; if the inventory level is above the base stock level, then nothing is ordered, and a "price discount" is offered.

While this pricing rule is the solution to a rather complicated dynamic problem, we can deduce a rather simple rule of thumb. Since the retail price at $S(p)$ is equal to retail price just below $s(p)$, the slope of the pricing function is approximately $\frac{K}{(S(p)-s(p))^2}$ when inventories are above $s(p)$.[6] While the pricing function above $s(p)$ is not linear, it appears that a linear rule closely approximates the optimal rule.

In figures 14 and 15 we plot a simulated path of inventories and prices.

---

[6]The downward sloping $s(p)$ band can be deduced in figure 8 by connecting the kinks in the pricing functions at $s(p)$. To trace out the $S(p)$ band, connect points to right of the $s(p)$ such that the retail price equals the retail price just below $s(p)$.

Figure 12: Actual versus Predicted Prices for Product 4 with Price Discrimination.

Figure 13: Actual versus Predicted Prices for Product 4 with No Price Discrimination.


Figure 14: Simulated inventory data from the estimated model for product 4.

Figure 15: Censored (solid line) and uncensored (dotted line) purchase prices, p_t^b from a simulation for product 4.
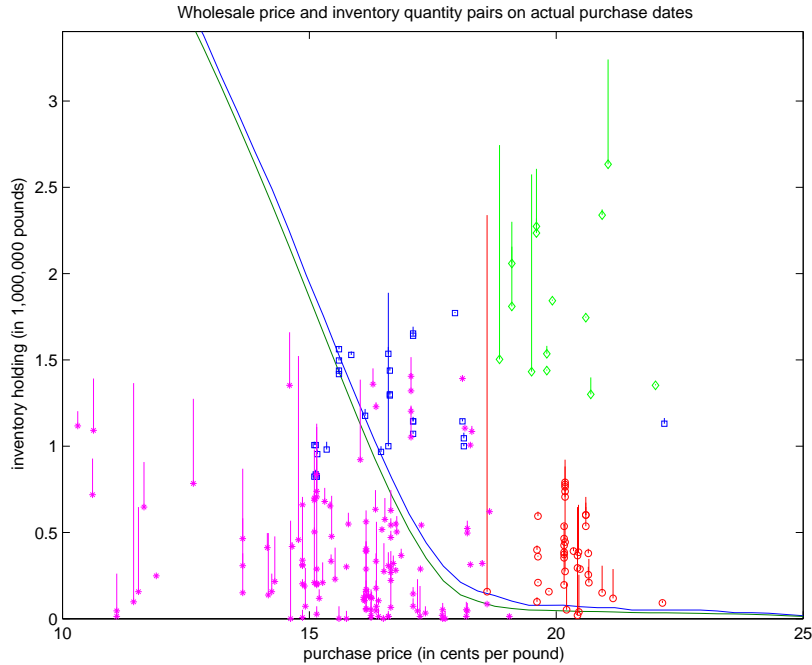
27

Figure 16: Scatterplot of purchase price and inventory holding pairs on the 212 purchase dates. for product 4. The length of each vertical line measures the size of the purchase. The downward-sloping curves are the $S(p)$ and $s(p)$ bands from the model.

Figure 16 illustrates the main failure of the model on the "quantity side". In this figure, the downward sloping curves are the (S, s) bands from the model for product 4. We superimpose a scatterplot of the purchase price and inventory level pairs we observe on 212 days the firm purchased product 4. We then draw vertical lines from each one of these (p, q) pairs to the point $(p, q + q^b)$, so the length of each vertical line represents the size of the purchase q^b. If our model was true, each of the observed (p, q) pairs would be to the southwest of the $s(p)$ band. Further, the top of each of the vertical lines would just touch the $S(p)$ band. Clearly this is not what we see.

If we divide the sample into four periods, we can see evidence of shifting (S, s) bands. For the initial period from July 1, 1997 to March 27, 1998, we mark the (p, q) pairs in figure 16 with circles. During this period prices were around 20 cents per pound, and the large purchase on March 27, 1998 at 18.5 cents per pound discussed in section 2.1 looks consistent with sharply downward-sloping (S, s) bands. In the second period, from March 28, 1998 to December 31, 1998, we denote the (p, q) pairs with diamonds. Prices during this period return range from 19 to 22 cents, but the (S, s) bands have shifted upward. Why the shift? During this second period, sales were unusually high, averaging between 40,000 and 50,000 pounds

per day. It appears there was an outward shift in the demand schedule for steel during this period that our model completely misses. But then with the large increase in imported steel in 1999, the demand schedule appears to shift in dramatically. During this third period which we date from January 1, 1999 to October 8, 1999 and mark with squares, average daily sales fell to 10,000 pounds per day – just 1/4 of what they were the previous year. It appear the (S,s) bands shifted in as the firm reduced its inventory in response to this demand shock despite falling prices. The fourth regime from October 9, 1999 to the end of the sample is denoted with stars. During this period, sales are relatively stable and larger purchases are associated with lower prices.

In table 4 we report the average coefficients and standard errors from 30 simulations of the two models. Both models capture the presence of quantity discounts as the coefficients on Quantity of Demand are negative, though neither are as large in absolute value as we see in the data. Also the coefficients on Material Cost and Most Recent Purchase Cost are not as large as we see in the data. Finally, the model predicts a negative coefficient on the Inventory Level while we observe a positive coefficient in the data.

Finally, we use simulations from the estimated model to deduce the relative importance of price discrimination for the overall profitability of the firm. By substituting the law of motion for inventories (3) into the firm's per-period profit function, the discounted present value of the firm's profits can be expressed as:

$$
\sum_{t=0}^{T}\beta^t\pi(p_t^b,p_t^s,q_t^r,q_t+q_t^b) = \sum_{t=0}^{T}\beta^t(p_t^s-p_t^b)q_t^s+q_0p_0^b+\sum_{t=1}^{T}\beta^t(p_t^b-(1+r)p_{t-1})q_t -
$$
$$
\sum_{t=0}^{T}\beta^t I(q_t^b)K - \sum_{t=0}^{T}\beta^t\gamma I\{q^d>(q_t+q_t^b)\} - \sum_{t=0}^{T}\beta^t c^h(q_t+q_t^b,p_t). \quad (30)
$$

The first term on the right hand side of equation (30) can be interpreted as the discounted present value of the markup paid by the firm's retail customers over the current wholesale price while the third term can be interpreted as the discounted present value of the capital gains or loss from holding the steel from period $t-1$ into period t. The fourth, fifth and sixth terms are the discounted present values of the order costs, lost sales penalty, and the holding costs incurred by the firm over the sample period.

To decompose the firm's actual profits we simulate the buy price path between purchase dates via importance sampling. That is, for each interval between successive purchase dates, we simulate buy price paths that are consistent with the estimated law of motion (24) and the observed purchase prices at the beginning and end of the interval. Since our theory implies that the firm places an order anytime the quantity falls below the order threshold, $s(p)$, we truncate the simulated price process by rejecting any

variable	Product 2				Product 4			
	Simulations		Data		Simulations		Data	
	point estimate	standard error	point estimate	standard error	point estimate	standard error	point estimate	standard error
Constant	1.09716	0.03596	0.5953	0.0695	0.94753	0.03111	1.1807	0.0656
Quantity of Demand	-0.00063	0.00017	-0.0047	0.0009	-0.00065	0.00019	-0.0019	0.0007
Transportation Cost	—	—	-0.6032	0.1620	—	—	-0.5996	0.1040
Material Cost	0.01826	0.04654	0.7506	0.0282	0.12708	0.04102	0.5925	0.0262
Most Recent Purchase Cost	0.03458	0.03671	0.2354	0.0241	-0.00140	0.03526	0.2436	0.0219
Inventory Level	-0.00038	0.00002	0.0233	0.0038	-0.00017	0.00001	0.0214	0.0018
No of Observations	911		1492		888		1484	
R^2	0.64		0.81		0.52		0.78	

Table 4: OLS Estimation Results on Sell Prices from 30 Model Simulations and Coefficients Reported in Table 1.

	Firm's Actual Performance			Model With Price Discrimination			Model Without Price Discrimination		
markup	$364,166	(17,573)	87%	$510,539	(26,156)	69%	$501,709	(26,394)	69%
capital gain	56,100	(18,910)	13%	228,684	(43,920)	31%	226,422	(44,700)	31%
holding cost	-106,656	(0)		-112,085	(43,723)		-106,760	(42,507)	
lost sales	-0-	(0)		-28	(13)		-40	(17)	
order costs	-2,654	(0)		-2,711	(282)		-2,330	(271)	
total profits	310,954	(X,XXX)		624,398	(53,806)		619,001	(54,806)	

Table 5: Profit Decomposition For Product 2 Equation (30)

The firm's profits cover the 1500 days studied and are discounted back to the start of the sample period, July 1, 1997. The profits from the model simulations are the discounted sum over 1500 periods. The profit numbers reported are the average across 100 simulations. The numbers in parentheses are the standard deviations from the 100 simulations. Total profits are the sum of the first four rows.

	Firm's Actual Performance			Model With Price Discrimination			Model Without Price Discrimination		
markup	$443,382	(22,996)	85%	$151,275	(8,200)	32%	$146,554	(8,729)	31%
capital gain	79,894	(24,695)	15%	337,528	(70,033)	68%	333,272	(69,738)	69%
holding cost	-91,301	(0)		-120,701	(50,192)		-115,619	(47,904)	
order costs	-0-	(0)		-205	(117)		-239	(131)	
order costs	-1,551	(0)		-1,441	(304)		-1,300	(280)	
total profits	430,424	(2,565)		366,456	(X,XXX)		362,668	(58,702)	

Table 6: Profit Decomposition For Product 4 Equation (30)

The firm's profits cover the 1500 days studied and are discounted back to the start of the sample period, July 1, 1997. The profits from the model simulations are the discounted sum over 1500 periods. The profit numbers reported are the average across 100 simulations. The numbers in parentheses are the standard deviations from the 100 simulations. Total profits are the sum of the first four rows.

paths such that $q_t < s(p_t)$ for any draw within the simulated paths. This simulation method is discussed in more detail in Hall and Rust (2003).

For a given interpolated price series, we decomposed the firm's profits using the actual data for q, q_t^s, and q_t^o and our point estimates for r, K, μ_p, ϕ_p and σ_p. In tables 5 and 6 we report the average decomposition from 100 simulated buy price paths.

5   Conclusions

References

[1] Chan, H. M. (2002) "Empirical Essays on Bargaining and Price Discrimination: A Case Study of a Steel Intermediary" Ph.D. dissertation, Yale University.

[2] Duffi e, D. and K.J. Singleton (1993) "Simulated Moments Estimation of Markov Models of Asset Prices" *Econometrica* Vol 61, No 4, pp 929-952.

[3] Federgruen, A. and A. Heching (1999) "Combined Pricing and Inventory Control Under Uncertainty" *Operations Research* Vol 47, No. 3, pp. 454-475.

[4] Hall, G. and J. Rust (2000a) "An Empirical Model of Inventory Investment by Durable Commodity Intermediaries" *Carnegie-Rochester Conference Series on Public Policy*, Vol 52, pp 171-214.

[5] Hall, G. and J. Rust (2000b) "The (S, s) Rule is an Optimal Trading Strategy in a Class of Commodity Price Speculation Problems" manuscript, Yale University.

[6] Hall, G. and J. Rust (2003) "Simulated Minimum Distance Estimation of a Model of Optimal Commodity Price Speculation with Endogenously Sampled Prices" manuscript, Yale University.

[7] Hansen, L. P. (1982) "Large Sample Properties of Generalized Method of Moments Estimators" *Econometrica* Vol 50, pp 1929-1954.

[8] Karlin, S. and R. Carr (1962) "Prices and Optimal Inventory Policy" in *Studies in Applied Probability and Management Science*, K. Arrow, S. Karlin, H. Scarf (ed.), Stanford, CA: Stanford University Press.

[9] Lee, B. and B.F. Ingram (1991) "Simulation Estimation of Time-Series Models" *Journal of Econometrics* Vol 47, pp 197–205.

[10] Myerson, R. (1981) "Optimal Auction Design" *Matthematics of Operations Research*, Vol 6, No 1, pp. 58-73.

[11] Polatoglu, H. and I. Sahin (2000) "Optimal Procurement Policies Under Price-Dependent Demand" *International Journal of Production Economics*, Vol 65, pp. 141-171.

[12] Rust, J. and G. Hall (2003) "Middle Men versus Market Makers: A Theory of Competitive Exchange" *Journal of Political Economy*, Vol 111, No 2, pp 353-403.

[13] Scarf, H. (1960) "The Optimality of (S, s) Policies in the Dynamic Inventory Problem" In *Mathematical Methods in the Social Sciences* K. Arrow, S. Karlin and P. Suppes (ed.), Stanford, CA: Stanford University Press, 196–202.

[14] Thomas, L. (1974) Price and Production Decisions with Random Demand" *Operations Research*, Vol. 22, No. 3, pp. 513-518.

[15] Whiten, T. (1955) "Inventory Control and Price Theory" *Managment Science*, Vol 2. pp. 61-80.

[16] Zettlemeyer, F., F. Scott Morton, and J. Silva-Risso (2003) "Inventory Fluctuations and Price Discrimination: The Determinants of Price Variation in Car Retailing" manuscript, Yale School of Management.

[17] Zuckerman, G., A. Davis and S. McGee (2001) "Before and After: Why Cantor Fitzgerald Can Never Re-Create What It Once Was" *Wall Street Journal*, October 26. p. A1.

Appendix 1: The Optimal Bargaining Strategy is Not to Bargain

This appendix shows that the optimal selling strategy of a middleman in a commodity market with randomly arriving customers is under certain assumptions "isomorphic" to the optimal auction design problem analyzed by Myerson (1981). We show that the optimal sales price derived in equation (2) of section 3 is identical to the optimal reserve price in Myerson's optimal auction when there is only a single buyer.

A1.1. Review of Myerson's Result

In Myerson's framework the seller values a single indivisible object to be sold at c. There is a single buyer who values the object at r. However although the buyer knows the seller's valuation, the seller does not know the valuation r of the buyer. Instead the seller has beliefs that the buyer's valuation is a realization from a cumulative distribution function $F(r)$ with density $f(r)$. Myerson defined the buyer's *virtual valuation* $v(r)$ as

$$v(r) = r - \frac{1 - F(r)}{f(r)} \tag{31}$$

Myerson showed that the optimal selling mechanism (in the case of multiple buyers with the same distribution of valuations F (in the "regular case" where the virtual valuation function is montonically increasing in r) is to award the object to the buyer i with the highest virtual valuation $v_i(r_i)$. The price that the winning buyer would pay is the second highest valuation among the buyers, provided this is higher than a reservation price, p^s, given by $p^s = v_i^{-1}(c)$. Using the fact that the reservation price is the inverse of the virtual valuation evaluated at the seller's valuation of the object, we have

$$c = v_i(v_i^{-1}(c)) = v_i(p^s) = p^s - \frac{1 - F(p^s)}{f(p^s)} \tag{32}$$

Rearranging this equation we have the following expression for the seller's optimal reserve price

$$p^s = c + \frac{1 - F(p^s)}{f(p^s)} \tag{33}$$

Comparing this equation to the equation for the optimal retail price p^s that we derived above (see equation 2), we see that both have the same form.

In the case of only a single bidder, there is no second highest bid. Myerson showed that the optimal selling strategy in this case is to sell the object to the buyer if the buyer's virtual valuation $v(r)$ exceeds the seller's valuation c. The price that the buyer pays in this case is

$$p^s = \min\{p | v(p) \ge c\} \tag{34}$$

Under regularity conditions that ensure that the virtual valuation $v(p)$ is a monotonically increasing function of p, then v is invertible, and the expression for p^s becomes $p^s = v^{-1}(c)$, and the reservation price has the same representation as a markup over the seller's cost c that we presented in equation (33) above.

A1.2 Isomorphism Between Myerson's Model and the Bargaining Subgame

Now let us consider a bit more formally how the bargaining stage game between a customer and the middleman is isomorphic to the optimal selling strategy derived by Myerson when there is only a single buyer. It may appear that in our problem the seller is selling a continuously divisible quantity q^d where is in Myerson's analysis there is only a single indivisible item to be sold. However recall that we assumed that the buyer has an exogenously specified demand for a predetermined quantity q^d of the commodity. Given this, the buyer's only objective is to secure this quantity at the lowest possible price. The problem is therefor isomorphic to the problem of trying to purchase a single indivisible item at the lowest possible price. The fact that the actual quantity that is purchased is a continuous quantity is irrelevant to the analysis under our assumptions. Indeed, even when the middleman cannot fulfill the full order, the quantity to be sold $\min[q, q^d]$ does not affect the middleman's decision about how to set the optimal price. We can see from equation (21) in section 3 that the quantity to be sold $\min[q, q^d]$ factors out of the middleman's optimization problem as a multiplicative scaling constant that does not directly affect the choice of an optimal price. However $\min[q, q^d]$ does affect the middleman's choice of sale price indirectly since it affects the shadow valuation of the quantity to be sold.

Thus, the fact that the middle man is selling a continuous quantity to the buyer is irrelevant to the analysis of the optimal selling strategy since we assumed that the buyer has an exogenously determined need for the quantity q^d and we have ruled out the possibility of bargaining over price-quantity bundles: the only case where customer purchases only a part of his desired purchase quantity p^d is where $q < q^d$, and the sale causes the middleman to "stock out". The middleman's shadow value $c(q^d, q, p^b, x)$ reflects any "restocking costs" but is treated as exogenous from the standpoint of the bargaining problem with the buyer. Also as noted above, we assume that the buyer knows $c(q^d, q, p^b, x)$ so that the problem is one of one sided incomplete information that Myerson considered.

A second difference between our bargaining stage game and the selling problem that Myerson considered is that the "valuations" in our problem are "reservation values" and thus are derived valuations from an underlying optimal search problem. However this does not affect the isomorphism, since just as a buyer in Myerson's problem is willing to buy an item if its price is less than the buyer's valuation, a customer in our problem is willing to buy $\min[q, q^d]$ from the middleman if the price the middleman quotes is less

than their reservation problem. In both our case and Myerson's case, buyers (customers) wish to buy the object (commodity) at the lowest possible price. Conversely the seller (middleman) wishes to sell the object (commodity) at the highest possible price. Thus the incentives of the buyer and seller in Myerson's model are exactly the same as the incentives of the middleman and the customer in our formulation of the 'bargaining stage game"

A final difference is that in Myerson's framework the seller has beliefs given by an (unconditional) distribution function $F(r)$ of the buyer's valuation r. In our problem the middleman has beliefs given by a conditional distribution $F(r|q^d, l, z, p^b, x)$ of the customers reservation value r. However these differences are inconsequential: if we identify the buyer's valuation r with the customer's reservation value r, and the seller's beliefs about the buyer's distribution of values $F(r)$ with the conditional distribution $F(r|q^d, l, z, p^b, x)$, we see the problems are isomorphic and our formula for the optimal take it or leave it price in equation (22) is the same as the equation for the optimal reservation price (33).

A sufficient condition guaranteeing the optimality of the take it or leave it price is that the virtual valuations are strictly monotonically increasing functions of the buyer's valuation, r, and that the buyer's virtual valuation at r, $v(r) \geq c$, where c is the seller's valuation of the object. It is not difficult to show that the second order condition for optimization defining the optimal take it or leave it price for the seller, $p^s(c)$, implies this sufficient condition for Myerson's result. The optimal take it or leave it price is

$$p^c(c) = \underset{p}{argmax} \ [1 - F(p)][p - c] \tag{35}$$

The first order condition for this problem is

$$0 = -f(p^s(c))(p^s(c) - c) + 1 - F(p^s(c)) \tag{36}$$

The second order necessary condition for the optimum must be satisfied:

$$0 > -f'(p^s(c))(p^s(c) - c) - 2f(p^s(c)) \tag{37}$$

Solving for $p^s(c) - c$ from the first order condition and rearranging terms, the second order condition is equivalent to

$$\frac{f'(p^s(c))[1 - F(p^s(c))]}{[f(p^s(c))]^2} + 2 > 0 \tag{38}$$

Totally differentiating the first order condition for $p^s(c)$ we find that $d/dc[p_s(c)] > 0$, so that the optimal take it or leave it price is a strictly monotonically increasing function of the seller's valuation of it, c. This

means $p^s(c)$ is an invertible function. Further, the second order condition in equation (38) above must hold for each c, so let $c_0 = p^{s-1}(c)$. For this value of c_0 the second order condition reduces to

$$\frac{f'(c)[1-F(c)]}{[f(c)]^2} + 2 > 0. \tag{39}$$

However this implies that the virtual valuation is a strictly montonically increasing function of the valuation, c, since the derivative of the virtual valuation is

$$d/dc[v(c)] = 2 - \frac{1-F(c)}{f(c)^2}f'(c). \tag{40}$$

If the virtual valuation is strictly monotonic, then it is invertible and the optimal reservation price is given by equation (33) above. Recall that $v^{-1}(c)$ is the take it or leave it price the seller charges. So a buyer with valuation r will accept this take it or leave it price if and only if $r \geq v^{-1}(c)$. But by monotonicity of the virtual valuation, this is equivalent to

$$v(r) \geq c \tag{41}$$

which is the condition Myerson derived for awarding the object to the buyer (see the formula and discussion in the middle of page 70 of Myerson's article).

We summarize this discussion with the following

Lemma: (Isomorphism between optimal auction and optimal take it or leave it price when) *The optimal take it or leave it price (i.e. the solution to optimization problem in equation (35) above) is equivalent to the seller's reservation price in the optimal auction with a single buyer. Furthermore, the second order optimality conditions imply that a buyer will accept the optimal take it or leave it price if and only if the sufficient condition is satisfied for awarding the object to the buyer in Myerson's optimal auction.*

A1.3 Discussion and Caveats

We have shown that under some assumptions about the nature of the bargaining stage game between a customer and the middleman, the problem is isomorphic to the bargaining problem between a buyer and a seller of an indivisible object that Myerson considered in his "Optimal Auction Design" paper. Of course, if we relax some of our assumptions, then the isomorphism is no longer exact and it is not necessarily the case that the optimal selling strategy for the middleman in the bargaining stage game is to quote a take it or leave it price p^r given in equation (22). For example if some customers are repeat customers, it will likely be optimal for the seller to condition on past information about the customer and take this into account in setting a price. We have simply assumed that either that there are no repeat customers, or

that the middleman does not "remember" their past sales. Neither of these assumptions is true for the steel middleman considered in our empirical analysis. The middleman does quote take it or leave it prices to his customers, but at the same time, he also promises quantity discounts and customer rebates when sales over a specific interval of time (e.g. a quarter or a year) are sufficiently high. Clearly, these pricing schemes are designed to encourage "customer loyalty" and are not included in our analysis here.

A key assumption underlying our "isomorphism" result is the assumption that the customer knows the middleman's shadow value of inventory. This would require knowledge of the middleman's inventory q_t and the wholesale price that the middleman can purchase steel at, p_t, and it seems unlikely that a customer would know this information. So the bargaining between the middleman and the customer is likely to be a two-sided game of incomplete information.

Another assumption is that there is no simultaneous bargaining or pricing over price and quantity. The steel middleman offers quantity discounts for purchases on any specific occasion (i.e. based on only the current purchase ocasiopn and independent of the quantity purchased over a specified period of time as discussed above). However if the middleman does not have sufficient quantity on hand to fill the customer's desired order size, there may be bargaining in the form of additioinal discounts to compensate for the fact that the buyer will have to engage in additional search to acquire the residual quantity $q^d - q$ that could not be filled by the middleman. For the products we study, it is rarely the case that customers are not able to buy the full amount they demanded (i.e. where an ensuring stockout occurred on the next business day). It is much more likely that when there was not sufficient stock on hand to satisfy the customer's desired demand quantity, the middleman acquires additional inventory on short notice so that the full order could be met, even if this meant paying a substantial premium to acquire the residual quantity from a nearby competing middleman.

Thus, while we can imagine that once we relax several of our restrictive assumptions, there may be alternative selling/pricing strategies in the bargaining stage game than the rather simple take it or leave it pricing strategy that we have considered here. However we think that the use of such strategies is quite reasonable, and while there may be some modified version of the pricing rule that take into account additiona information (e.g. the past history of purchases) it is hard to imagine an alternative selling mechanism that would yield significantly higher expected profits than a take it or leave it pricing strategy, appropriately adjusted to accomodate the extra information (e.g. on the history of purchases for repeat customers) that we have not considered here. Our objectve was to appeal to Myerson's seminal work to justify why a take it or leave it pricing strategy is a reasonable thing to do. It would appear to be a much more difficult

problem to characterize the optimal selling strategy under weaker assumptions, yet it is not at all clear that there is an obviously superior mechanism.

There is also a very practical reason why it appears nearly optimal to use take it or leave it price quotes: time constraints. Multiple round bargaining and haggling can be very time and energy consuming. The middleman trades in over 2000 separate steel products. If the middlman devoted more than a few minutes to bargaining to negotiate a mutually acceptable price, there would be a need to hire additional sales agents to assist in the bargaining/sales process. Yet Myerson's work suggest that there is no clear advantage to bargaining, and given that it can considerably redeuce the time and energy necessary to consummate a transaction there seems to a second, independent reason why it may be a good idea to precommitt not to engage in bargaining and haggling over price.

One final argument against the empirical relevance of quoting take it or leave it prices to customers is that it may be hard for a firm to precommitt to this strategy if the potential surplus per transaction is very high. When one precommitts to a take it or leave it price quote, one has to accept that there is a significant chance that the customer will "walk" even though the likelihood that there is significant positive surplus to be had from transacting with this customer. In different industries, such as automobiles, where the number of transactions made per middleman (e.g. individual car salesman) is much fewer, it may be significantly more difficult to precommitt to letting the customer walk out the door if they did not find the salesman's first price quote sufficiently attractive.

Thus, while we are aware that the assumptions under which the optimal take it or leave it price can be shown to optimal via an isomorphism with Myerson''s optimal auction design paper are rather strong and unrealistic, we think it is worthwhile to point out the result since it gives added support and intuition about why a take it or leave it price may be an optimal thing to do under certain conditions.