

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1228

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

AN EMPIRICAL MODEL OF INVENTORY INVESTMENT
BY DURABLE COMMODITY INTERMEDIARIES

George Hall and John Rust

June 1999

An Empirical Model of Inventory Investment by Durable Commodity Intermediaries

George Hall and John Rust*
Yale University

June 1999

Abstract

This paper introduces a new detailed data set of high-frequency observations on inventory investment by a U.S. steel wholesaler. Our analysis of these data leads to six main conclusions: orders and sales are made infrequently; orders are more volatile than sales; order sizes vary considerably; there is substantial high-frequency variation in the firm's sales prices; inventory/sales ratios are unstable; and there are occasional stockouts. We model the firm generically as a *durable commodity intermediary* that engages in commodity price speculation. We demonstrate that the firm's inventory investment behavior at the product level is well approximated by an *optimal trading strategy* from the solution to a nonlinear dynamic programming problem with two continuous state variables and one continuous control variable that is subject to frequently binding inequality constraints. We show that the optimal trading strategy is a *generalized (S, s) rule*. That is, whenever the firm's inventory level q falls below the *order threshold* $s(p)$ the firm places an order of size $S(p) - q$ in order to attain a *target inventory level* $S(p)$ satisfying $S(p) \geq s(p)$, where p is the current spot price at which the firm can purchase unlimited amounts of the commodity after incurring a fixed order cost K . We show that the (S, s) bands are decreasing functions of p , capturing the basic intuition of commodity price speculation, namely, that it is optimal for the firm to hold higher inventories when the spot price is low than when it is high in order to profit from "buying low and selling high." We simulate a calibrated version of this model and show that the simulated data exhibit the key features of inventory investment we observe in the data.

Keywords: commodities, inventories, dynamic programming

JEL classification: D21, E22

*Corresponding author: John Rust, Department of Economics, Yale University, 37 Hillhouse Avenue, P.O. Box 208264 New Haven CT 06520-8264, phone: (203) 432-3569, fax: (203) 432-6323, e-mail: jrust@gemini.econ.yale.edu. This paper was prepared for the *Carnegie-Rochester Conference Series on Public Policy*. John Rust is grateful to the Alfred P. Sloan Foundation for research support. We thank numerous seminar participants for helpful comments, particularly Terry Fitzgerald, James Kahn, Valerie Ramey, and Herbert Scarf.

1 Introduction

This paper formulates and solves a dynamic model of optimal inventory investment by durable commodity intermediaries. Commodity intermediaries are companies whose business is to stock-pile quantities of homogeneous durable goods such as steel, lumber, coal, etc. These firms do minimal production processing, and make profits via inter-temporal speculation, purchasing bulk quantities of durable commodities at competitive world spot market prices and subsequently selling their inventories to customers at a mark-up.

We study a new database from one such intermediary, a U.S. steel wholesaler or, in industry lingo, a steel service center. This firm has offered us a unique opportunity to undertake detailed observations of its operations on an ongoing basis by providing us with daily data on purchases and sales of each of the 2,200+ separate products that it sells. We know the firm's initial inventory holdings starting in July 1, 1997, allowing us to calculate inventory holdings for each product on a daily basis for 20 consecutive months. We also have highly confidential data on the identities of each of the firm's customers, the prices they were charged, and the quantities they purchased. Our analysis of these data yields six main conclusions:

1. Orders and sales are made infrequently.
2. Orders are more volatile than sales.
3. There is considerable variability in order levels.
4. There is no stable inventory/sales relationship.
5. Inventory stockouts and near stockouts occur regularly, especially during regimes of low inventory holdings.
6. There is considerable high-frequency variation in sales prices.

We observe all six facts at the individual product level. We observe facts 2, 3, and 6 at the firm level. To explain these facts we solve a dynamic programming model which treats each product as an independent "profit center". Using this model we ask whether the firm's behavior can be accurately approximated by the optimal trading strategy implied by the model's solution.

In the model, the spot price $\{p_t\}$ of the commodity 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^o to order at the spot price p_t . There is a fixed transaction cost K for placing any order, so the firm will only place sufficiently large orders for which the incremental change in expected profits exceeds K . In all other respects we model the firm as behaving passively. That is, we assume that the firm does not attempt to bargain with customers or price discriminate. Instead the firm quotes an exogenously specified markup over the current spot price p_t , and receives a random realized demand q_t^d which is filled on a “first come, first served” basis subject to the constraint that quantity sold cannot exceed stock on hand $q_t + q_t^o$.

The firm’s optimal speculative investment strategy is the solution to an infinite horizon dynamic programming problem. This problem is isomorphic to the problem of optimal inventory management that has been extensively studied in the Operations Research literature. Although a number of existing models in this literature allow the costs of “producing” new inventory to evolve stochastically, we are not aware of a previous study that is directly relevant to the problem faced by speculative investor or a durable commodity intermediary who has the ability to purchase (versus produce) new inventory at a constant marginal cost p_t which changes stochastically from period to period according to a Markov transition density $g(p_{t+1}|p_t)$.

The fact that our model involves a non-convex fixed transaction cost (adjustment cost) K suggests that the most directly relevant predecessor to our work is the theory of optimal inventory investment developed by Arrow, Harris and Marschak (1951) and Scarf (1960). Extending a classic result by Scarf (1960) characterizing the optimal inventory investment strategy as an (S, s) rule, Hall and Rust (1999) proved that the optimal inventory investment strategy continues to take the (S, s) form when the spot price p_t represents the marginal cost of production that evolves stochastically. In this case the optimal solution takes the form of a *generalized (S, s) rule* in which S and s are functions of p . The function $s(p)$ is the *order threshold* and the function $S(p)$ is the *target inventory level* satisfying $S(p) \geq s(p)$. Under an (S, s) rule, the optimal order size is zero whenever the current inventory level q exceeds $s(p)$. However when q falls below $s(p)$ the firm places an order of size $S(p) - q$, restoring inventory levels to the target level $S(p)$. The magnitude of the gap between s and S depends on the magnitude of fixed costs of ordering new inventories: if $K = 0$ then $s(p) = S(p)$, otherwise $s(p) < S(p)$.

In our example both $s(p)$ and $S(p)$ are decreasing functions of p , capturing the basic intuition of commodity price speculation, namely, that it is optimal for the firm to hold higher inventories when the spot price is low than when it is high. In effect it is a prescription for how best to

profit from “buying low and selling high.” Under the optimal policy the firm exploits low spot order price opportunities by making large purchases. The firm can make capital gains on its inventory holdings once the price rises. However the firm faces a risk that if prices remain low for a protracted period, some or all of its expected speculative profits will be dissipated by the interest opportunity costs and physical costs of storing the commodity. Interest opportunity costs are an increasing function of the spot price of steel. Further, demand tends to be lower when prices are high. This implies that both $S(p)$ and $s(p)$ are small when p is high, reflecting the firm’s desire to maintain a relatively low level of inventories when demand is low and holding costs are high. As a result when p is high, q is relatively small and stockouts occasionally occur. Via a numerical simulation, we show that our simple model of optimal commodity price speculation implies the key stylized facts of inventory investment that we observe in the steel data. In particular, we find that in our simulated data set orders are infrequent, order quantities are more variable than sales, inventory/sales ratios vary dramatically, stock-outs occur when spot prices are high, and inventory holdings follow “saw-tooth” trajectories similar to those we observe for individual steel products.

While the main focus of this paper is to explain the high-frequency behavior of a single firm, the issues addressed may be of interest to economists studying movements of aggregates at lower, particularly business cycle, frequencies. In general, recessions can be characterized as periods of inventory liquidations. While in the U.S. inventory investment averages less than one-half of one percent of GDP, during a typical recession the reduction in inventory investment accounts arithmetically for about 50 percent of the reduction in GDP (Ramey and West, 1997). So if we want to understand business cycles, it is important to understand inventory investment behavior, and as we show below, commodity intermediaries account for a large share of aggregate inventory investment.

In the U.S., commodity intermediaries are classified in the merchant wholesale trade sector of the economy (SIC Major Groups 50 and 51). As a group, the wholesale trade sector comprises between 6.5 and 7.0 percent of GDP, and this sector holds about 26% of the total outstanding stock of inventories.¹ The wholesale trade sector is decomposed into a durable goods sector (SIC Major Group 50) and a nondurable goods sector (SIC Major Group 51). About 2/3 of the stock of wholesale trade inventories are held by establishments within the durable goods sector, with

¹The remaining stock of inventories is held by either manufacturers or retailers.

the remaining 1/3 held by establishments in the nondurable goods sector.

Steel service centers are classified within the durable goods sector of wholesale trade.² There are 5,000 such firms located throughout the U.S. with a high concentration in the Great Lakes region. These firms currently hold between 7 and 8 million tons of steel in inventory. Out of the 127 million tons of steel consumed in the U.S. in 1998, about 29 million tons (23 percent) was shipped through steel service centers. This makes steel service centers the largest single customer group of the ultimate suppliers, the steel mills.

Section 2 provides a brief review of the existing literature on inventory investment. Section 3 presents the steel inventory data and summarizes the six main conclusions from our empirical analysis that we will attempt to explain with a simple dynamic programming model of inventory investment. Section 4 presents the model. Section 5 displays numerically computed solutions and stochastic simulations of a calibrated example of the model. Section 6 compares our firm level data to more aggregated data. Section 7 summarizes our findings.

2 Background

There is an extensive literature on the role of commodity storage from an aggregate perspective (see, e.g. Working, 1949 and Williams and Wright, 1991); however we are unaware of more detailed micro-oriented studies of individual agents participating in these markets. Although the main ideas behind the role of commodity storage have been around for many years, only relatively recently have economists attempted to deduce the implications of this model for commodity prices. A stylized version of the dynamic rational expectations commodity storage model, (e.g. Deaton and Laroque, 1992 or Miranda and Rui, 1997) posits that the aggregate supply of a commodity is produced inelastically, with the supply evolving according to some stochastic process $\{z_t\}$. There is a stationary demand function $D(p)$, so in the absence of storage, equilibrium prices evolve according to the stochastic process $\{D^{-1}(z_t)\}$. However if we assume a storage technology exists with a “convenience yield” $c_t = c(x_t)$ (equal to the immediate benefit from having one additional unit of the commodity in storage net of the costs of storing it, where x_t is a vector of state variables affecting the costs and benefits of storage), then competition by commodity intermediaries and

²The four-digit SIC code for steel service centers is 5051; their NAICS code is 42151.

speculators should cause prices to satisfy the equation

$$p_t = \max \left[D^{-1}(z_t), c(x_t) + \beta E\{p_{t+1} | p_t, x_t, z_t\} \right], \quad (1)$$

where $\beta = 1/(1 + r)$. This functional equation (1) defines $p_t = p(x_t, z_t)$ as the unique fixed point to a contraction mapping. Deaton and Laroque, Miranda and Rui and others have solved this functional equation numerically and have analyzed the implications of storage for the time series behavior of commodity prices. This work has shown that many of the observed properties of commodity prices, including skewness, occasional price spikes (i.e. sharp price increases as opposed to price decreases), and high autocorrelations can be explained as a result of competitive storage even if the fundamental “forcing process” $\{z_t\}$ is assumed to be independent and identically distributed *IID*. To our knowledge there is no “microfoundations” derivation for the intertemporal equilibrium relationship (1). However the argument is that if prices did not satisfy this relationship, speculators would buy or sell the commodity to equate current and expected future prices net of storage/carrying costs. We attempt to cast some insight into this using our micro model of commodity intermediaries. According to the functional equation, price spikes occur during aggregate stockouts; otherwise speculators succeed in stabilizing prices, preventing sharp increases or decreases in commodity prices during times of production surpluses or shortages. The theory suggests that sudden crashes in commodity prices should not occur, since this would induce speculators to purchase and store the commodity for subsequent resale.

The steel service center we study is precisely one of the “speculators” implicit in the commodity storage model. However the recent collapse in commodity prices in the aftermath of the 1997 Asian financial crisis calls into question the power of inventory speculation in preventing the steep price declines that occurred during 1998. The physical costs of storing commodities such as steel are presumably very small and the rate of depreciation of steel is close to zero. However the interest opportunity costs of storing these commodities can be substantial, a fact that seems to have been overlooked in the commodity storage literature. It is reasonable to suppose that speculators will not buy large quantities of a commodity in the aftermath of a price crash if they expect it to be followed by a sustained recession that would limit their ability to resell the commodity at attractive prices in the future. This observation underscores the importance of extending the commodity storage model by building more detailed models of the speculators underlying these models, including the commodity intermediaries we study here.

There is also an extensive literature of macro-level models of inventories which assume short-run increasing marginal costs to holding inventories. The workhorse model of this literature is the linear quadratic (LQ) model introduced by Holt, Modigliani, Muth and Simon (1960). The standard LQ model implies that production (orders) should be smoother than sales. Since this implication is almost always rejected empirically, a variety of modifications have been made. For example many authors augment these models with an “accelerator term” in the profit function which is essentially a quadratic penalty function from deviating from a fixed “target” inventory/sales ratio. This target is treated as an unknown parameter to be estimated (e.g. Blanchard, 1983; West, 1986; and Kashyap and Wilcox, 1993). Kahn (1987, 1992) justifies an inventory/sales ratio target by explicitly incorporating costly stock-outs. Bils and Kahn (1996) further justify targeting such a ratio by modeling sales as an increasing function of the available inventories. A second modification is to assume that firms operate on flat or even decreasing regions of their short-run marginal cost curves. Ramey (1991), Bresnahan and Ramey (1994), and Hall (1997) provide evidence that firms may often operate in such regions. Third, Blinder (1986b) and Miron and Zeldes (1988) and others have added cost shocks in the form of input price shocks, while others such as Eichenbaum (1984, 1989) have added cost shocks in the form of unobservable technology shocks. In these cost-shock models inventories are used to smooth production costs rather than the level of production. These modifications have improved the ability of the LQ model to explain aggregate inventory time series, although as we will show in the next section it has some severe handicaps in its ability to explain our product-level data.

Dynamic micro-level models of inventory investment incorporating a fixed cost to ordering were pioneered by Arrow, Harris, and Marschak (1951) and Scarf (1959). Scarf was the first to prove that the optimal policy is of the (S, s) form. In the simplest case, the firm chooses an order limit point s , and an upper inventory point S . The firm places no orders until inventories fall to s or below, whereupon the firm places an order to reset the inventory level to S . Blinder (1981), Caplin (1985), and Fisher and Hornstein (1998) argue that explicitly modeling fixed costs at the firm level helps explain inventory behavior at the aggregate level.

Despite extensive research in the area of inventory investment, a satisfactory model to explain this important time series has not yet been written down and solved. Even models which appear capable of explaining the basic features of the data have clear flaws. For example attempts to estimate macro models of inventory investment often yield parameter estimates of the wrong sign.

Some of the problems may stem from a lack of high-quality data on production and inventories. Fair (1989) suggests that the observation that production is more volatile than sales is just a figment of poorly constructed data. Miron and Zeldes (1989) demonstrate that there is substantial measurement error in both the monthly manufacturing and inventory investment data. The absence of high quality inventory data at the macro-level motivates us to study this issue at the firm level. In their survey of the inventory literature for the *Handbook of Macroeconomics* Ramey and West (1997) “advocate more plant and firm-level studies, although gathering such data requires substantial work.” (p. 47).

3 Data

A U.S. steel service center (referred to below as the “firm”) provided us detailed data on every transaction it undertook between July 1, 1997 to February 26, 1999 (434 business days) for each of the 2200+ individual products that it sells. 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. The firm’s records provide data on the level of inventories for each product at the beginning and end of each month. Using the inventory accumulation identity we can track the firm’s inventory holdings for each day within the month. Also since we observe the prices at which this firm buys and sells steel, we also have a near-perfect measure of the mark-ups charged to customers. Finally since we meet regularly with company executives, we are able to eliminate any discrepancies in the transaction and inventory data. This is an exceptionally clean dataset.

The firm records transactions on the day the steel either enters or leaves one of its warehouses. Although the firm does receive some commitments for sales in advance, most of their sales are delivered within 24 hours of the commitment, and 95 percent of their orders are filled within five days. Indeed, the primary service this wholesaler provides is having the goods on hand and being able to deliver them to the customer on short notice. While back-orders do occasionally occur, we study products which customers can assume the firm will have on hand. We do not have data on when the firm makes an order to purchase steel. From discussion with company executives we know that some of their orders are made weeks in advance (up to 12 weeks when purchasing foreign steel), while some purchases are made with only a day or two notice. In this paper we

assume the relevant time period is one business day.

Although this company offers over 2200 products, tables 1 and 2 provide summary statistics for prices and quantities for eighteen of their most important products which are considered baseline products within the industry. These products serve as key indicators from which the prices of other products are calculated, and display the characteristic features that we see for many other products. For reasons that will become clear subsequently, these products are also of interest because none involve any actual production processing beyond storage and resale. Finally, we chose relatively high volume products for which the firm made at least four orders during the sample period. Figure 1 plots an indicator of the firm's aggregate inventory holdings, the sum (in pounds) of the inventories for each of these eighteen products. Figure 2 plots the inventory/sales ratio measured as "days supply" which we define as the level of current inventories divided by the average daily sales rate for the previous 30 business days.³ Figures 5 - 16 plot daily time series for inventories, days-supply, and spot order and sales prices, for products 2, 4 and 13 in tables 1 and 2. These figures also contain three dimensional scatterplots of purchase quantities as a function of current inventory and order prices.

Our analysis of these data can be summarized in six main conclusions:

1. *Orders and sales are made infrequently.* In the second column of table 1, we report the number of days in which each product enters one of the firm's warehouses. We have selected some of the highest volume products this firm deals in; nevertheless, orders are rarely made. Sales are made more frequently as can be seen from column (5) of table 1 and from the absence of long flat segments in the inventory graphs. However even for product 2, the product with the most frequent sales, sales are made less than 3/4 of the days in the sample. Note also that the periodicity between successive orders is highly variable.
2. *Orders are more volatile than sales.* The last column in the bottom row of table 2 reports the ratio of the standard deviation of aggregate orders to the standard deviation of aggregate sales. This ratio is 9.2, which shows that for this firm orders are substantially more volatile than sales. Columns (2), (4), (6), and (8) of table 2 report the unconditional means and standard deviation of orders and sales. But since sales and orders are made infrequently, we

³Computing days-supply using future sales instead of past sales does not change the qualitative features of any of the graphs in this paper.

also report in columns (3), (5), (7), and (9) the means and standard deviation conditional on an order or sale occurring. Not surprisingly, since orders are made less frequently than sales, the average order size is larger than the average sales size. As is found in many other studies, the standard deviation of orders is larger than the standard deviation of sales. This holds for all eighteen products; see column (10). Note that extremely large sales are rare events as can be seen from the relatively small number of large discontinuous downward jumps in inventory levels in the time series plots.

3. *There is considerable variability in order levels.* In table 2, we can see that for all but four of the eighteen products, conditional on an order occurring, the standard deviation of order size (column (5)) is larger than mean order size (column (3)). This fact can also be seen graphically in our plots of the data for products 2, 4, and 13 in figures 5 - 16. Figures 5, 9, and 13 display the time path of the inventory holdings for these three products. These figures display a “saw-tooth” pattern for inventory holdings with intervals during which inventory levels steadily decrease due to sales to customers punctuated by periodic large orders that replenish inventory holdings. Thus, inventory holdings can be characterized as a jump process with a negative drift due to numerous small sales, and periodic discontinuous upward jumps due to a relatively small number of large orders.

However the firm also makes many small orders. This is apparent in figures 6, 10, and 14, which display scatterplots of order size as a function of current inventory holding and the order price. In general, these three graphs illustrate that the lower the price and the lower the level of inventories, the larger the order. But a striking feature of these figures is the number of small orders – especially when inventories and the order price are high. Also note that in figure 6 most of the orders for product 2 lie in a band between 19.00 and 19.50. The tendency for order size to increase rapidly as a function of order price suggests that the firm’s demand for product 2 is highly elastic. This suggests that inventory holdings are quite sensitive to the spot price of steel, a conclusion that is confirmed from an inspection of the time series for inventories and order prices in figures 5 and 7, 9 and 11, and 13 and 15, respectively. Comparing these graphs vertically, we see that the biggest upward jumps in inventories generally occur when the (interpolated) order price series hits historical lows. However our ability to make clear inferences about this is hampered by the fact that we only

observe spot prices for these products on the days the firm places orders for steel. Thus we cannot be sure that the actual spot price series may actually have been even lower between the successive dates at which large purchases occurred. However we have indirect evidence of the importance of price shocks from aggregate price indices such as the example displayed in figure 3. At least for the last three quarters of 1998, the steady decrease in steel prices are matched by steady increases in inventory levels as we can see from figure 4 which plots the inventory/sales ratios for several independent measures of carbon plate (i.e our firm level data and aggregate industry holdings of carbon plate).

4. *There is no stable inventory/sales relationship.* Figures 8, 12, and 16 display the inventory/sales ratio in terms of days-supply. As in the case of the aggregate days-supply series, these three inventory/sales ratios fluctuate widely and in the case of products 4 and 13 appear to have multiple "regimes" with high and low inventory/sales ratios.

This finding is not inconsistent with the well-documented fact in the inventory literature that there is considerable persistence in the deviations in the inventory-sales relationship from its long-run mean (e.g. Feldstein and Auerbach, 1976; Blinder, 1986a; and Ramey and West, 1997). The mean of the days-supply series of the firm's aggregate inventory holdings plotted in figure 2 is 66 days. So for the first 240 business days of the sample the firm is below its long-run mean, and for the second 200 business days the firm is above its long-run mean. This could be interpreted as considerable persistence in the inventory-sales relationship; however for reasons we discuss below it does not appear that the firm is targeting a constant inventory-to-sales ratio and just slowly adjusting toward it.

5. *Inventory stockouts and near stockouts occur regularly, especially during regimes of low inventory holdings.* From figures 8, 12, and 16, we can see that the firm often allows inventories to fall to a level below 5 days worth of sales. Moreover, for product 13, the firm was completely stocked-out (i.e. had zero inventories) for 24 days during the time period.
6. *There is considerable high-frequency variation in the sales price, with large changes in sales prices on successive sale dates.* This firm is clearly charging some customers higher prices than others, a fact readily acknowledged by company executives. While we do not attempt to model the firm's pricing decisions in this paper, this feature of the data motivates our

desire to do future work analyzing dynamic models of endogenous price setting and price discrimination. See Athreya (1999) for an exploratory empirical analysis of the determinants of price variation among different customers, products, and time periods.

We now consider whether any of the standard models of inventories outlined in section 2 are capable of explaining the six main facts listed above.

1. **(S, s) models.** The saw-tooth pattern of the inventory series is clearly reminiscent of an (S, s) policy, which also predicts intervals of steady declining inventories (due to sales to customers) interspersed by occasional upward jumps in inventories (due to new orders by the firm). While the saw-tooth pattern of inventory holdings in figure 1 is suggestive of an (S, s) policy, closer analysis reveals that the firm's behavior cannot possibly be described by a standard (S, s) rule where S and s are fixed time-invariant constants. Under such a policy the firm places an order of size $S - s$ when its current inventory q falls below the lower order threshold s . This implies that whenever the firm places an order we should see inventories replenished to the same target level S . However it is clear from figure 1 that the amount of inventory the firm holds after each order varies widely over time. Also, in the absence of large discontinuous downward jumps in inventories resulting from large sales (e.g. in limiting continuous-time versions of the (S, s) inventory model where sales follow a diffusion process), all orders should be of the same size $S - s$. It is clear from figure 1 that the size of the firm's orders vary widely over time. Finally, the frequent number of stockouts also casts doubt on the empirical validity of the continuous time diffusion version of the (S, s) rule, which predicts that in the absence of jumps in the demand process that with probability 1 inventories will remain in the interval (s, S) . When there are positive fixed costs of ordering, $s > 0$, and the only way inventories can fall below this level is if there are discontinuous jumps in demand. On the other hand, if fixed costs of ordering inventories were 0, then the firm should place new orders each day to maintain the target inventory level S . In either case stockouts should not occur under the standard (S, s) model. Thus, we conclude that this firm's behavior is inconsistent with the predictions of the standard (S, s) inventory model.
2. **Production smoothing models.** Our finding that orders are on average 9 times more variable than sales shows that this firm's behavior is inconsistent with the predictions of

standard production-smoothing models. These models imply that the variance of production should be lower than the variance of sales. Of course, one can question the relevance of the production smoothing model for studying the behavior of this firm since it does a minimal amount of actual production processing. Although this firm does have a small assembly line that “levels” steel coil (i.e. it unwinds the coil and chops it into rectangular sheets), the firm’s main “production” activity for many of its other products such as heavy steel plate and pipe simply involves placing new orders to replace inventory at a time-varying “marginal cost” p_t , the spot price of steel on day t . There are no costs of stopping, idling, and restarting the “assembly line” for these latter products, so that the theory predicts that there is far less incentive to attempt to smooth production (which in this case simply amounts to placing new orders for steel).⁴ Indeed, to the extent that there are fixed costs to placing orders, it would appear that it is optimal for the firm to do the opposite of production-smoothing, namely to make relatively infrequent large orders rather than frequent small orders. We conclude that the standard versions of the production-smoothing model cannot provide a plausible empirical model for this firm.

3. ***LQ models.*** A particularly popular type of production smoothing model is the *LQ* model, which is the standard framework for modeling inventories in the macro literature. Unfortunately our analysis suggests that the *LQ* model has severe deficiencies at the micro level, particularly for describing the product-level inventory holdings of this firm. The *LQ* model ignores the frequently binding constraint that orders must be non-negative and is therefore unable to explain the observation that on most days orders are zero. Even if we were to interpret the *LQ* model’s predictions of negative orders as representing “desired orders” and use Tobit-style censoring to map negative desired orders to the observed order of zero, we believe that the linear laws of motion for the state variables in *LQ* models would have a hard time approximating the mass point at zero that we observe in the distributions of quantity ordered and sold.

4. ***LQ models with inventory/sales ratio targets.*** In order to explain the widely observed fact that production is more volatile than sales, the standard *LQ* production smoothing mod-

⁴However Abel (1983) finds in a model with a production lag, stock-outs, and endogenous pricing the variance of sales exceeds the variance of production even if the cost of producing are linear.

els have been augmented to include a target inventory/sales ratio and a quadratic penalty for deviating from this target (e.g. Blanchard, 1983). Although the assumption that the firm has a fixed target inventory/sales ratio is not derived from first principles, under certain circumstances tacking on such a term to the firm's cost function yields optimal policies for which production is more variable than sales. However our data provide little support for the hypothesis that the firm has a fixed inventory/sales target. A simple inspection of figure 2 shows that the inventory/sales ratio is extremely variable, beginning with a "low inventory regime" during which the firm has only a month's supply on hand, followed by a "high inventory regime" when it has more than 5 month's supply on hand. While some of the rise in the days-supply series is due to a drop in sales during the last two months of the sample period, much of this dramatic increase appears to be due to significant declines in the spot price of steel over the entire period. In simple terms, this firm appears to be engaging in commodity price speculation, attempting to "buy low and sell high". This strategy implies that the firm should buy large quantities of steel when prices are low, holding it for subsequent resale when prices are higher. Such a strategy is inconsistent with maintaining a fixed inventory/sales ratio.

Our discussions with company executives lead us to conclude that maintaining a stable inventory-to-sales ratio is a rather low priority for the firm. When we asked the general manager whether the firm targeted an inventory-to-sales ratio, he stated that he prefers to carry under 60 days-supply worth of inventory. When we asked him why then he kept making large purchases of steel even when his days-supply exceeded 100 days, he stated that explicit adherence to this rule of thumb would keep him from exploiting good buying opportunities. Perhaps more importantly the only times the general manager or the CEO of the firm discussed inventory-to-sales ratios with us was when we brought it up. It is not a statistic they compute on a regular basis or have in front of them when making purchasing decisions.

Our analysis of the firm's product level data suggests that cost shocks — which in this case are mainly changes in the spot price at which the firm acquires steel inventories — could be the key explanation for the observation that orders are more volatile than sales. A second explanation is the fact that this firm does not do any actual production processing for the products we have

studied, and a third explanation is the existence of positive fixed costs associated with placing new orders for steel. We believe the first explanation is the key to understanding the large variation in inventory holdings over our sample period. As we can appreciate from figure 3 the spot price of steel is likely to be one of the most volatile of the cost shocks facing this firm, whereas the other production and storage costs are unlikely to have varied much over this period. Conversations with company executives do not give us any reason to believe that the fixed costs associated with ordering steel are large, and no reason to suppose that they should have changed over our sample period. Similarly, storage costs appear to have been nearly constant over our sample period. The labor and depreciation costs associated with operating the warehouses in which the steel inventories are stored are small in comparison to the main cost of storage, the opportunity cost of capital as measured by the short term interest rate. The interest rate has been fairly constant over our sample period, and there haven't been any changes in the physical production/storage technology that we are aware of. On the other hand the firm's major "cost of production", the spot price of steel, has declined fairly dramatically for many of its products including carbon plate products as we have seen in figure 3. Many of these price declines are a consequence of reduced world-wide steel demand following the Asian crisis together with possible "dumping" of steel by foreign producers in Russia, Japan, Brazil, and other countries.

More sophisticated econometric and economic modeling is required in order to assess the relative importance of the different explanations of the observation that orders are more volatile than sales. A major problem is created by the fact that we only observe spot prices for the firm's products on the days it placed orders, resulting in infrequent observations of spot prices at irregular time intervals. Due to econometric problems arising from endogenous sampling of these spot price processes, we have been careful not to draw any conclusions about the high frequency behavior of steel prices by simply interpolating our endogenously sampled spot price series. In future work we will develop estimators that correct for this endogenous sampling problem, but in the meantime we have focused our analysis on characterizing the main facts about inventory stocks, orders, and sales for which problems of endogenous sampling problems do not arise. Our analysis has lead us to reject all of the main models that have been used to model inventory behavior in the existing literature.

In the next section we formulate and solve a dynamic programming model of inventory investment by durable commodity intermediaries, in which the optimal policy is a generalization of

the classic (S, s) rule with (S, s) bands that are declining functions of the current spot price of steel. This suggests that many of the stylized facts we have observed for this firm, particularly the observation that orders are more variable than sales and the instability in inventory/sales ratios, could be a consequence of an optimal inventory speculation strategy on the part of the firm. We confirm this in section 5 by presenting simulations of a calibrated version of this model that show that the predicted behavior of this model is strikingly similar to the behavior of this firm. In particular simulated data from the model exhibits 5 of the 6 main features that we have observed in the product level data for this firm.

4 The Model

Our model is in the tradition of the dynamic (S, s) model pioneered by Arrow *et. al.* (1951) and Scarf (1959). We extend their models to allow the spot market price at which the firm purchases the commodity to follow a Markov process. The uncertainty and serial correlation in spot prices imply that a simple (S, s) strategy with fixed S and s thresholds is generally no longer optimal. The optimal inventory investment strategy in our extended model is a function of the spot market price for the commodity p as well as inventory on hand q . However we find that a *generalized (S, s) rule* is optimal. The firm’s optimal trading strategy consists of a pair of *functions* $S(p)$ and $s(p)$ satisfying $s(p) \leq S(p)$. The lower band $s(p)$ is the firm’s *order threshold*, i.e. it is optimal for the firm to order inventory whenever $q \leq s(p)$. The upper band $S(p)$ is the firm’s *target inventory level*, i.e. whenever the firm places an order to replenish its inventory, it orders an amount sufficient to insure that inventory on hand (the sum of the current inventory plus new orders) equals $S(p)$.

Furthermore, the (S, s) bands are generally monotonically declining functions of p , reflecting the simple logic of commodity price speculation, namely to “buy low and sell high”. Low spot prices present an opportunity for the intermediary to make an expected profit by purchasing the commodity when it is cheap, storing it, and subsequently selling it at a higher price. While we assume that the firm could sell the commodity immediately with a positive expected mark-up over the current spot price, most of its profits are obtained from selling the commodity in subsequent periods when the gross of markup prices at which the intermediary sells to its customers have “recovered”. It follows that the firm’s desired holdings of the commodity are larger when spot

prices are low than when spot prices are high.

Under certain circumstances the generalized (S, s) rule takes the form of a “bang-bang” strategy with price “thresholds”: whenever the spot price falls below a price threshold the firm makes a speculative “bet” by placing large orders for steel. This results in large, infrequent discontinuous increases in inventory levels during periods of unusually low “bargain prices” in the spot market, behavior. This behavior is consistent with the observed instabilities and “regime shifts” in the inventory/sales ratio that we observed in our steel intermediary data. It is suboptimal for the intermediary to set a fixed, time-invariant inventory/sales target as is typically assumed in LQ models since this impedes the firm’s ability to profit from buying low and selling high. Indeed when spot prices are sufficiently high the firm’s desired inventory holdings can fall to nearly zero and the incidence of stockouts rises precipitously. The high sales revenues and high opportunity costs of inventory holding during high price “regimes” make it optimal for the firm to liquidate rather than replenish its inventory holdings. Once fully liquidated, the firm optimally chooses to forego inventory investment until spot prices revert to lower levels, even though this comes at a high cost in terms of sacrificed sales revenue and a steep increase in the incidence of stockouts.

We derive these results from a relatively simple dynamic programming model of a generic durable commodity intermediary. These intermediaries do not undertake any physical production processing: their main function is to buy the durable good at spot prices, store it, and sell it subsequently at a markup. We make a number of strong simplifying assumptions about the operations of these intermediaries that we hope to relax in future work. Our first simplification is a *decentralization hypothesis* that allows us to model the inventory investment decisions for each product traded by the intermediary separately. This separation is formally justified under the assumption that there are no technological interdependencies (storage externalities or joint capacity constraints) for the different products the intermediary carries, and the strong assumption that the price processes for the different products are conditionally independent. Under these assumptions it is easy to show that the firm’s multi-product inventory investment problem decomposes into independent subproblems: essentially each individual product becomes a separate sub-firm or “profit center” which can be modeled in isolation from the others.

We need this assumption to break the “curse of dimensionality” associated with the firm’s dynamic programming problem. In the absence of decentralization, a “central planner” within the firm would have to solve a single 4400+ dimensional dynamic programming problem (since each of

the firm’s 2200+ products requires a minimum of two continuous state variables, p and q). Since the complexity of continuous-state and continuous-control DP problems increases exponentially fast in the number of state and control variables, it is clear that such a problem would be far too large to solve using current hardware and software. However under our decentralization hypothesis, the firm’s problem decomposes into 2200+ lower dimensional DP problems, each of which is tractable. Thus the decentralization hypothesis makes it feasible for us to model the entire firm by simply summing the optimal trading rules for each individual product. There are interesting questions about how this firm decentralizes its operations in practice (many of the sales and pricing decisions for individual products are delegated to the firm’s sales agents), but we do not have space here to offer more in depth consideration of these issues.

We abstract from difficult issues connected with modeling endogenous price setting and price discrimination and assume that the firm charges a fixed markup over the current spot price of the commodity. We also abstract from taxes and the details of the firm’s financial policy: these issues will be discussed in more detail below. Finally, we abstract from delivery lags and assume that the firm cannot backlog unfilled orders. Thus, whenever demand exceeds quantity on hand, the residual unfilled demand is lost. This fundamental “opportunity cost” motivates the firm to incur inventory holding costs, even in the absence of any stockout penalty capturing the “goodwill costs” of lost future sales due to alienated customers.

We model the intermediary as making decisions about buying and selling a durable commodity in discrete time. For concreteness, we consider a model with daily time intervals, matching the intervals in our data set. The state variables for the firm are (p_t, q_t) where q_t denotes the inventory on hand at the start of day t , and p_t denotes the per unit spot price at which the intermediary can purchase the commodity at day t . We assume $\{p_t\}$ evolves according to an exogenous Markov process with transition density $g(p_{t+1}|p_t)$. At the start of day t the intermediary observes (p_t, q_t) and places an order $q_t^o \geq 0$ for immediate delivery of the commodity at the current spot price p_t . We assume that the intermediary sets a uniform sales price to its customers, p_t^s , via an exogenously specified markup rule over the current spot price p_t :

$$p_t^s = f(p_t) + \epsilon_t, \quad E\{\epsilon_t|p_t\} = 0. \tag{2}$$

As a first approximation, we assume the firm uses a linear markup rule, $f(p_t) = \alpha_0 + \alpha_1 p_t$, where α_0 and α_1 are positive constants.

After receiving q_t^o and setting p_t^s , the intermediary observes the quantity demanded, q_t^d . We assume that the distribution of q_t^d depends on the spot price p_t , reflecting a stochastic form of downward sloping demand. Let $H(q_t^d|p_t)$ denote the distribution of realized customer demand. We assume that H has support on $[0, \infty)$ with at most one mass point at $q^d = 0$ and is regular in the sense that for any continuous, bounded function G , the function $EG(p, q)$ is a twice continuously differentiable function of its arguments where EG is given by:

$$EG(p, q) = \int G(p, q, q^d)H(dq^d|p). \quad (3)$$

We allow H to have a mass point at 0, reflecting the event that the intermediary receives no customer orders on a given day t . Let $h(q^d|p)$ be the conditional density of sales given that $q^d > 0$. This is a density with support on the interval $(0, \infty)$. Let $\eta(p) = H(0|p)$ be the probability that $q^d = 0$. Then we can write H as follows:

$$H(q^d|p) = \eta(p) + [1 - \eta(p)] \int_0^{q^d} h(q'|p)dq'. \quad (4)$$

As noted above, we assume that there are no delivery lags and unfilled orders are not backlogged. This eliminates the need to carry additional state variables describing the status of pending deliveries and backlogged orders. We also assume that the firm does not behave strategically with regard to its sales to its customers. In addition to charging an exogenously specified markup as in equation (2), the firm does not withhold any inventory for future sale when there is a current demand for it. Thus, we assume that the intermediary meets the entire demand for its product in day t subject to the constraint that it can not sell more than the quantity it has on hand, the sum of beginning period inventory q_t and new orders q_t^o , $q_t + q_t^o$. Thus the intermediary's realized sales to customers in day t , q_t^s , is given by

$$q_t^s = \min [q_t + q_t^o, q_t^d]. \quad (5)$$

We assume the commodity is completely durable and not subject to physical depreciation. Therefore the law of motion for start of period inventory holdings $\{q_t\}$ is given by:

$$q_{t+1} = q_t + q_t^o - q_t^s. \quad (6)$$

Since the quantity demanded has support on the $[0, \infty)$ interval, equation (5) implies that there is always a positive probability of unfilled demand $q_t^s < q_t^d$. We let $\delta(p, q + q^o)$ denote the probability

of this event:

$$\delta(p, q + q^o) = 1 - H(q + q^o|p). \quad (7)$$

Whenever $q_t^d > q_t^s$, equations (5) and (6) imply that a *stockout* occurs, i.e. $q_{t+1} = 0$. Of course, the firm can minimize the probability of a stockout by insuring that quantity on hand, $q + q^o$, is sufficiently high. It is interesting to ask whether it would ever be optimal for the firm to set $q + q^o = 0$, which *maximizes* the probability of a stockout. This can be optimal in our model if spot prices and holding costs are sufficiently high.

We define the intermediary's expected sales revenue $ES(p, q, q^o)$ by:

$$\begin{aligned} ES(p, q, q^o) &= E\{p^s q^s | p, q, q^o\} \\ &= E\{p^s | p\} E\{q^s | p, q, q^o\} \end{aligned} \quad (8)$$

where:

$$E\{p^s | p\} = f(p) \quad (9)$$

and:

$$E\{q^s | p, q, q^o\} = [1 - \eta(p)] \left[\int_0^{q+q^o} q^d h(q^d|p) dq^d + \delta(p, q + q^o)[q + q^o] \right]. \quad (10)$$

A key property to notice about the ES function is that it is symmetric in its q and q^o arguments: from the definitions given above we see that ES can be written as $ES(p, q + q^o)$. Thus, expected sales revenue depends only on the total quantity on hand $q + q^o$, rather than upon the separate values of q and q^o . This symmetry is a key to the proof of the optimality of the generalized (S, s) policy.

We turn now to specifying the per period profit function, which requires some additional assumptions about taxes and the intermediary's financial policy. We appeal to the Modigliani-Miller Theorem to argue that in the absence of taxes, borrowing constraints, and other capital market imperfections, the intermediary's inventory investment policy should be unaffected by its financial policy. This allows us to abstract from the details of how the intermediary actually finances its inventory holdings and allows us to conclude that regardless of whether its inventory holdings are financed by debt or retained earnings, the intermediary incurs an interest opportunity cost of inventory holdings equal to $r_t p_t (q_t + q_t^o)$ in day t where r_t denotes the spot interest rate at date t . However we model the intermediary as an entrepreneur whose personal intertemporal discount factor is $\beta \in (0, 1)$ which may not equal the current market discount factor $1/(1 + r_t)$.

This implies that the owner would like to borrow when β is less than $1/(1+r_t)$ and lend otherwise. Thus, financial policy does affect the firm's expected discounted profits even in the absence of taxes, borrowing constraints, and other capital market imperfections. Since the steel company will not disclose information about their financial policy, we assume the intermediary finances its inventory holdings out of retained earnings, incurring an opportunity cost of maintaining inventory level q_t equal to $r_t p_t q_t$. Furthermore, we assume r_t is fixed; $r_t = r$ for all t .⁵

We assume the intermediary incurs a cost of ordering inventory given by a function $c^o(q^o)$ which may be discontinuous at $q^o = 0$ but is twice continuously differentiable for $q^o > 0$. The discontinuity in c^o at $q^o = 0$ reflects possible fixed costs of placing orders. For concreteness, we will assume a simple fixed order cost,

$$c^o(q^o) = \begin{cases} F & \text{if } q^o > 0 \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

This specification could be easily generalized to account for per unit shipping costs and quantity discounts. However in order to derive the optimality of a generalized (S, s) policy we need to assume that the derivative of c^o is constant for $q^o > 0$. For simplicity we assume this derivative is 0 in what follows below.

We assume that the intermediary incurs a physical storage cost $c^h(q)$ of holding inventory level q , where c^h is nondecreasing and twice continuously differentiable. The intermediary perceives a "goodwill cost" $\gamma \geq 0$, where γ represents the present value of lost profits from customers who switch to alternative suppliers in the event that $q^d > q + q^o$. Finally the intermediary has a maximum storage capacity equal to $\bar{q} \leq \infty$. Thus the intermediary's single-period profits π is given by:

$$\pi(p_t, p_t^s, q_t^s, q_t, q_t^o) = p_t^s q_t^s - r p_t (q_t + q_t^o) - c^o(q_t^o) - c^h(q_t + q_t^o) - p_t q_t^o - \gamma I\{q_t^s = q_t + q_t^o\}. \quad (12)$$

Notice that our assumptions imply that the profit function π is symmetrical in its q_t and q_t^o arguments and can be written as $\pi(p_t, p_t^s, q_t^s, q_t + q_t^o)$.

The intermediary's inventory investment behavior is governed by the decision rule:

$$q_t^o = q^o(p_t, q_t), \quad (13)$$

⁵The assumption of constant interest rates can be easily relaxed as far as the theoretical presentation of the model is concerned, however it does lead to an extra state variable that complicates the numerical solution of the model. In future work we plan to include r_t as a state variable to study the sensitivity of inventories to changes in interest rates, a topic of interest in studies of the role of inventories in macroeconomic fluctuations.

where the function q^o is the solution to:

$$V(p_t, q_t) = \max_{q^o} E \left\{ \sum_{j=t}^{\infty} \beta^{(j-t)} \pi(p_j, p_j^s, q_j^s, q_j^o + q_j^s) \middle| p_t, q_t \right\}, \quad (14)$$

The value function $V(p, q)$ is given by the unique solution to Bellman's equation:

$$V(p, q) = \max_{0 \leq q^o \leq \bar{q} - q} [W(p, q + q^o) - pq^o - c^o(q^o)], \quad (15)$$

where:

$$W(p, q) \equiv [ES(p, q) - rpq - c^h(q) - \gamma\delta(p, q + q^o) + \beta EV(p, q)]. \quad (16)$$

and EV denotes the conditional expectation of V given by:

$$\begin{aligned} EV(p, q) &= \eta(p) \int_{p'} V(p', q) g(p'|p) dp' \\ &+ [1 - \eta(p)] \delta(p, q) \int_{p'} V(p', 0) g(p'|p) dp' \\ &+ [1 - \eta(p)] \int_{p'} \int_0^q V(p', q - q') h(q'|p) g(p'|p) dq' dp'. \end{aligned} \quad (17)$$

The optimal decision rule $q^o(p, q)$ is given by:

$$q^o(p, q) = \underset{0 \leq q^o \leq \bar{q} - q}{\operatorname{argmax}} [W(p, q + q^o) - pq^o - c^o(q^o)], \quad (18)$$

Hall and Rust (1999) proved the following theorem, which includes Scarf's (1960) characterization of the optimality of (S, s) as a special case. The key to the proof is the same as in Scarf's (1960) theorem, the property of K -concavity.

Definition: A function $f : R^+ \rightarrow R$ is K -concave if and only if for all $q \in R^+$ and all $z \geq 0$ and all $b \geq 0$ satisfying $q - b \geq 0$ we have:

$$f(q + z) - K \leq f(q) + \frac{z}{b} [f(q) - f(q - b)]. \quad (19)$$

Intuitively, a (nonconcave) function is K -concave if the secant approximation to $f(q + z)$ given on the right hand side of equation (19) exceeds $f(q + z)$ less the constant K . Clearly a concave function is 0-concave, and thus K -concave for all $K \geq 0$. Hall and Rust (1999) prove that if $W(p, q + q^o) - pq^o$ is K -concave in q^o , then the optimal inventory policy is an (S, s) rule. So it suffices to establish sufficient conditions for $W(p, q + q^o) - pq^o$ to be K -concave. There are two key lemmas that are used to establish this result: 1) the Bellman operator maps K -concave functions

into K -concave functions, and 2) pointwise limits of K -concave functions are K -concave. For details of the proof of the following Theorem, see Hall and Rust (1999).

Theorem: Consider the function $W(p, q + q^o)$ defined in equation (16), where W is defined in terms of the unique solution V to Bellman's equation (15). If $W(p, q + q^o)$ is a K -concave function of q^o for any p , then the firm's optimal inventory investment policy $q^o(p, q)$ takes the form of a generalized (S, s) rule. That is, there exist a pair of functions (S, s) satisfying $S(p) \geq s(p)$ where $S(p)$ is the desired or target inventory level and $s(p)$ is the inventory order threshold, i.e.

$$q^o(p, q) = \begin{cases} 0 & \text{if } q \geq s(p) \\ S(p) - q & \text{otherwise} \end{cases} \quad (20)$$

where $S(p)$ is given by:

$$S(p) = \operatorname{argmax}_{0 \leq q^o \leq \bar{q} - q} [W(p, q^o) - pq^o] \quad (21)$$

and the lower inventory order limit, $s(p)$ is the value of q that makes the firm indifferent between ordering and not ordering more inventory:

$$s(p) = \inf_{q \geq 0} \{q | W(p, q) - pq \geq W(p, S(p)) - pS(p) - F\}. \quad (22)$$

We conclude this section by noting that if the decision rule take the form of a generalized (S, s) policy, the value function is linear in q with slope equal to p when $q < s(p)$. To see this, we simply substitute the form of the generalized (S, s) policy (20) into the formula for V in Bellman's equation (15) to obtain:

$$V(p, q) = \begin{cases} W(p, S(p)) - p[S(p) - q] - F & \text{if } q \leq s(p) \\ W(p, q) & \text{otherwise} \end{cases} \quad (23)$$

Thus, V takes the form $V(p, q) = \gamma(p) + pq$ for $q \leq s(p)$, which shows that the "shadow price" of an extra unit of inventory is p . The intuition for this simple result is straightforward: if the firm has an extra unit of q when $q \leq s(p)$ then it needs to order one fewer unit in order to attain its target inventory level $S(p)$. The savings from ordering one fewer unit of inventory is simply the current spot price of the commodity, p . When $q > s(p)$ the shadow price of inventory is no longer equal to p . We do know that since $q = S(p)$ maximizes $W(p, q) - pq$, we must have $\partial W(p, q)/\partial q = p$ when $q = S(p)$. If W were strictly concave, $\partial W(p, q)/\partial q > p$ when $q \in (s(p), S(p)]$ and $\partial W(p, q)/\partial q < p$ when $q \in (S(p), \bar{q}]$. Thus, there is a kink in V function at the inventory order threshold, $q = s(p)$. which is inconsistent with the assumption that W is strictly concave in q . However the result does hold under the weaker condition that W is K -concave in q .

5 A Calibrated Example

To illustrate the behavior implied by our model we solved a discrete approximation of (15) numerically under the following assumptions. We assumed that the daily interest rate is time-invariant and equal to $r = .05/261$.⁶ We assumed the firm uses the sales price markup rule $p_t^s = 0.9 + 1.06p_t$ and spot prices $\{p_t\}$ evolve according to a truncated lognormal $AR(1)$ process:

$$\log(p_{t+1}) = \mu_p + \lambda_p \log(p_t) + \epsilon_t \quad (24)$$

where $\mu_p = .06$, $\lambda_p = .98$, and $\{\epsilon_t\}$ is an *i.i.d.* $N(0, \sigma_p^2)$ sequence, with $\sigma_p^2 = 8.6510^{-5}$. The upper and lower truncation bounds on this process were chosen to be (16, 25) which are beyond the minimum and maximum spot purchase prices observed in our sample or in long run simulations of the untruncated version of this process.

We choose the function form and parameters values for the price process to qualitatively match the histograms of the transaction prices generated by simulations from our model with the histograms of the transaction prices observed in the data. Recall that we only observe prices on days the firm purchases steel so we have infrequently and irregularly sampled price series in which the sampling is made by the firm endogenously. Since estimators that correct for this endogenous sampling problem do not exist, we fit the price process visually rather than employing a formal econometric criterion. The current parameterization of equation (24) yields an order price process with an invariant distribution with mean of 20.5 cents per pound and a standard deviation of 1.00 cents per pound. Given the markup rule, the mean and standard deviation of the sell price process are 22.6 and 1.06, respectively. The means of these price processes are in the range of means reported in table 1. The standard deviations are below those reported in table 1; but again, we are silent on the issue of price discrimination.

We assumed that quantity demanded, q_t^d , is a mixed truncated lognormal distribution conditional on p_t . That is, with probability .5 $q_t^d = 0$, and with probability .5 q_t^d is a draw from a truncated lognormal distribution with location parameter $\mu_q(p) = 4.43 - .7 \log(p_t)$ and standard deviation parameter $\sigma_q = 1.081$.⁷ These parameters yield a stationary distribution for q_t^d (conditional on $q_t^d > 0$) with conditional mean equal to 18.3 and conditional standard deviation equal

⁶We assumed there are $365 - (2 \times 52)$ business days in a year.

⁷This downward sloping demand curve in conjunction with the linear markup rule implies $\frac{p_t^s}{p_t}$ decreases as p_t increases and q_t decreases. So markups are procyclical. We have experimented with various forms of the markup rule and have found that the basic results of the model do not depend on procyclical markups.

to 28.2. The units of the quantity variables are in 1,000's of pounds. The first two moments of the quantity demanded process are in the range of the moments reported in columns (7) and (9) in table 2.

We assumed that goodwill costs of stockouts γ and physical holding costs were zero, $c^h(q_t) = 0$, and that the fixed order cost is equal to \$50, i.e. $c^o(0) = 0$ and $c^o(q^o) = \$50$ if $q^o > 0$. Finally, we assumed that the firm owner's personal subjective discount factor was given (on a daily basis) by $\beta = 1/(1 + .05/261)$; so $\beta = 1/(1 + r)$.

We solved for the optimal inventory investment rule by the method of policy function iteration which computes a discrete approximation to the value function $V(p, q)$ as the unique fixed point to the Bellman equation, (15). We used a uniform discretization of the (p, q) state space to approximate the continuous DP problem by the solution to a finite state problem with 750 grid points (15 in the p dimension and 50 in the q dimension). The grid points are evenly spaced along the p dimension. Along the q dimension, the distance between the grid points increases as q increases. Thus the grid point are more densely spaced in the region where there is more curvature in the decision rule. Although the state variables were discretized, we treated the control variable q^o as a continuous variable subject to the constraint that $0 \leq q^o \leq \bar{q} - q$. Policy iteration is not guaranteed to converge in continuous choice problems such as this one; but for this example, the algorithm converged in 39 iterations. Using the values computed at these 750 grid points we produced continuous approximations to the value function and decision rule via multi-linear interpolation.

As can be seen from Bellman's equation (15), the policy improvement step requires the solution of a constrained optimization problem involving the two functions $ES(p, q)$ and $EV(p, q)$, each of which is a conditional expectation of functions of two continuous variables (sales, $p^s q^s$, and the value function, $V(p, q)$). Since no analytic solutions to these conditional expectations exist, we resorted to numerical integration. We experimented with two different methods of numerical integration, a "quadrature" approach that approximates EV by a probability weighted sum:

$$\hat{E}V(p, q) = \frac{1}{N_p} \frac{1}{N_q} \sum_{i=1}^{N_p} \sum_{j=1}^{N_q} I\{q_j \leq q\} \hat{V}(p_i, q - q_j) h(p_j|p_i) g(p_i|p) \quad (25)$$

where $h(q_j|p_i)$ is a discretized approximation to the conditional probability density $h(q|p, q)$ and $g(p_i|p)$ is a discretized approximation to the transition probability density $g(p'|p)$. Further adjustments to this formula were made in order to allow $\hat{E}V(p, q)$ to account for mass points on

stockouts and zero sales as in equation (17). A second method of approximating EV was a “quasi monte carlo, probability integral transform method” (MC-PIT) given by

$$\hat{EV}(p, q) = \frac{1}{N} \sum_{i=1}^N \hat{V}(\tilde{p}_i, q - \tilde{q}_i) \quad (26)$$

where $\{\tilde{p}_i, \tilde{q}_i\}$ are draws from the density $h(q'|p', q)g(p'|p)$ computed from uniformly distributed draws $\{\tilde{u}_{1,i}, \tilde{u}_{2,i}\}$ from the unit square, $[0, 1]^2$ via the probability integral transform method. Instead of using pseudo-random random draws for $\{\tilde{u}_{1,i}, \tilde{u}_{2,i}\}$ we obtained acceleration using *Generalized Faure sequences*, also known as *Tezuka sequences*. Using number-theoretic methods (see, e.g. Neiderreiter 1992, or Tezuka, 1995), one can prove that for certain classes of integrands, the convergence of monte carlo methods based on deterministic *low discrepancy sequences* is $O(\log(N)^d/N)$ where d is the dimension of the integrand and N is the number of points. This rate of convergence dominates the rate of convergence of carlo methods converge at rate $O_p(1/\sqrt{N})$. These favorable rates of convergence have been observed in practice, see e.g. Papageorgiou and Traub (1997).⁸ The density $h(q'|p, q)$ is the conditional density of q' given that $q' \leq q$,

$$h(q'|p, q) = \frac{h(q'|p)}{1 - \delta(p, q)} \quad (27)$$

where $\delta(p, q) = \Pr\{q' > q|p\} = 1 - H(q|p)$. As in the quadrature method, we adjusted the MC-PIT formula (26) to account for mass points corresponding to stockouts and zero sales. We found that the optimal order size q^o was sensitive to the way the functions ES and EV are approximated. It was critical to use methods that provide accurate approximations both their levels and their derivatives, since the latter determine the first order conditions for a constrained optimum for q^o . In regions where $EV(p, q)$ is nearly linear in q , small inaccuracies in the estimated derivatives can create oscillations in the approximations to EV causing the approximate solution to have multiple local maxima. Without very careful safeguarding of the uni-dimensional optimization algorithm for computing optimal order size, the algorithm could get stuck on a local maximum, generating large instabilities in the estimated value of q^o . The solutions are also sensitive to the discretization of the p and q axes, and the number of points used in the discretization. Through a fair amount of experimentation we have developed numerical procedures that we trust. In particular the two different approximation methods for computing \hat{ES} and \hat{EV} discussed above produced nearly identical results.

⁸We are grateful to Joseph Traub for providing the `FINDER` software co-authored with A.F. Papageorgiou that generated the low discrepancy sequences used in this study.

Figures 17-20 present the optimal decision rule q^o as a function of p and q and the associated expected sales, value functions, and (S, s) bands. Note that our solution technique does not exploit our prior knowledge about the form of the decision rule. The computed value function is nearly a linear function of q . At low inventory levels (in regions the firm is expecting to buy steel), $V(p, q)$ is decreasing in p , whereas at high values of q , (in regions the firm is expecting to not buy but just sell steel) V is increasing in p . The kink at $s(p)$ is not apparent at this level of resolution. These results are consistent with the discussion in the previous section. The optimal decision rule is decreasing in both p and q , although it generally decreases faster in p than in q . In particular when $q^o(p, q) > 0$, $\partial q^o(p, q)/\partial q = -1$ which is consistent with the prediction of the generalized (S, s) rule that $q^o(p, q) = S(p) - q$.

Figure 19 shows the generalized $(S(p), s(p))$ bands implied by our model. The graph of the function s is the curve on the (q, p) plane where the $q^o(p, q)$ surface intersects the plane at $q^o = 0$. The graph of S , is the curve on the (q, p) plane where the $q^o(p, q)$ surface intersects the plane at $q = 0$. These bands are plotted in figure 20 to make it easier to compare them. Due to the fixed costs of ordering (\$50), the $S(p)$ band is strictly above the $s(p)$ band although the difference between the two bands decreases as the price increases. In other words, the order size at s is a decreasing function of the price.

As can be seen in figure 20, when the price is near the lower truncation price bound (16 cents per pound), the constraint on maximum storage capacity \bar{q} (5 million pounds) becomes binding.⁹ This makes sense because the firm knows prices cannot go any lower. The firm cannot make a capital loss on any steel purchased at the lower bound. Nevertheless, this boundary issue is not a major concern since the firm very rarely ever observes prices in this region; this region is over 4 standard deviations away from the mean of the price process.

Figures 21- 24 present the results from a single stochastic simulation of the DP model for 434 periods. At first glance, the simulated series look quite similar to the actual data. Figure 21 shows the time series for inventory levels, and there appears to be multiple regimes. During the first 275 days of the simulation, inventory levels are centered around 200,000 pounds. This average level matches the average level of inventory holdings for product 13 and the first 200 days of product 4. Starting around day 275, the firm enters a “high inventory regime” with the simulated inventory levels reaching a peak over 1,500,000 pounds. This peak is consistent with observed levels of

⁹In figure 19 we plot the decision rule for prices between 17 and 24 to make the graph more easily readable.

inventories for products 2 and 4. During this high inventory regime, days supply reaches almost 350. The transition from the low to high inventory regime occurs when the order price falls below a threshold value. Later, as prices begin to rise (from day 360 to the end of the simulation) the firm lets its inventory holdings gradually fall.

The high- and low-regime property of the optimal inventory holdings can be seen from the decision rule, $q^o(q, p)$. In figure 19, $q^o(q, p)$ is sharply decreasing in p when $q^o(q, p) > 0$. This occurs for two reasons. First, the firm takes advantage of low order prices to build up inventories knowing that it is likely to capture a capital gain on its inventory holdings when prices rise. Second, the firm faces a downward sloping demand curve for its product; so when the price falls, q^d rises and the firm will hold more inventories to accommodate the increase in demand.

The simulation results are consistent with this intuition. Figure 23 presents the censored and uncensored order and sales price series. In this graph, the solid line is the “censored transaction price process” analogous to the one we observe in our dataset. For convenience, we superimposed a linear interpolation of the times and prices at which simulated orders took place on the underlying uncensored “latent price process” $\{p_t\}$ which is plotted as a dotted line in figure 23. Under an optimal ordering strategy, the firm appears to have an uncanny ability to predict turning points in spot prices, with most orders occurring at local minimum points of the realized trajectory for $\{p_t\}$. When prices hit a record low around days 285 and 360, the firm placed several very large orders that ushered it into a “high inventory regime” between days 260 and 434.

In this simulation the firm sold steel on 210 days at average price of 22.67 during the simulation period and purchased steel on 26 days at an average price of 20.04. The average order size was 116,000 pounds with a conditional standard deviation of 62.3. These implied moments from the model are consistent with the moments we observe in the data. Finally the ratio of the standard deviation of orders to the standard deviation of sales for this simulation is 2.4. So the model does imply that orders are more volatile than sales. The particular realization we presented is typical, and not designed to make our model look good. Longer simulations also generate similar results.

These results are qualitatively similar to the actual inventory time series for our firm in figures 5-16. Our DP model display regime shifts in the inventory levels and days supply of inventory with little evidence of a single fixed inventory/sales target; however, we have not systematically searched over the parameter space to ensure that our DP model captures the full volatility and magnitude in these regime shifts. In our individual product data, we also see very large increases

inventory levels occurring when prices hit what appear to be record lows. But we do not see the either very large or very small individual orders. In particular the large increase in inventories around day 350 is spread across four orders. Moreover comparing figures 6, 10, and 14 with figure 22, we see that the DP model generates fewer small size orders than we observe in the data. This suggests that perhaps the fixed order cost is too large; however when we set the fixed cost to zero, we get the counterfactual result that with prices are high, the firm tightly matches orders to sales, ordering almost every period an amount equal to last period's sales. Finally the model does imply occasional stockouts. In the simulation, the firm stocked out on day 108 when quantity demanded was unusually large (over 1 million pounds) and current inventories were relatively low (around 250,000 pounds).

We conclude that cost shocks in the form of serially correlated spot prices in the steel market is the principal explanation for the observed volatility in inventory/sales ratios and the fact that orders are more volatile than sales. We believe this simple model provides a promising starting point for more rigorous estimation and testing using more advanced econometric methods.

6 Aggregation

It is natural to ask whether the firm we study is representative of other durable commodity intermediaries. We address this issue in figure 3 which presents a monthly price index for carbon plate constructed by *Purchasing Magazine*. The data are from January, 1987 to February, 1999. We deflated this index by the PPI-all commodities so the units are in 1982 cents per pound.¹⁰ Note that at the end of the sample the real price of carbon plate is at (at least) a twelve-year low.

Figure 4 plots the firm's days-supply for product 2, a specific type of carbon plate. We also plot the aggregate days-supply of carbon plate for member firms of the Steel Service Center Institute (SSCI).¹¹ Finally we plot the days supply for establishments in the SIC 505 sector (wholesale trade: metal and minerals, except petroleum). All three data series are monthly, and we plot three-month centered-moving averages. Since the mean of the SIC 505 data is one half the mean of SSCI and individual firm data, the scale for the SSCI and firm-level data is the left-hand side axis, and the scale for the SIC 505 data is one the right-hand side axis.

¹⁰Deflating this price index by the CPI would not change any of analysis presented below.

¹¹The SSCI is an industry organization which among other things collects data on member firms' shipments and inventory holdings.

For the sample period of our firm-level dataset (July, 1997 to February, 1999) the more aggregated data appear to be consistent with both our firm level data and the implications of our theory. During this period carbon plate prices fell to record lows and inventory levels at all three levels of aggregation rose significantly. This suggests that the firm’s strategy of placing speculative bets is not atypical of metal wholesalers. We would observe similar results if we were to aggregate the simulated inventory holdings of different simulated firms. While there are idiosyncratic demand shocks that will be averaged out over firms in the simulation, their behavior is affected in a similar way by the common “cost shock” $\{p_t\}$. To the extent that these price series are affected by macroeconomic factors such as the Asian crisis, we have a simple explanation for the role of inventory investment as a propagating mechanism in the business cycle. It would not be difficult to add other “macro shocks” to our model. For example, rather than allowing the interest to be constant, we could allow $\{r_t\}$ to evolve stochastically, say according to a Markov process. We would then be able to study the impact of monetary policy on inventory investment, determining features such as the interest elasticity of inventory investment. This is a topic for future work, however.

We note that the aggregate data present several interesting challenges to try to explain using the model developed in this paper. For example the large swings observed in price of carbon plate seem superficially at odds with the predictions of our model and the commodity storage literature more generally. In particular the latter literature implies that the price process should satisfy the arbitrage condition in equation (1). Our model implies a similar condition

$$p = \frac{\partial ES}{\partial q}(p, S(p)) - rp - \frac{\partial c^h}{\partial q}(S(p)) + \beta \frac{\partial EV}{\partial q}(p, S(p)). \quad (28)$$

The first terms $\partial ES(p, S(p))/\partial q - rp - \partial c^h(S(p))/\partial q$ constitute the “convenience yield” net of holding costs of adding an extra unit of inventory, the analog of the term $c(x_t)$ in the commodity storage model in equation (1). In our case, the convenience yield equals the increase in expected sales of having an extra unit of inventory and the holding costs are the sum of the interest opportunity costs rp plus the marginal physical holding costs $\partial c^h(S(p))/\partial q$. The second term, $\beta \partial EV(p, S(p))/\partial q$, is the expected discounted shadow price of an extra unit of inventory. However as we noted above, V is essentially linear in q with slope p , so $\beta \partial EV(p, S(p))/\partial q$ is the analog of the term $\beta E\{p_{t+1}|p_t, x_t, z_t\}$ in equation (1). Large swings in prices in and of themselves do not contradict either (1) or (28), but intermediaries such as the one we study should tend to dampen

price swings by buying when prices are low and selling off accumulated inventory when prices are high.

It is striking to note that even with 5,000 steel service centers in the U.S., each one presumably solving a dynamic programming problem similar to one presented above, the real price of carbon plate rose 70 percent from early-1987 to mid-1988 only to fall 50 percent by mid-1992. A very puzzling feature is that during the 1988-1989 period prices for carbon steel hit a record high – but so did days-supply both at the steel service center industry level and at the three digit SIC level. According to our model, if intermediaries viewed the prices during this period as being in a temporary “high price regime”, they should have been reducing rather than increasing their inventory holdings. Furthermore during the early 1990s as price fell, so did days supply, a result that is also hard to explain using our model. Of course there may have been demand shocks in the steel market during this period that we are currently unaware of and that might need to be incorporated in a more realistic model. We hope to address these issues more fully in future work.

7 Concluding Remarks

This paper has presented a new data set containing high quality, high frequency observations on product-level inventory investment by a U.S. steel wholesaler. Our empirical analysis yielded six conclusions about inventory investment and price setting by this firm: 1) orders are more volatile than sales, 2) orders are made infrequently, 3) there is considerable volatility in order levels, 4) there is no stable inventory/sale relationship, 5) there is considerable volatility in sales prices consistent with price discrimination, and 6) inventory stockouts occur relatively frequently, especially during periods of high commodity prices when inventory holdings are low.

We argued that the standard versions of the (S, s) model, production smoothing models, and LQ models with target inventory/sales ratios are incapable of explaining these facts. We introduced a new model of optimal inventory speculation by durable commodity intermediaries and showed that the optimal inventory investment strategy takes the form of a generalized (S, s) rule where the S and s bands are declining functions of the spot price of the commodity. Simulations of a calibrated version of our DP model suggest that the firm’s behavior at the product level can be well approximated by an optimal trading strategy. We employed a continuous-state version of Howard’s policy iteration algorithm to solve a two-dimensional nonlinear infinite horizon dynamic

programming problem with continuous state and control variables that are subject to frequently binding inequality constraints. The predicted behavior from the generalized (S, s) rule appears to explain a number of different aspects of inventory investment behavior by our steel wholesaler, including highly variable inventory/sales ratios and occasional stockouts during low inventory regimes when the spot price for steel is relatively high.

In future work we plan to undertake more rigorous econometric estimation and testing of our generalized (S, s) model which will account for difficult problems of “dynamic selectivity bias” arising from endogenous sampling of the prices at which the firm purchases inventory. We also plan to extend the model to allow for additional state and control variables such as the firm’s sales price p_t and the interest rate r_t . The former will allow us to study endogenous price determination and price discrimination, whereas the latter will allow us to study the impact of monetary policy on inventory investment as a potential propagating mechanism for business cycles. In doing so, we will need to address some difficult issues connected with the curse of dimensionality underlying the solution of high dimensional DP problems such as the one considered in our paper. Recent progress in this area by Rust (1997, 1998) and Rust, Traub, and Woźniakowski (1998) make us optimistic about the prospect for solving these larger and more realistic models.

In future work we plan to develop more realistic models that relax some of the strong simplifying assumptions in our model. This includes our assumption that the current spot price p_t is a sufficient statistic for the distribution of per period retail demand. We want to allow for the impact of “macro demand shocks” and the possibility that the firm’s demand in period t , q_t^d also depends on its realized value in previous periods. More ambitiously, we would like to model equilibrium determination of prices in durable commodities markets with three different types of agents: producers, retail consumers, and intermediaries. We want to determine whether the fundamental functional equation in the rational expectations commodity price model of Williams and Wright, equation (1), can be derived from microfoundations in a market where informational frictions and transactions costs lead to considerable price dispersion and potential market inefficiency despite the standard nature of the product.

References

- [1] Abel, A. (1983) “Inventories, Stock-Outs and Production Smoothing” *Review of Economic Studies*, 52: 283-293.
- [2] Arrow, K.J., Harris, T., and Marschak, J.(1951) “Optimal Inventory Policy” *Econometrica*, 19-3: 250–272.
- [3] Athreya, R. (1999) “An Empirical Study of Price Discrimination by a Durable Commodity Intermediary” manuscript, Department of Economics, Yale University.
- [4] Bills, M., and Kahn J. (1996) “What Inventory Behavior Tells Us about Business Cycles” manuscript, University of Rochester.
- [5] Blanchard, O. (1983) “The Production and Inventory Behavior of the American Automobile Industry” *Journal of Political Economy*, 91: 365-400.
- [6] Blinder, A. (1981) “Retail Inventory Investment and Business Fluctuations” *Brookings Papers on Economic Activity*, 2: 443-505.
- [7] Blinder, A. (1986a) “More on the Speed of Adjustment in Inventory Models” *Journal of Money, Credit and Banking*, 18: 355-365.
- [8] Blinder, A. (1986b) “Can the Production Smoothing Model of Inventory Behavior be Saved?” *Quarterly Journal of Economics*, 101: 431-453.
- [9] Bresnahan, T., and Ramey, V. (1994) “Output Fluctuations at the Plant Level” *Quarterly Journal of Economics*, 109: 593-624.
- [10] Caplin, A. (1985) “The Variability of Aggregate Demand with (S, s) Inventory Policies” *Econometrica*, 53: 1395-1409.
- [11] Deaton, A., and Laroque G. (1992) “On the Behavior of Commodity Prices” *Review of Economic Studies* 59: 1-23.
- [12] Eichenbaum, M. (1984) “Rational Expectations and the Smoothing Properties of Inventories of Finished Goods” *Journal of Monetary Economics*, 14: 71-96.

- [13] Eichenbaum, M. (1989) "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment" *American Economic Review*, 79: 853-64.
- [14] Fair, R. (1989) "The Production-Smoothing Model is Alive and Well" *Journal of Monetary Economics*, 24: 353-370.
- [15] Feldstein, M., and Auerbach, A. (1976) "Inventory Behavior in Durable Goods Manufacturing: The Target-Adjustment Model" *Brookings Papers on Economic Activity*, 2: 351-396.
- [16] Fisher, J., and Hornstein, A. (1998) " (S, s) Inventory Policies in General Equilibrium" manuscript, Federal Reserve Bank of Chicago.
- [17] Hall, G. (1997) "Nonconvex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants" Cowles Foundation Discussion Paper 1169.
- [18] Hall, G., and Rust, J. (1999) "Commodity Storage at the Firm Level", manuscript, Yale University.
- [19] Holt, C., Modigliani, F., Muth, J., and Simon, H. (1960) *Planning Production, Inventories and Work Force*. Englewood Cliffs, N.J.:Prentice-Hall.
- [20] Judd, K. (1998) *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- [21] Kahn, J. (1987) "Inventories and the Volatility of Production" *American Economic Review*, 77: 667-679.
- [22] Kahn, J. (1992) "Why is Production more Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory Holdings" *Quarterly Journal of Economics*, 107: 481-510.
- [23] Kashyap, A., and Wilcox, D. (1993) "Production and Inventory Control at the General Motors Corporation During the 1920's and 1930's" *American Economic Review*, 83: 383-401.
- [24] Miranda, M., and Rui, X. (1997) "An Empirical Reassessment of the Nonlinear Rational Expectations Commodity Storage Model" manuscript, Ohio State University, forthcoming, *Review of Economic Studies*

- [25] Miron, J., and Zeldes, S. (1988) “Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories” *Econometrica*, 56: 877-908.
- [26] Miron, J., and Zeldes, S. (1989) “Production, Sales, and the Change in Inventories: An Identity That Doesn’t Add Up” *Journal of Monetary Economics*, 24: 31-51.
- [27] Neiderreiter, H. (1992) *Random Number Generation and Quasi-Monte Carlo Methods*. CBMS-NSF Regional Conference Series in Applied Mathematics, Philadelphia, SIAM.
- [28] Papageorgiou, A.F., and Traub, J.F. (1997) “Faster Evaluation of Multidimensional Integrals” *Computational Physics*, 11: 574–578.
- [29] Ramey, V. (1991) “Nonconvex Costs and the Behavior of Inventories” *Journal of Political Economy*, 99: 306-334.
- [30] Ramey, V. and K. West (1997) “Inventories” NBER working paper 6315, December, forthcoming in the *Handbook of Macroeconomics*.
- [31] Rust, J. (1997a) “Using Randomization to Break the Curse of Dimensionality” *Econometrica*, 65: 487-516.
- [32] Rust, J. (1997b) “A Comparison of Policy Iteration Methods for Solving Continuous-State, Infinite-Horizon Markovian Decision Problems Using Random, Quasi-random, and Deterministic Discretizations” manuscript, copies available at Economics Working Paper Archive <http://econwpa.wustl.edu/eprints/comp/papers/9704/9704001.abs>
- [33] Rust, J., Traub, J., and Woźniakowski, H. (1998) “No Curse of Dimensionality for Contraction Fixed Points Even in the Worst Case” manuscript.
- [34] Scarf, H. (1959) “The Optimality of (S, s) Policies in the Dynamic Inventory Problem” In *Mathematical Methods in the Social Sciences*. eds. K. Arrow, S. Karlin and P. Suppes. Stanford, CA: Stanford University Press.
- [35] Tezuka, S. (1995) *Uniform Random Numbers: Theory and Practice* Dordrecht, Netherlands: Kluwer.

- [36] Van Roy, B. Bertsekas, D.P., Lee, Y., and Tsitsiklis, J.N. (1997) “A Neuro-Dynamic Programming Approach to Retailer Inventory Management” manuscript, MIT Laboratory for Information and Decision Systems.
- [37] West, K. (1986) “A Variance Bounds Test of the Linear Quadratic Inventory Model” *Journal of Political Economy* 94: 374-401.
- [38] Williams, J.C. and B. Wright (1991) *Storage and Commodity Markets.*, New York: Cambridge University Press.
- [39] Working, H. (1949) “Theory of Price and Storage” *American Economic Review* 39: 1254–62.

product (1)	# order		mean		std		# sell		mean		std		std(order price)/	
	days (2)	order price (3)	order price (4)	std (4)	days (5)	sell price (6)	sell price (7)	std (7)	std(order price) (8)					
1	46	20.23	2.80	2.80	213	22.24	1.99	1.40						
2	61	19.54	1.27	1.27	314	22.12	1.01	1.27						
3	4	19.27	0.22	0.22	114	21.93	1.43	0.16						
4	60	19.83	1.41	1.41	286	22.11	1.23	1.15						
5	26	20.20	1.58	1.58	88	22.21	1.21	1.30						
6	38	20.05	1.63	1.63	190	22.64	1.38	1.19						
7	13	20.57	3.48	3.48	46	22.36	1.58	2.21						
8	9	21.01	3.20	3.20	38	23.53	1.01	3.17						
9	23	21.25	2.52	2.52	95	23.63	1.05	2.41						
10	47	21.96	2.88	2.88	176	23.86	1.16	2.49						
11	8	21.98	2.84	2.84	11	23.69	0.75	3.76						
12	21	21.82	2.99	2.99	66	24.14	1.03	2.90						
13	31	21.58	3.10	3.10	97	24.17	1.19	2.61						
14	21	21.44	2.19	2.19	40	24.36	1.47	1.49						
15	24	21.66	2.48	2.48	45	24.53	1.93	1.29						
16	11	20.90	2.56	2.56	15	25.22	1.01	2.52						
17	4	24.78	2.90	2.90	7	25.34	0.68	4.24						
18	5	23.99	0.24	0.24	9	26.71	1.10	0.22						

Table 1: First and Second Moments of Prices

There are 434 business days in the sample period. Column (2) reports the number of days the firm made one or more orders. Likewise column (5) reports the number of days one or more sales were made. Columns (3), (4), (6), and (7) are in cents per pound.

product (1)	mean order (2)	mean (o o>0) (3)	std order (4)	std (o o>0) (5)	mean sale (6)	mean (s s>0) (7)	std sale (8)	std (s s>0) (9)	std(o o>0)/ std(s s>0) (10)
1	8.61	81.43	45.78	118.95	5.99	12.29	10.29	11.82	10.06
2	24.34	173.54	122.75	287.52	19.78	27.23	23.80	24.01	11.97
3	1.70	184.41	18.31	51.44	3.19	12.19	10.04	16.64	3.09
4	25.63	185.78	149.33	365.74	21.49	33.04	39.71	45.21	8.09
5	2.99	50.05	21.12	72.78	2.52	12.45	7.15	11.39	6.39
6	9.33	106.83	47.45	125.36	8.34	19.19	14.21	16.02	7.83
7	1.53	51.15	12.63	54.95	1.55	14.65	5.49	9.73	5.65
8	1.12	54.27	10.26	49.65	1.07	12.98	4.13	7.23	6.87
9	5.05	95.46	30.73	98.04	3.18	14.56	7.43	9.35	10.48
10	14.33	132.64	65.96	158.15	11.06	27.33	19.90	23.14	6.83
11	0.51	27.77	4.08	12.86	0.41	16.41	2.72	5.62	2.29
12	4.33	89.61	28.55	98.28	3.14	21.04	9.23	13.99	7.03
13	6.68	93.80	36.42	103.66	5.78	26.21	17.66	29.74	3.48
14	3.64	75.30	24.00	82.67	2.23	24.26	9.18	19.75	4.19
15	5.50	99.65	35.03	115.60	3.47	33.54	15.90	38.22	3.02
16	2.83	111.98	23.53	92.61	1.03	29.73	6.28	17.59	5.27
17	0.95	102.92	13.94	118.84	0.32	19.60	3.46	20.58	5.78
18	1.56	135.91	22.11	173.77	0.54	26.14	3.72	0.00	∞
aggregate	120.62	274.70	507.52	738.79	95.10	101.90	81.63	80.30	9.20

Table 2: First and Second Moments of Quantities

Columns (2)-(9) are in 1,000's of pounds.




Figure 1: Aggregate inventory holdings for the eighteen products studied.




Figure 2: Aggregate days-supply for the eighteen products studied (in business days).




Figure 3: Price index of carbon plate steel from *Purchasing Magazine* deflated by the PPI.




Figure 4: Three-month moving average of days-supply for product 2 (dashed line), days-supply for aggregate carbon plate of SSCI firms (solid line), and days-supply for all firms in the SIC 505 sector (dotted line). The units for the firm's holding of product 2 and the SSCI companies holdings are on the left-hand side axis; for the SIC 505 sector the units are on the right-hand side axis.




Figure 5: Times series plot of the inventory for product 2.




Figure 6: Size of purchases for product 2 as a function current inventory holdings and the buy price.




Figure 7: Order prices (solid line) and sell prices (dashed line) for product 2. For the order price series, the size of the marker is proportional to the size of the purchase.

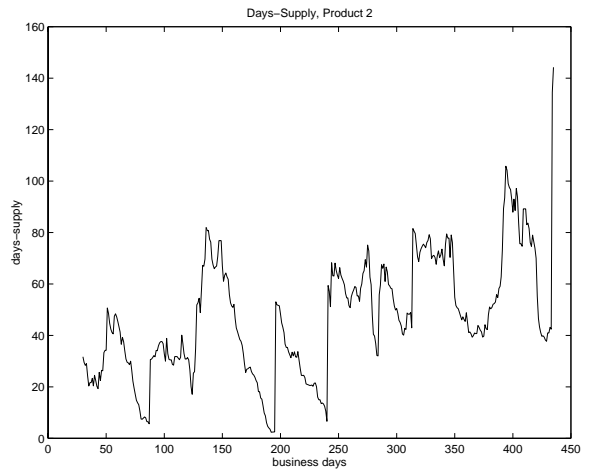


Figure 8: Days-supply of inventory for product 2 (in business days).




Figure 9: Times series plot of the inventory for product 4.




Figure 10: Size of purchases for product 4 as a function current inventory holdings and the buy price.




Figure 11: Order prices (solid line) and sell prices (dashed line) for product 4. For the order price series, the size of the marker is proportional to the size of the purchase.




Figure 12: Days-supply of inventory for product 4 (in business days).




Figure 13: Times series plot of the inventory for product 13.




Figure 14: Size of purchases for product 13 as a function current inventory holdings and the buy price.




Figure 15: Order prices (solid line) and sell prices (dashed line) for product 13. For the order price series, the size of the marker is proportional to the size of the purchase.




Figure 16: Days-supply of inventory for product 13 (in business days).

Expected Sales Function for Inventory Problem




Figure 17: Expected sales revenue, ES for the calibrated example.

Value Function for Inventory Problem




Figure 18: The value function, $V(q, p)$ for the calibrated example.

Optimal Decision Rule for Inventory Problem




Figure 19: Decision rule, $q^o(q, p)$, for the calibrated example.

S-s Bands for Inventory Problem




Figure 20: $S(p)$ and $s(p)$ for the calibrated example.




Figure 21: Simulated inventory holdings




Figure 22: Simulated orders as a function current inventory holdings and buy price.



Figure 23: Censored (solid line) and Uncensored (dotted line) order and sales prices from the simulation.

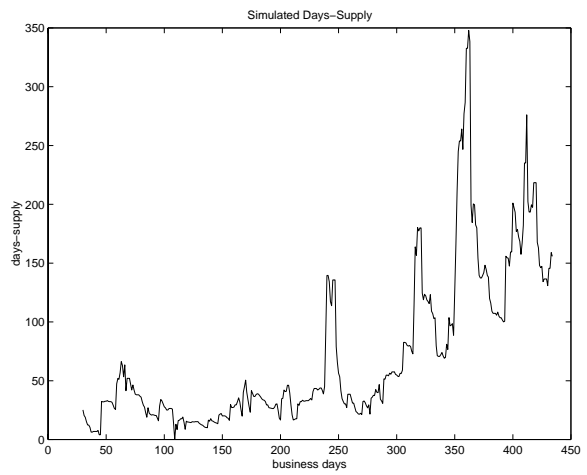


Figure 24: Simulated days-supply of inventory (in business days).