

An (S, s) Model of Commodity Price Speculation

George Hall and John Rust*
Yale University

Revised, November 1999

Abstract

This paper presents a generic theory of speculative inventory investments by *durable commodity intermediaries*. Our model is motivated from an empirical analysis of operations of daily observations on inventories, sales, and purchases of over 2,300 individual products by a U.S. steel wholesaler. 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 considerable day-to-day variability in sales prices; inventory/sales ratios are unstable; and there are occasional stockouts. We derive the firm's optimal trading strategy (on a daily level for individual products under a separability assumption) as the solution to an infinite horizon dynamic programming problem with two continuous state variables and one continuous control variable that is subject to frequently binding inequality constraints. We prove that the optimal trading strategy takes the form of 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, speculation, 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 extends results and arguments presented in Hall and Rust (1999a). We gratefully acknowledge financial support from a National Science Foundation grant, SES-9905145. John Rust is grateful to the Alfred P. Sloan Foundation for financial support of the project "Does Computational Complexity Limit Knowledge in Economics?" We thank numerous seminar participants for helpful comments, particularly Terry Fitzgerald, James Kahn, Valerie Ramey, and Herbert Scarf.

1 Introduction

This paper introduces a theory of optimal inventory investment and price speculation by *durable commodity intermediaries*. These companies, also known as wholesalers, stockpile homogeneous durable goods such as steel, lumber, coal, etc. and typically do only a minimal amount of actual production processing (value added) prior to resale. Instead they make a substantial share of their profits via price speculation, purchasing bulk quantities of durable commodities from producers and other intermediaries and subsequently selling their inventories to retail customers at a mark-up. In very simple terms, the objective of these firms is to “buy low and sell high.” We formulate the firm’s problem as an infinite horizon dynamic programming problem, and characterize the form of its optimal trading strategy. In the process we forge new connections between the previously separate literatures on optimal inventory investment and the rational expectations commodity storage model.

Our theory was motivated by an empirical analysis of a new database from one such intermediary, a U.S. steel supplier, or in industry lingo – a steel service center. This firm supplied us with twenty-eight months of confidential daily data on purchases, sales, and inventory holdings of over 2,300 separate products. 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 considerable day-to-day and within-day variation in sales prices.
5. There is no stable inventory/sales relationship.
6. Inventory stockouts and near stockouts occur regularly, especially during regimes of low inventory holdings.

We observe all six facts at the individual product level. We observe facts 2, 3, and 4 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 from the solution to a dynamic programming problem.

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, we prove 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 $s(p) < S(p)$.

For our model 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 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.

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 and provides sufficient conditions for the optimal inventory investment policy to take the form of an (S, s) rule. 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 makes some concluding comments.

2 Background

The purpose of this section is to make links between the model we introduce in this paper and three different literatures: 1) the market-oriented literature on the rational expectations commodity storage model, 2) the macro-oriented literature on inventory investment, and 3) the micro-oriented literature on optimal inventory theory. We show that a simple generalization of the classical (S, s)

inventory model pioneered by Arrow, Harris, and Marschak (1951) and Scarf (1959) results in a theory of optimal commodity price speculation that can explain the stylized facts of inventory holdings that we observe for our steel intermediary. In addition our model may be a first step towards establishing the micro-foundations for price dynamics in commodities markets populated by speculators and intermediaries.

There is an extensive literature on commodity storage, although all of the studies that we are aware of approach the topic from an aggregate perspective that abstracts from modeling the individual agents in these markets (see, e.g. Working, 1949 and Williams and Wright, 1991). This literature focuses on price determination and abstracts from the determination of quantities, whereas our focus is to provide a micro level model of the evolution of quantities in order to provide a theory of the daily movements in inventories at the product level for our steel firm. However as we will see, there is a close connection between quantities and prices in our model, which provides first step towards a partial equilibrium model in which both prices and quantities are endogenously determined.

While the main ideas behind the role of commodity storage have been around for many years, economists have not attempted to deduce the detailed implications of this model for commodity prices until relatively recently, spurred by dramatic improvements hardware and software that have enabled numerical solution of these models. We briefly review the standard version of the dynamic rational expectations commodity storage model that has been studied by Deaton and Laroque, 1992 and Miranda and Rui, 1997. This model posits that the aggregate supply z_t of a commodity is produced inelastically according to an *IID* process. There is a stationary demand function $D(p)$, so in the absence of storage, equilibrium prices are also an *IID* 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 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 $\beta = 1/(1+r)$. The functional equation (1) is also a contraction mapping with a unique solution $p_t = p(z_t, x_t)$. 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 such as skewness in price distribution, occasional price spikes (i.e. sharp price increases during periods of aggregate stockouts), and high autocorrelations can be explained as a result of competitive storage even if the fundamental “forcing process” $\{z_t\}$ is *IID*.¹ Although the intertemporal equilibrium relationship (1) has not been derived from micro-foundations as an equilibrium outcome resulting from the interaction of producers, end consumers and intermediaries in commodity markets, one can argue 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. Price spikes occur only during aggregate stockouts (when $p_t = D^{-1}(z_t) > c(x_t) + \beta E\{p_{t+1}|x_t, z_t\}$); 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 intermediary that we study in this paper is precisely one of the “speculators” implicit in the commodity storage model. Rather than assume that prices satisfy the functional equation (1), this paper treats the spot price process for the commodity $\{p_t\}$ as an exogenous stochastic process and focuses on deriving the intermediary’s optimal inventory investment strategy. This leads us to a type of “Euler equation” describing the intermediary’s desired inventory holdings that is similar to, although not exactly the same as, the equilibrium condition (1) assumed in the commodity price model. However due to price shocks and transactions costs involved in restocking, the intermediary will rarely be holding its desired level of inventories. As a result, the current spot price p_t will typically be higher or lower than its expected discounted value net of storage costs, $c(x_t) + \beta E\{p_{t+1}|x_t, z_t\}$.

Furthermore, the recent collapse in commodity prices following the 1997 Asian crisis raises further doubts about the empirical relevance of the commodity storage model: one might have expected that actions of speculators should have largely averted the fairly steep sudden decline in prices that we observed. Although we might expect storage to have a limited impact on the

¹Deaton and Laroque had difficulty matching the high serial correlations observed for perishable commodities such as bananas. However Miranda and Rui showed that a simple extension of their framework to allow for a more flexible specification of the “convenience yield” of holding inventories can result in simulated solutions whose autocorrelations match the data.

price path for a perishable commodity following a persistent decline in demand, it is much harder to explain these declines in non-perishable commodities such as steel where the physical costs of storage are very small and the rate of depreciation 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. Moreover, speculators may be unwilling to 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 need to extend the commodity storage model by building more detailed models of the inventory investment decisions of the agents underlying these models — with particular attention to speculation by commodity intermediaries.

This paper also makes contact with an extensive literature on inventory investment. Much of the macro literature has focused on the linear quadratic (LQ) model introduced by Holt, Modigliani, Muth and Simon (1960). The standard LQ model implies that that production should be smoother than sales. In the case of an intermediary, “production” corresponds to purchases of additional inventory on the spot market. As noted above the implication that purchases are smoother than sales is decisively rejected by our data. In order to explain the observation that production (orders) are more variable than sales, a standard trick is to augment the LQ models with an “accelerator term” that 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 targeting an inventory/sales ratio 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 (1999) provide evidence that firms may often operate in such regions. Third, Blinder (1986) and Miron and Zeldes (1988), Durlauf and Mancini (1995) 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. Thus inventories are used to smooth production costs rather than the level of production. These modifications have improved the ability of LQ models to explain aggregate inventory time series, although as we will show in the next section we have doubts about 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). In these models, 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 place 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. Unfortunately, as we show in the next section, plots of our product level inventory data decisively reject the classical versions of the (S, s) inventory model. In particular, we observe the firm placing orders at widely different inventory levels (suggesting that there is no fixed value for s), and observe wide variations in inventory levels following a restocking (suggesting that there is no fixed value for S). However we introduce a simple extension of the (S, s) inventory model that allows the cost of ordering new inventory to depend on the current spot price p_t . We assume that spot prices follow a Markov process and prove that under certain conditions the optimal inventory strategy takes the form of a generalized (S, s) rule where the S and s bands are functions of the current spot price of steel. We show that this simple modification enables us to build a theory of speculative inventory investment that accounts for the main stylized facts of inventory investment behavior that we observe in the product level data for our steel firm.

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

This section provides an initial exploratory analysis of data on the operations of a particular steel service center. This analysis motivates the model introduced in section 3. However we believe that most of the conclusions that we draw from this exploratory analysis, and the model that we develop to explain them, are generic in the sense that they apply to a wider class of commodity intermediaries than steel service centers.

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.² 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 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.³ There are 5,000 such firms located throughout the U.S. with a high concentration in the midwest 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.

One of these steel service centers (referred to below as the “firm”) provided us detailed data on every transaction it undertook between July 1, 1997 to November 8, 1999 (595 business days) for its 2300+ individual products. For each transaction we observe the quantity (number of units and/or weight in pounds) of steel bought or sold, the sales price, the shipping costs, and the identity of the buyer or seller. 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 (both in physical units and in dollar value) 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

²The remaining stock of inventories is held by either manufacturers or retailers.

³The four-digit SIC code for steel service centers is 5051; their NAICS code is 42151.

and inventory data. This is an exceptionally clean dataset.

The firm records transactions on the day the steel either enters or leaves one of the 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 2300 products, tables 1 and 2 provide summary statistics for prices and quantities for eighteen of their more 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 redistribution. 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.⁴ Figures 3 - 14 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.

⁴Computing days-supply using future sales instead of past sales does not change the qualitative features of any of the graphs in this paper.

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 8.4, 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 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 relatively 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 the order size (column (5)) is larger than the 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 3 - 14. Figures 3, 7, and 11 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 4, 8, and 12, 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 in figure 4 that most of the orders for product 2 lie in the *order price* 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 3 and 5, 7 and 9, and 11 and 13, 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.

4. *There is no stable inventory/sales relationship.* Figures 6, 10, and 14 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.
5. *Inventory stockouts and near stockouts occur regularly, especially during regimes of low inventory holdings.* From figures 6, 10, and 14, 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 day-to-day and within-day 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.

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 all 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 validity of the standard (S, s) policy, which predicts that (in the absence of jumps) that inventories should remain in the (interval (s, S)). When there are positive fixed costs of ordering, $s > 0$, so inventories should not fall below this level. 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 8.4 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” of p_t equal to the spot price of steel on day t . There are no costs of stopping, idling,

and restarting an assembly line for these latter products, so it appears that in such case there is far less incentive to attempt to smooth production (which in this case simply amounts to placing new orders for steel).⁵ 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 orders are usually 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 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 models 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. This variation

⁵However 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.

does not appear to be due to non-stationarity in sales, but rather due to significant declines in the spot price of steel over this 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 analysis of the firm’s product level data suggests that cost shocks — which in this case are mainly due to 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. The spot price of steel is clearly 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. Besides the labor and depreciation costs associated with operating the factory building in which the steel inventories are stored, the main cost of storage is 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. 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 simple interpolations of our endogenously sampled spot price series. In Hall and Rust (1999b) we develop estimators that correct for this endogenous sampling problem, but in the current paper 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 next section we formulate and solve a dynamic programming model of speculative inventory investment by durable commodity intermediaries that extends the classic (S, s) inventory investment model by allowing the cost of ordering steel (marginal cost of production) to equal the current price p_t which varies stochastically over time. We prove that the optimal policy in this model takes the form of a generalized (S, s) rule, where the S and s bands 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 qualitatively similar to the behavior of this firm. In particular simulated data from the model exhibit the 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 . We show that under fairly general conditions 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 $S(p)$.

We show that the (S, s) bands are monotonically declining functions of p , reflecting the simple logic of commodity 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.

We derive these results from a relatively simple dynamic programming model of a generic durable commodity intermediary. In our model the intermediary does not undertake any physical production processing: its 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 the intermediary that we intend 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 4600+ dimensional dynamic programming problem (since each of the firm’s 2300+ 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 the decentralization hypothesis, the firm’s problem decomposes into 2300+ lower dimensional DP problems, each of which is tractable. Thus under the decentralization hypothesis it becomes feasible for us to model the entire firm. There are interesting questions about how this firm actually decentralizes its operations in practice (many of the sales and pricing decisions for individual products are delegated to the firm’s sales agents, but their responsibilities are divided by region rather than by particular products), 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 by assuming 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 :

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

For concreteness, in our model below we assume 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 of the

commodity by the intermediary's customers, 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 a mass point at $q^d = 0$, 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:

$$H(q^d|p) = \eta(p) + [1 - \eta(p)] \int_0^{q^d} h(q'|p) dq'. \quad (3)$$

Assumption 1: $h(q|p) > 0$ for all $p > 0$ and $q > 0$.

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

$$q_t^s = \min [q_t + q_t^o, q_t^d]. \quad (4)$$

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

$$q_{t+1} = q_t + q_t^o - q_t^s. \quad (5)$$

Since the quantity demanded has support on the $[0, \infty)$ interval, equation (4) 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:

$$\delta(p, q + q^o) = 1 - H(q + q^o|p). \quad (6)$$

Whenever $q_t^d > q_t^s$, equations (4) and (5) 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:

$$\begin{aligned} ES(p, q, q^o) &= E\{p^s q^s | p, q, q^o\} \\ &= E\{p^s | p\} E\{q^s | p, q, q^o\} \end{aligned} \tag{7}$$

where:

$$E\{p^s | p\} = f(p) \tag{8}$$

and:

$$E\{q^s | p, q, q^o\} = [1 - \eta(p)] \left[\int_0^{q+q^o} q^d h(q^d | p) dq^d + \delta(p, q + q^o)[q + q^o] \right]. \tag{9}$$

A key property to notice about the function ES 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 . The other key property of ES is given in Lemma 1 below.

Lemma 1: $ES(p, q)$ is a strictly monotonically increasing and concave function of q for each $p > 0$.

Proof: From equation (7) it is sufficient to show that $E\{q^s | p, q\}$ is strictly monotonically increasing and concave in q . Taking first and second derivatives we have:

$$\begin{aligned} \partial E\{q^s | p, q\} / \partial q &= [1 - \eta(p)] \delta(p, q) > 0 \\ \partial^2 E\{q^s | p, q\} / \partial q^2 &= -[1 - \eta(p)]^2 h(q | p) < 0. \end{aligned} \tag{10}$$

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 about the way 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 in our model the intermediary is an entrepreneur whose personal intertemporal discount factor $\beta \in (0, 1)$ 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 .⁶

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 assume the firm faces a fixed cost of placing orders:

Assumption 2: *The firm's order cost function c^o is given by:*

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

This specification can be easily modified to account for constant per unit shipping costs. However if there are quantity discounts that lead to a non-constant derivative of c^o with respect to q^o , then a generalized (S, s) policy may not be optimal. For notational simplicity we assume that any per unit shipping costs are already embodied in the spot price p , so that at least in this respect the simplified specification of order costs given in assumption 2 involves no loss of generality.

Assumption 3: *The intermediary incurs a per period physical storage cost $c^h(q)$ of holding inventory level q , where c^h is a strictly convex function of q .*

Assumption 4: *The intermediary has a maximum storage capacity equal to $\bar{q} \leq \infty$.*

Under these assumptions, the intermediary's single-period profits π is given by:

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

⁶The 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 interest rate fluctuations, a topic of interest in studies of the role of inventories in macroeconomic fluctuations.

Notice that our assumptions imply that the profit function π is also 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:

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

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:

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

and EV denotes the conditional expectation of V given by:

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

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

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

Note that we invoke the inf operator in the definition of the optimal decision rule in equation (18) to handle the case where there are multiple maximizing values of q^o . We effectively break the tie in such cases by defining $q^o(p, q)$ as the *smallest* optimal of the optimizing values of q^o .

Definition 0: An (S, s) policy is a decision rule of the form:

$$q^o(p, q) = \begin{cases} 0 & \text{if } q \geq s(p) \\ S(p) - q & \text{otherwise} \end{cases} \quad (19)$$

where S and s are functions satisfying $S(p) \geq s(p)$ for all p .

Candidate functions for the upper and lower bands of the generalized (S, s) policy can be defined in terms of the optimal decision rule $q^o(p, q)$. The upper band $S(p)$ is defined as the optimal order quantity when the firm has no inventory on hand:

$$S(p) = q^o(p, 0). \quad (20)$$

The lower band $s(p)$ is the smallest value of q such that desired inventory investment is 0:

$$s(p) = \inf_{0 \leq q} \{q | q^o(p, q) = 0\}. \quad (21)$$

It is not difficult to show that desired inventory investment at the upper $S(p)$ band is 0: $q^o(p, S(p)) = 0$. Since $s(p)$ is the smallest value of q satisfying $q^o(p, q) = 0$ it follows that $s(p) \leq S(p)$.

The definitions of the (S, s) bands may appear somewhat circular, since $q^o(p, q)$ is defined in terms of (S, s) in equation (19), whereas the (S, s) bands are defined in terms of $q^o(p, q)$ in equations (20) and (21). However there is no circularity since both formulations are equivalent characterizations of the optimal decision rule. We now discuss sufficient conditions for the optimality of a generalized (S, s) rule. There are two key properties needed to establish this result: 1) *symmetry*, and 2) *K-concavity*.

The symmetry property has already been discussed above: it is simply a statement of the fact that the variables q and q^o do not enter as separate arguments in the value function W given in (16): rather they enter as the sum $q + q^o$ as shown in equation (18). The symmetry property is a consequence of our timing assumptions: since new orders of steel arrive instantaneously, the firm's expected sales, inventory holding costs, and expected discounted profits only depend on the sum $q + q^o$, representing inventory on hand at the beginning of the period after new orders q^o have arrived. It follows that if the firm is holding less than its desired level of inventories $S(p_t)$ at the start of day t , it will only have to order the amount $q^o(p, q) = S(p) - q$ in order to achieve its target inventory level $S(p)$. Another way to see this is to note that when it is optimal for the firm to order, the optimal order level solves the first order condition:

$$\partial W(p, q + q^o) / \partial q^o - p = 0. \quad (22)$$

Assuming that W is strictly concave in q (an assumption we will relax shortly), it is clear that there will be a unique value of $q + q^o$ that solves equation (22) for any value of p . Call this solution $S(p)$:

$$\partial W(p, S(p)) / \partial q^o = p. \quad (23)$$

Then we have $q + q^o = S(p)$, or $q^o(p, q) = S(p) - q$.

The function $W(p, q)$ may not necessarily be strictly concave but under assumption 5 below, we can show that W is K -concave as a function of q for each fixed p . Using the K -concavity property we can prove that whenever $q \geq s(p)$, it is not optimal to order: $q^o(p, q) = 0$. When $q < s(p)$ the symmetry property implies that $q^o(p, q) = S(p) - q$ as discussed above. This is the basic argument underlying the proof of the optimality of the generalized (S, s) policy. K -concavity is a generalized form of concavity introduced by Scarf (1960) in his original proof of the optimality of the (S, s) rule.⁷

Definition 1: 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:

$$f(q + z) - K \leq f(q) + \frac{z}{b} [f(q) - f(q - b)]. \quad (24)$$

Intuitively, a (nonconcave) function is K -concave if the secant approximation to $f(q + z)$ given on the right hand side of equation (24) exceeds $f(q + z)$ less the constant K . Clearly a concave function is 0-concave, and thus K -concave for all $K \geq 0$. A basic result, proven in the appendix, is 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. In the appendix we do this in several steps. 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. To prove the first lemma we require the following assumption about the conditional expectation operator:

Assumption 5: The conditional expectation operator EV preserves monotonicity and K -concavity, i.e., if V is non-decreasing and K -concave in q for each p , then $EV(p, q)$ is non-decreasing and K -concave in q for each p where EV is defined in equation (17).⁸

⁷Actually, Scarf introduced the concept of K -convexity and showed how this could be used to establish the optimality of an (S, s) rule when the inventory model was formulated as a cost minimization problem. However just as the negative of a convex function is concave, the negative of a K -convex function is K -concave. In our case we are interested in solving the firm's profit maximization problem. Since costs enter negatively in our formulation, it follows that K -concavity rather than K -convexity is the relevant concept in our case.

⁸We have been unable to find easily verifiable "lower level" assumptions on the conditional distribution of demand, $H(q|p)$, that are sufficient to establish Assumption 5. We simply note that our computational results suggest that Assumption 5 does hold for a wide range of parameter values for our truncated lognormal specification of H as we have verified in the calibrated solutions in the next section. One reason why it is difficult to find simple low level conditions on H for the conditional expectation operator to be K -concavity preserving is that $EV(p, \max[0, q - \bar{q}^d])$ is the expectation of a composition of a convex function $\max[0, q - \bar{q}^d]$ and a K -concave

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). Under assumptions 1-5 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.

$$q^o(p, q) = \begin{cases} 0 & \text{if } q \geq s(p) \\ S(p) - q & \text{otherwise} \end{cases} \quad (25)$$

where $S(p)$ is given by:

$$S(p) = \operatorname{argmax}_{0 \leq q^o \leq \bar{q} - q} [W(p, q^o) - pq^o] \quad (26)$$

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:

$$s(p) = \inf_{q \geq 0} \{q | W(p, q) - pq \geq W(p, S(p)) - pS(p) - K\}. \quad (27)$$

Corollary: Under assumptions 1-5, the value function V is linear with slope p on the interval $[0, s(p)]$:

$$V(p, q) = \begin{cases} W(p, S(p)) - p[S(p) - q] - K & \text{if } q \in [0, s(p)] \\ W(p, q) & \text{if } q \in (s(p), \bar{q}]. \end{cases} \quad (28)$$

The details of the proof of Theorem 1 are in the appendix. The proof is a straightforward extension of Scarf's original proof, although the details of our proof follow the exposition in Bertsekas (1995). The corollary is easily established by substituting the generalized (S, s) rule in equation (25) into the definition of V in equation (15).

Note that equation (28) shows that the "shadow price" of an extra unit of inventory is p when $q < s(p)$. However for $q > s(p)$ the shadow price is generally not equal to p except at the target inventory level $S(p)$ as can be seen from equation (23). For $q \in [s(p), S(p))$ we have $\partial W(p, q)/\partial q > p$ and for $q \in (S(p), \bar{q}]$ we have $\partial W(p, q)/\partial q < p$. Thus, the value function is not concave in q , but as the proof of the theorem shows, it is K -concave in q . 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 $q > s(p)$ it is not optimal to order and the shadow price of an extra unit of inventory is no longer equal to p . We leave a more detailed analysis of lower level conditions sufficient to establish assumption 5 to a future analysis.

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 is strictly concave, this implies that $\partial W(p, q)/\partial q > p$ 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, $q = s(p)$. As we can see from formula (16) this kink is also present in the expected value function $W(p, q)$. However in our numerical examples below we find that there is only a small discontinuity in the partial derivative $\partial W(p, q)/\partial q$ at $q = s(p)$, so that $W(p, q)$ is approximately strictly concave in q .

We now consider how p affects the (S, s) bands. Intuitively we would expect both of the $S(p)$ and $s(p)$ functions to be downward sloping in p a result of the simple logic of commodity price speculation, namely to buy a large amount of the commodity when prices are low in order to profit from subsequent sales of its inventory holdings when prices are high. However it turns out to be rather difficult to demonstrate that this simple intuition is valid analytically. Differentiating the Euler equation for $S(p)$, (23), we obtain

$$S'(p) = \frac{1 - \partial^2 W(p, S(p))/\partial q \partial p}{\partial^2 W(p, S(p))/\partial q^2} \quad (29)$$

Since $S(p)$ maximizes $W(p, q) - pq$, the second order conditions for a maximum imply that $\partial^2 W(p, q)/\partial q^2 < 0$ at $q = S(p)$. So a sufficient condition for $S(p)$ to be downward sloping in p is that $\partial^2 W(p, S(p))/\partial q \partial p < 1$. We can derive a formula for $s'(p)$ in cases where $s(p) > 0$ since in that case we have the equality

$$W(p, s(p)) - ps(p) = W(p, S(p)) - pS(p) - K \quad (30)$$

which can be totally differentiated to obtain a formula for $s'(p)$. Since the formula for $s'(p)$ is rather complicated and unenlightening, we omit it. Suffice it to say that we have been unable to determine general sufficient conditions for either $S(p)$ or $s(p)$ to be downward sloping in p , but have verified that they are downward sloping in the numerical example in the next section.

We conclude this section by considering how our results relate to the functional equation for prices (1) assumed in the literature on the commodity storage model. The first order condition (23) determining the target inventory level $S(p)$ can be viewed as an individual level ‘‘Euler equation’’ that can be viewed as an agent-level version of (1). To see this, note that the first order condition (23) can be rewritten as

$$p = \frac{\partial ES}{\partial q}(p, S(p)) - rp - \frac{\partial c^h}{\partial q}(S(p)) + \beta \frac{\partial EV}{\partial q}(p, S(p)). \quad (31)$$

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 $\partial c^h(S(p))/\partial q$. Now consider the second term, $\beta \partial EV(p, S(p))/\partial q$, in the “Euler equation” (31). This is the expected discounted shadow price of an extra unit of inventory, and is the analog of the term $\beta E\{p_{t+1}|p_t, x_t, z_t\}$ in the commodity storage model (1). As noted above, $\partial V(p, q)/\partial q = p$ for $q < s(p)$, and at $q = S(p)$. Even though $\partial V(p, q) > p$ for $p \in [s(p), S(p))$, we also have $\partial V(p, q)/\partial q < p$ for $q \in [S(p), \bar{q}]$, so it is plausible that “on average” $\partial EV(p, S(p))/\partial q$ is close to the conditional expectation $E\{p_{t+1}|p_t\}$ that is assumed in the commodity storage model, a result that we verify in the numerical results in the next section. Thus, while individual level Euler equations do not exactly coincide with the equilibrium functional equation assumed in the commodity storage model, the two equations appear to be close.

We imagine a commodity market with a large number of intermediaries, some of whom may not be placing orders on any given day, but some may be placing orders and in doing so, choose order levels to satisfy the Euler equation (31). To the extent that different intermediaries trading in the commodity market have common convenience yields, it appears that their actions could be driving commodity prices to satisfy the functional equation (1) assumed in the commodity storage literature. We will pursue this idea in future work, closing the market by explicitly modeling the endogenous determination of commodity prices resulting from the interaction of the three types of agents in these markets: commodity producers, end consumers, and the middleman as intermediaries.

5 A Calibrated Example

To illustrate the behavior implied by our model we solved a discrete approximation of (15) numerically using parameter values selected by an informal calibration procedure. In particular we selected parameters such that specific moments constructed from simulated data “matched” the moments the actual data. We concede that it would be better to have a more formal way of selecting the parameters, such as picking those that maximize an explicit statistical criterion. In

Hall and Rust (1999b), we discuss two different approaches for doing just this, and are currently in the process of implementing these more formal estimation procedures.

We assumed that the daily interest rate is time-invariant and equal to $r = .05/261$.⁹ 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:

$$\log(p_{t+1}) = \mu_p + \lambda_p \log(p_t) + \epsilon_t \tag{32}$$

where $\mu_p = .06$, $\lambda_p = .98$, and $\{\epsilon_t\}$ is an *IID* $N(0, \sigma_p^2)$ sequence, with $\sigma_p^2 = 3.94 \times 10^{-4}$. The upper and lower truncation bounds on this process were chosen to be (13, 29) 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. These values yield a order price process with an invariant distribution with mean of 20.0 cents per pound and a standard deviation of 2.00 cents per pound. Given the markup, the mean and standard deviation of the sell price process are 22.1 and 1.11, 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 $\mu_q(p) = 5.50 - .7 \log(p_t)$ and standard deviation parameter $\sigma_q = 1.4$. These parameters yield a stationary distribution for q_t^d (conditional on $q_t^d > 0$) with conditional mean equal to 25.0 and conditional standard deviation equal to 25.0. 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 is \$100, the physical holding costs are zero, $c^h(q_t) = 0$, and that the fixed order cost is equal to \$75, i.e. $c^o(0) = 0$ and $c^o(q^o) = \$75$ 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 parameterized policy iteration (PPI). This PPI algorithm amounts to the following iterative procedure:

⁹We assumed there are $365 - (2 \times 52)$ business days in a year.

1. Approximate the value function $V(p, q)$ with a finite linear combination of basis functions.
2. Discretize the state space into a finite number of (p, q) pairs.
3. Using equation (18), compute the optimal decision rule $q_i^o(p, q)$ at each of the discretized (p, q) pairs. Note that although we discretized the state variables, we treat the control variable q_i^o as a continuous variable subject to the constraint that $0 \leq q_i^o \leq \bar{q} - q$.
4. Perform a policy iteration step. That is compute

$$V_i(p, q) = E \left\{ \sum_{j=0}^{\infty} \beta^j \pi(p_j, p_j^s, q_j^s, q_i^o(p_j, q_j) + q_j^s) \mid p, q \right\}. \quad (33)$$

5. Regress the updated value function, $V_i(p, q)$, on the discrete set of p and q 's to compute a new parameterized approximation of the value function.
6. Iterate over i on steps 3–5 until the coefficients on the parameterized approximation of the value function converge.

For the example presented in this paper we approximated the value function by a complete set of Chebychev polynomials of degree 3 in p and q . We discretized the state space with 99 grid points (9 in the p dimension and 11 in the q dimension). The grid points are fixed at the Chebychev zeros, so the grid points are more heavily weighted toward the boundaries of the state space. Policy iteration is not guaranteed to converge in continuous choice problems such as this one; but for this example, this algorithm converged in 22 iterations.

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:

$$E\hat{V}(p, q) = \frac{1}{N_p} \frac{1}{N_q} \sum_{i=1}^{N_p} \sum_{j=1}^{N_q} I\{q_j \leq q\} \hat{V}(p_i, q - q_j) h(p_j | p_i) g(p_i | p) \quad (34)$$

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 that $E\hat{V}(p, q)$ reflects that 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

$$\hat{EV}(p, q) = \frac{1}{N} \sum_{i=1}^N \hat{V}(\tilde{p}_i, q - \tilde{q}_i) \quad (35)$$

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. Niederreiter 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), whereas traditional monte carlo methods converge at rate $O_p(1/\sqrt{N})$. These favorable rates of convergence have been observed in practice, see e.g. Papageorgiu and Traub (1997).¹⁰ The density $h(q'|p, q)$ is the conditional density of q' given that $q' \leq q$,

$$h(q'|p, q) = \frac{h(q'|p)}{1 - \delta(p, q)} \quad (36)$$

where $\delta(p, q) = \Pr\{q' > q|p\} = 1 - H(q|p)$. As in the quadrature method, we adjusted the MC-PIT formula (35) to account for mass points corresponding to stockouts and zero sales. We found that the optimizing solutions for q^o were 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 the value function is nearly flat in q^o , small inaccuracies in the estimated derivatives can create 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 different approximation methods for computing \hat{ES} and \hat{EV} produced nearly identical results.

Figures 15-18 present the optimal decision rule q^o as a function of p and q and the associated expected profit function, value function and (S, s) bands. Note that our solution technique does

¹⁰We are grateful to Joseph Traub for providing the *FINDER* software co-authored with A.F. Papageorgiu that generated the low discrepancy sequences used in this study.

not exploit our prior knowledge about the form of the decision rule. The computed value function appear to be nearly linearly increasing in current inventory 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 . 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 17 shows the generalized $(S(p), s(p))$ bands implied by our model. The set of order limit points, $s(p)$, is the curve on the (q, p) plane where the $q^o(p, q)$ surface intersects the plane at $q^o = 0$. The set of target inventory points, $S(p)$, is the curve on the (q_o, p) plane where the $q^o(q, p)$ surface intersects the plane at $q = 0$. These bands are plotted in figure 18. Due to the fixed costs of ordering (\$75), 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.

Figures 19- 22 present the results from a single stochastic simulation of the DP model for 595 periods. Figure 19 shows the time series for inventory levels, and it is apparent there are multiple regimes. During the first 100 days of the simulation, as the price of steel average around 21 cents per pound, the firm maintains very low levels of inventories. Despite the stockout penalty, γ , the firm has no steel for sale on 26 of these first 100 periods. During the 74 days for which it does have steel for sale, it maintains only about 2 days-supply worth of steel.

Starting around day 100, the firm enters a “high inventory regime” with the simulated inventory levels reaching a peak over 2,000,000 pounds. This peak is consistent with observed levels of inventories for products 2 and 4. During this high inventory regime, days supply exceeds 250. 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 100 to day 300) the firm lets its inventory holdings gradually fall until inventories hit almost zero. This pattern then repeats itself. For the next 200 days then firm holds very low levels of inventories and stock-outs are frequent. But then around day 490 as the spot price falls below 18 cents per pound and the firm once again begins building up its inventories.

The high- and low-regime property of the optimal inventory holdings can be seen from the

decision rule, $q^o(q, p)$. In figure 17, $q^o(q, p)$ 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 will be able 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, quantity demanded, q^d , rises and the firm will hold more inventories to accommodate the increase in demand. Moreover, the optimal strategy implied by the model is to aggressively price speculate. When the spot price exceeds 21 cent per pound, the profit maximizing strategy is to not purchase any new steel regardless of the current level of inventories. Despite the presence of a positive markup, the likelihood of incurring a capital loss should the spot price fall dominates the profit maximization calculation.

The simulation results are consistent with this intuition. Figure 21 presents the censored and uncensored order and sales price series. In this graph, the solid line is the analogue of what we observe in our dataset, we linearly interpolate between the prices at which transactions took place; the dotted line includes the unobserved prices at which no transactions occurred. During the high inventory regimes (e.g. days 100-300, 490-550, and 570-595) the firm opportunistically bought at the troughs. This result is of course sensitive to our specification of the law of motion of the price process, equation (32).

The exogenous price and quantity demanded processes implied that the firm sold steel on 288 days at average price of 22.2 during the simulation period. The decision rule dictated that the firm purchased steel on 83 days at an average price of 18.2. The average order size was 82,150 pounds. And the conditional standard deviation of the order size was 157,000. 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 7.9. So the model does imply that orders are more volatile than sales. Longer simulations generate similar results.

These results are also qualitatively similar to the actual inventory time series for our firm in figures 3-14. 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 orders occurring when prices hit what appear to be record lows. However comparing

figures 4, 8, and 12 with figure 20, we see that the DP model generates fewer small size orders at low prices than we observe in the data. This suggests that perhaps the fixed order cost is too large; however setting the fixed cost to zero, exacerbates the counterfactual result that when prices are high, the firm holds very low levels of inventories and tightly matches orders to sales. The model implies frequent stockouts; in the simulation presented, the firm stocks out (holds zero inventories) on 55 days.

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.

Finally, from the simulations we can deduce the importance to the firm's profits of the markup relative to the capital gains and losses from the firm's optimal speculation strategy. Note that by substituting the law of motion for inventories, equation (5), into the firm's one-period profit function, equation (12), the profit function can be rewritten as

$$\pi_t = (p_t^s - p_t)q_t^s + (p_t - (1+r)p_{t-1})q_t - rp_t(q_t + q_t^o) - c^o(q_t^o) - c^h(q_t + q_t^o) \quad (37)$$

The first term of (37) can be interpreted as the markup over the current spot price while the second term can be interpreted as a capital gains or loss from holding the steel from period $t - 1$ into period t . Using this decomposition on simulations of 10,000 periods, we find that according to our model, about 44 percent of the firm's revenue are due the capital gains component while the remaining 56 percent are attributed to the markup. We also resolved the model setting the markups to zero (so $p_t^s = p_t \forall t$). We then simulated the model with no markups for 10,000 periods. For the model with positive markups, the present value of the firm's profits is \$1,430,700; in the case with zero markup, the present value of the firm's profits is just \$492,500. The value of the firm falls by almost two-thirds if we eliminate the markup. While the model predicts that the firm should aggressively speculate on the price of steel, the returns to price speculation account for less than half of the firm's revenue.

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 23 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.¹¹ Note that at the end of the sample the real price of carbon plate is at (at least) a twelve-year low.

Figure 24 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).¹² 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.

For the period from July, 1997 to May, 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). As discussed above, our model implies a similar condition,

¹¹Deflating this price index by the CPI would not change any of analysis presented below.

¹²The SSCI is an industry organization which among other things collects data on member firms’ shipments and inventory holdings.

equation (31). Large swings in prices in and of themselves do not contradict either (1) or (31), 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 developed a model price speculation and inventory investment by durable commodity intermediaries. The model was motivated by 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 day-to-day and within-day 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 showed 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 formulated and solved the problem of optimal inventory speculation by durable commodity intermediaries, providing sufficient conditions under which the optimal inventory investment strategy takes the form of a generalized (S, s) rule where S and s are declining functions of the

spot price of the commodity. Via simulations of a calibrated version of our DP model, we demonstrated that the firm's behavior at the product level can be well approximated by an optimal trading strategy. We employed a novel continuous 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. In Hall and Rust (1999b) we present two methods for addressing the difficult problem 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 interest rates on inventory investment. 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. Finally, we plan to study aggregate our micro model to study endogenous price determination for the commodity market as a whole, in order to determine whether the no arbitrage conditions assumed in the rational expectations commodity price model of Williams and Wright can be derived from microfoundations in a market where there is considerable price dispersion and frictions despite the apparent homogeneity of the product.

References

- [1] Abel, Andrew (1983) “Inventories, Stock-Outs and Production Smoothing” *Review of Economic Studies* 52, 283-293.
- [2] Arrow, K.J., Harris, T. and J. Marschak (1951) “Optimal Inventory Policy” *Econometrica* 19-3, 250–272.
- [3] Bertsekas, D.P. (1995) *Dynamic Programming and Optimal Control* Athena Scientific, Belmont, MA.
- [4] Bertsekas, D.P. and J.N. Tsitsiklis (1996) *Neuro-Dynamic Programming* Athena Scientific, Belmont, MA.
- [5] Bils, M., and J. Kahn (1996) “What Inventory Behavior Tells Us about Business Cycles” manuscript, University of Rochester.
- [6] Blanchard, O. (1983) “The Production and Inventory Behavior of the American Automobile Industry” *Journal of Political Economy*, 91, 365-400.
- [7] Blinder, A. (1981) “Retail Inventory Investment and Business Fluctuations” *Brookings Papers on Economic Activity*, 2, 443-505.
- [8] Blinder, A. (1986) “Can the Production Smoothing Model of Inventory Behavior be Saved?” *Quarterly Journal of Economics*, 101, 431-453.
- [9] Bresnahan, T., and V. Ramey (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] Durlauf, S.N. and Maccini, L.J. (1995) “Measuring Noise in Inventory Models” *Journal of Monetary Economics* **36** 65–89.
- [12] Deaton, A. and G. Laroque (1992) “On the Behavior of Commodity Prices” *Review of Economic Studies* 59, 1–23.

- [13] Eichenbaum, M. (1984) “Rational Expectations and the Smoothing Properties of Inventories of Finished Goods” *Journal of Monetary Economics*, 14, 71-96.
- [14] Eichenbaum, M. (1989) “Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment” *American Economic Review*, 79, 853-64.
- [15] Fair, R. (1989) “The Production-Smoothing Model is Alive and Well” *Journal of Monetary Economics*, 24, 353-370.
- [16] Fisher, J. and A. Hornstein (1998) “ (S, s) Inventory Policies in General Equilibrium” manuscript, Federal Reserve Bank of Chicago.
- [17] Hall, G. (1999) “Nonconvex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants” manuscript, Yale University, forthcoming in *Journal of Monetary Economics*.
- [18] Hall, G. and J. Rust (1999a) “An Empirical Model of Inventory Investment by Durable Commodity Intermediaries” manuscript, Yale University, forthcoming in *Carnegie-Rochester Conference Series on Public Policy*, 52.
- [19] Hall, G. and J. Rust (1999b) “Econometric Methods for Endogenously Sampled Time Series: The Case of Commodity Price Speculation in the Steel Market” manuscript, Yale University.
- [20] Holt, C. Modigliani, F. Muth, J. and H. Simon (1960) *Planning Production, Inventories and Work Force* Prentice-Hall, Englewood Cliffs, N.J.
- [21] Judd, K. (1998) *Numerical Methods in Economics* MIT Press, Cambridge, MA.
- [22] Kahn, J. (1987) “Inventories and the Volatility of Production” *American Economic Review*, 77, 667-679.
- [23] 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.
- [24] Kashyap, A., and D. Wilcox (1993) “Production and Inventory Control at the General Motors Corporation During the 1920’s and 1930’s” *American Economic Review*, 83, 383-401.

- [25] 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*
- [26] Miron, J. and S. Zeldes (1988) “Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories” *Econometrica*, 56, 877-908.
- [27] Miron, J. and S. Zeldes (1989) “Production, Sales, and the Change in Inventories: An Identity That Doesn’t Add Up” *Journal of Monetary Economics*, 24, 31-51.
- [28] Neiderreiter, H. (1992) *Random Number Generation and Quasi-Monte Carlo Methods*. CBMS-NSF Regional Conference Series in Applied Mathematics, Philadelphia, SIAM.
- [29] Papageorgiu, A.F. and J.F. Traub (1997) “Faster Evaluation of Multidimensional Integrals” *Computational Physics* **11-6** 574-578.
- [30] Ramey, V. (1991) “Nonconvex Costs and the Behavior of Inventories” *Journal of Political Economy*, 99, 306-334.
- [31] Ramey, V. and K. West (1997) “Inventories” NBER working paper 6315, December, forthcoming in the *Handbook of Macroeconomics*.
- [32] Rust, J. (1997) “Using Randomization to Break the Curse of Dimensionality” *Econometrica* **65-3** 487-516.
- [33] Rust, J. (1998) “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>
- [34] Rust, J. Traub, J. and H. Woźniakowski (1998) “No Curse of Dimensionality for Contraction Fixed Points Even in the Worst Case” manuscript.
- [35] Scarf, H. (1960) “The Optimality of (S, s) Policies in the Dynamic Inventory Problem” In *Mathematical Methods in the Social Sciences* K. Arrow, S. Karlin and P. Suppes (ed.), Stanford, CA: Stanford University Press, 196-202.

- [36] Tezuka, S. (1995) *Uniform Random Numbers: Theory and Practice* Dordrecht, Netherlands, Kluwer.
- [37] Van Roy, B. Bertsekas, D.P., Lee, Y. and J.N. Tsitsiklis (1997) “A Neuro-Dynamic Programming Approach to Retailer Inventory Management” manuscript, MIT Laboratory for Information and Decision Systems, forthcoming,
- [38] West, K. (1986) “A Variance Bounds Test of the Linear Quadratic Inventory Model” *Journal of Political Economy* 94, 374-401.
- [39] Williams, J.C. and B. Wright (1991) *Storage and Commodity Markets* Cambridge University Press, New York.
- [40] Working, H. (1949) “Theory of Price and Storage” *American Economic Review* 39, 1254–62.

Appendix Proof of Optimality of Generalized (S, s) rule.

Our proof is a straightforward extension of Scarf's original proof (see Scarf, 1960) as expositied in Bertsekas 1995. It turns out that Scarf's original proof extends to inventory problems with multiple state variables provided these additional variables enter the formulation of the problem in a way that maintains the symmetry and K -conavity properties discussed in section 4. For simplicity we extend Scarf's proof to the case where there is only one extra state variable besides the current quantity on hand, q_t . In the proof below the additional state variable is interpreted as the spot price p_t at which the intermediary orders new additions to inventory. However it is not difficult to show that an (S, s) rule is optimal under more general formulations where in addition to p_t there is a vector of state variables x_t that affect the firm's current period profit function or which affects the firm's beliefs about future spot prices, sales prices, or demand. Examples of these extra x_t variables include interest rates or other macro variables that affect spot prices or the firm's future demand for its product. In this case the (S, s) bands will be functions of the additional state variables x in addition to p . This would require carrying extra notation (e.g. $S(p, x)$ and $s(p, x)$) but would not affect the basic structure of the proof. We begin by stating four preliminary lemmas that will be required in our inductive proof of the optimality of the generalized (S, s) rule. Proofs of Lemmas 1 and 2 appear in Bertsekas (1995): the proofs of Lemmas 3 and 4 are straightforward and are omitted for brevity.

Lemma 1: *If $g(p, q)$ is K -concave in q , then it is also L -concave in q for all $L \geq K$.*

Lemma 2: *If $f(p, q)$ is K -concave in q and $g(p, q)$ is L -concave in q , then $\alpha f + \beta g$ is $\alpha K + \beta L$ -concave in q for any non-negative constants α and β .*

Lemma 3: *Pointwise limits of K -concave functions are K -concave.*

For the next lemma, we need to introduce the *Bellman operator* for the inventory problem:

Definition: *The Bellman operator Γ is a mapping from the space $C([0, \bar{p}] \times [0, \bar{q}])$ of bounded, continuous functions from the set $[0, \bar{p}] \times [0, \bar{q}]$ to R given by:*

$$\Gamma(V)(p, q) = \max_{0 \leq q^o \leq \bar{q} - q} \left[ES(p, q + q^o) - rp(q + q^o) - c^h(q + q^o) - pq^o - c^o(q^o) + EV(p, q + q^o) \right], \quad (38)$$

where ES and EV are defined in equations (17) and (7).

The next lemma is a key result, showing that symmetry and K -concavity of the function $W(p, q)$ implies the optimality of (S, s) for the corresponding decision rule $q^o(p, q)$.

Lemma 4: If $W(p, q)$ is K -concave in q , then the function $q^o(p, q)$ given by:

$$q^o(p, q) = \underset{q^o \in [0, \bar{q} - q]}{\operatorname{argmax}} [W(p, q + q^o) - pq^o] \quad (39)$$

is of the (S, s) form, i.e.

$$q^o(p, q) = \begin{cases} 0 & \text{if } q \geq s(p) \\ S(p) - q & \text{otherwise} \end{cases} \quad (40)$$

where $S(p)$ is given by:

$$S(p) = \underset{0 \leq q^o \leq \bar{q} - q}{\operatorname{argmax}} [W(p, q^o) - pq^o] \quad (41)$$

and $s(p)$ is defined by:

$$s(p) = \inf_{q \geq 0} \{q | W(p, q) - pq \geq W(p, S(p)) - pS(p) - K\}. \quad (42)$$

Proof: First note that K -concavity does not necessarily imply the uniqueness of $S(p)$ in equation (41). If there are multiple maximizers, we define $S(p)$ as the *smallest* solution to (41). Now if $q < s(p)$ we need to show that $q^o(p, q) = S(p) - q$. First we show that

$$W(p, S(p)) - p(S(p) - q) - K > W(p, q), \quad (43)$$

i.e. it is better to order $q^o = S(p) - q$ than $q^o = 0$. This result follows from a rearrangement of the definition of $s(p)$ in equation (42). Specifically, if $q < s(p)$ then we have

$$W(p, q) - pq < W(p, S(p)) - pS(p) - K, \quad (44)$$

which is equivalent to inequality (43). Now we show that there is no alternative positive investment level that is superior to $q^o(p, q) = S(p) - q$. That is, for any $z \geq 0$ we have

$$W(p, q + z) - pz - K \leq W(p, S(p)) - p(S(p) - q) - K. \quad (45)$$

This inequality is equivalent to the inequality

$$W(p, q + z) - p(q + z) \leq W(p, S(p)) - pS(p) \quad (46)$$

which holds due to the definition of $S(p)$ in equation (41). We complete the proof by showing that when $q \leq s(p)$ the optimal order is zero, i.e. $q^o(p, q) = 0$. First note that the linear function $-pq$ is 0-concave in q , so Lemmas 1 and 2 imply that $W(p, q) - pq$ is K -concave in q . If $q > S(p)$, we have by the definition of K -concavity

$$W(p, q + z) - p(q + z) - K \leq W(p, q) - pq + \frac{z}{S(p) - q} [W(p, q) - pq - W(p, S(p)) + pS(p)]. \quad (47)$$

The definition of $S(p)$ implies that the second term on the right hand side of the above inequality is non-positive, so we have

$$W(p, q + z) - p(q + z) - K \leq W(p, q) - pq, \quad (48)$$

which is equivalent to

$$W(p, q + z) - pz - K \leq W(p, q), \quad (49)$$

which implies that $q^o(p, q) = 0$. Next consider the case where $q \in (s(p), S(p)]$. By K -concavity of $W(p, q) - pq$ we have

$$W(p, S(p)) - pS(p) - K \leq W(p, q) - pq + \frac{S(p) - q}{q - s(p)} [W(p, q) - pq - W(p, s(p)) + ps(p)]. \quad (50)$$

If $s(p) > 0$ we have $W(p, S(p)) - pS(p) - K \leq W(p, s(p)) - ps(p)$ from the definition of $s(p)$ in equation (42). Substituting this inequality into the right hand side of inequality (50) and rearranging we have

$$\frac{S(p) - s(p)}{q - s(p)} [W(p, S(p)) - pS(p) - K] \leq \frac{S(p) - s(p)}{q - s(p)} [W(p, q) - pq]. \quad (51)$$

This is equivalent to

$$W(p, S(p)) - p(S(p) - q) - K \leq W(p, q). \quad (52)$$

Further, by the definition of $S(p)$ we have

$$W(p, q + z) - pz - K \leq W(p, S(p)) - p(S(p) - q) - K \leq W(p, q). \quad (53)$$

This implies that $q^o(p, q) = 0$. The final case to consider is $q = s(p)$. Using the definitions of $S(p)$ and $s(p)$ we have

$$W(p, s(p) + z) - pz - K \leq W(p, S(p)) - p(S(p) - s(p)) - K \leq W(p, s(p)), \quad (54)$$

i.e. $q^o(p, s(p)) = 0$.

Lemma 5 *If $W(p, q)$ is non-decreasing and K -concave, then the optimal value function $V(p, q)$ implied by the optimal (S, s) decision rule in (40) is K -concave. That is, the function $V(p, q)$ given by*

$$V(p, q) = \begin{cases} W(p, S(p)) - p[S(p) - q] - K & \text{if } q \in [0, s(p)] \\ W(p, q) & \text{if } q \in (s(p), \bar{q}]. \end{cases} \quad (55)$$

is K -concave.

Proof: We need to show that for all nonnegative q , z and b satisfying $q + z \leq \bar{q}$ and $q - b \geq 0$ we have:

$$V(p, q + z) - K \leq V(p, q) + \frac{z}{b}[V(p, q) - V(p, q - b)]. \quad (56)$$

Since $V(p, q)$ is linear in the interval $[0, s(p)]$ it is 0-concave in this region, and thus K -concave in this region by Lemma 1. Also, $V(p, q) = W(p, q)$ for $q \geq s(p)$. It follows immediately that V is K -concave in these regions, so inequality (56) holds when $q + z < s(p)$, or when $q - b > s(p)$. The only cases that need more detailed examination are when $q + z > s(p)$ and $q < s(p)$ or $q - b < s(p)$ since these involve points on either side of the convex kink in $V(p, q)$ at the point $q = s(p)$. When $q + z > s(p)$ and $q < s(p)$ we substitute the definition of V in equation (55) into the condition for K -concavity of V in inequality (56) to obtain:

$$W(p, q + z) - K \leq W(p, S(p)) - p(S(p) - q) - K + zp, \quad (57)$$

which is equivalent to the inequality

$$W(p, q + z) - p(q + z) \leq W(p, S(p)) - pS(p), \quad (58)$$

which holds due to the definition of $S(p)$ in equation (41). The final case to consider is when $q + z > s(p)$, $q > s(p)$ and $q - b < s(p)$. In this case, making use of the definition of $s(p)$ as the smallest value of q for which $W(p, q) = W(p, S(p)) - p(S(p) - q) - K$, and substituting this into the condition for the K -concavity of V yields:

$$W(p, q + z) - K \leq W(p, q) + \frac{z}{b}[W(p, q) - W(p, s(p) + p(s(p) - (q - b))]. \quad (59)$$

However the K -concavity of $W(p, q)$ implies that

$$W(p, q + z) - K \leq W(p, q) + \frac{z}{q - s(p)}[W(p, q) - W(p, s(p))]. \quad (60)$$

However by assumption, $b > q - s(p)$ and $q > s(p)$ so by the assumption that $W(p, q)$ is non-decreasing in q we have $W(p, q) \leq W(p, s(p))$. Thus we have that $p(s(p) - (q - b)) > 0$ and

$$\frac{z}{b}[W(p, q) - W(p, s(p))] \geq \frac{z}{q - s(p)}[W(p, q) - W(p, s(p))]. \quad (61)$$

Thus, inequalities (60) and (61) imply that the K -concavity condition, inequality (59) holds, so the K -concavity of V is verified.

Using arguments similar to those used in the proofs of Lemmas 1 to 4, it is not difficult to establish the final two results.

Lemma 6: *If $V(p, q)$ is non-decreasing in its second argument, then $\Gamma(V)(p, q)$ is also non-decreasing in its second argument.*

Lemma 7: *If $V(p, q)$ is K -concave in its second argument, then $\Gamma(V)(p, q)$ is also K -concave in its second argument.*

We now have all the results to prove the key result, Theorem 1, which established the optimality of the generalized (S, s) rule.

Proof of Theorem 1: The proof is by induction using Lemmas 1 to 7. Since the Bellman operator Γ is a contraction mapping, the value function $V = \Gamma(V)$ equals the limit of successive approximations of the Γ starting from an initial guess of 0. That is

$$V = \lim_{n \rightarrow \infty} V_n \equiv \Gamma^n(0). \quad (62)$$

where 0 denotes the zero function in $C([0, \bar{p}] \times [0, \bar{q}])$ and Γ^n denotes n -fold composition of the Bellman operator, i.e., $\Gamma^2(0) = \Gamma(\Gamma(0))$, $\Gamma^3(0) = \Gamma(\Gamma(\Gamma(0)))$ and so forth. We show below that if V_{n-1} is non-decreasing and K -concave, then $V_n = \Gamma(V_{n-1})$ is non-decreasing and K -concave. Lemma 4 implies that the fixed point $V = \Gamma(V)$ also has these properties since it is a uniform limit of non-decreasing and K -concave functions. Lemmas 1 and 2 imply that the function $W(p, q)$ defined in equation (16) is K -concave since W is the sum of a concave function ES , a linear function $-pq$, and the function EV which is K -concave by virtue of the K -concavity of V and Assumption 5. The K -concavity of W and Lemma 4 imply that the optimal policy must be of the (S, s) form.

The proof is completed by noting that $V_0 = 0$ is a non-decreasing, K -concave function so it is only necessary to verify the inductive hypothesis, i.e., if V_{n-1} is non-decreasing and K -concave, then V_n has these same properties. That is, we need to show that the Bellman operator Γ preserves the properties of monotonicity and K -concavity. Γ preserves monotonicity by Lemma 6, and Γ preserves K -concavity by Lemma 7.

product (1)	# order		mean		std		# sell		mean		std		std(order price)/ std(sell price)	
	days (2)	order price (3)	order price (4)	std (4)	days (5)	sell price (6)	sell price (7)	std (7)	std (8)					
1	55	19.62	4.44	281	21.42	2.81	1.40							
2	81	18.58	3.93	396	21.15	3.35	1.87							
3	18	16.75	6.24	146	21.35	3.02	1.56							
4	88	18.62	4.72	374	21.58	2.91	1.99							
5	34	19.05	5.79	110	21.62	4.14	1.41							
6	62	18.40	5.86	253	21.35	3.86	1.64							
7	23	18.78	6.31	65	22.50	4.44	1.44							
8	16	19.20	6.42	50	22.85	3.24	1.02							
9	24	21.06	3.31	120	22.88	4.38	1.23							
10	55	21.29	5.38	229	23.29	1.54	1.85							
11	8	21.98	2.84	14	23.13	4.03	1.02							
12	23	21.41	4.09	91	23.49	3.62	1.65							
13	37	20.79	5.97	117	23.60	4.37	0.50							
14	21	21.44	2.19	48	23.35	6.77	0.36							
15	24	21.65	2.46	58	24.41	2.40	1.07							
16	11	20.90	2.56	22	24.11	6.68	0.44							
17	4	24.77	2.95	9	26.71	1.10	6.83							
18	6	22.88	7.54	9										

Table 1: First and Second Moments of Prices

There are 595 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	6.61	71.67	39.47	111.42	5.36	11.41	9.28	10.69	10.42
2	18.79	138.26	105.52	257.08	16.24	24.50	21.80	22.69	11.33
3	2.04	67.66	17.19	75.10	2.77	11.32	8.90	15.10	4.97
4	20.44	138.40	128.37	310.15	18.92	30.40	35.28	40.65	7.63
5	2.58	45.19	18.43	64.35	2.11	11.43	6.52	11.16	5.77
6	8.80	84.58	43.43	109.00	7.80	18.52	15.37	19.04	5.72
7	1.39	35.94	11.10	45.09	1.46	13.41	5.12	9.00	5.01
8	1.12	41.81	9.28	40.01	1.02	12.61	3.88	6.44	6.21
9	3.70	91.96	26.35	97.41	3.38	16.77	15.09	30.20	3.23
10	10.79	116.91	56.72	151.02	9.76	25.40	18.08	21.31	7.09
11	0.37	27.77	3.49	12.86	0.33	13.89	2.52	9.43	1.36
12	3.26	84.38	24.52	95.41	3.08	20.42	8.83	12.79	7.46
13	5.13	82.56	31.31	98.16	4.80	24.64	15.67	27.88	3.52
14	2.65	75.30	20.56	82.67	1.87	23.21	8.12	18.14	4.56
15	4.01	99.65	30.02	115.60	3.09	31.77	14.45	35.38	3.27
16	2.07	111.98	19.28	92.61	0.93	25.10	5.83	18.11	5.11
17	0.69	102.92	11.91	118.84	0.35	22.87	3.61	19.81	6.00
18	1.18	117.62	18.92	161.75	0.39	26.14	3.19	0.00	∞
aggregate	95.62	211.86	436.61	631.27	83.65	87.93	75.60	75.04	8.41

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: Times series plot of the inventory for product 2.




Figure 4: Size of purchases for product 2 as a function current inventory holdings and the buy price.

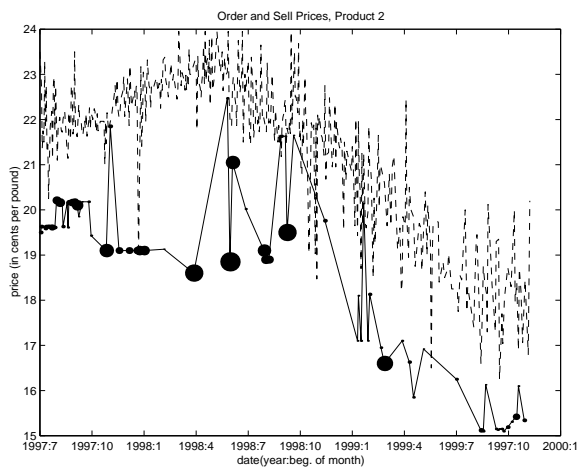


Figure 5: Order prices (solid line) and sell prices (dashed line) for product 2. For the order price series, the size of the marker is proportional to the size of the purchase.




Figure 6: Days-supply of inventory for product 2 (in business days).

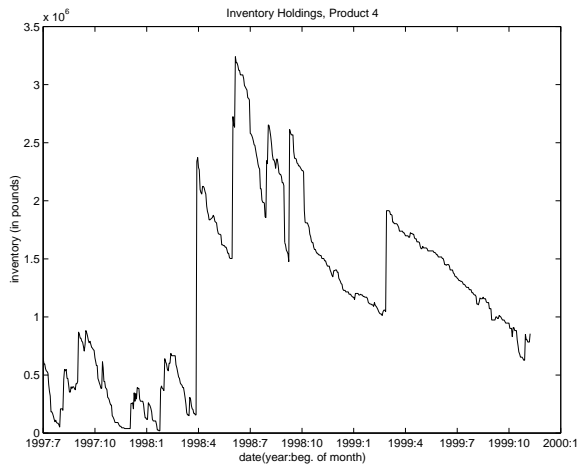


Figure 7: Times series plot of the inventory for product 4.

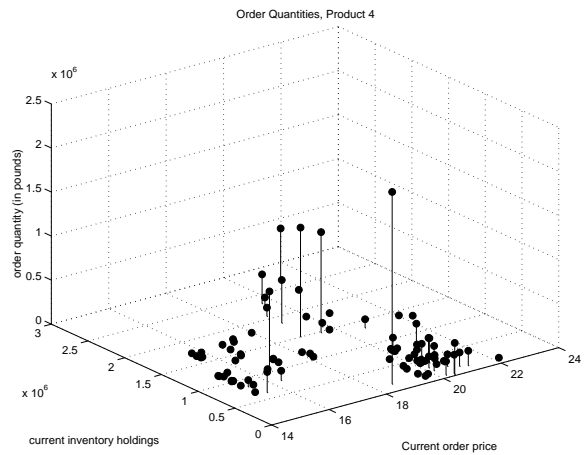


Figure 8: Size of purchases for product 4 as a function current inventory holdings and the buy price.




Figure 9: 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.

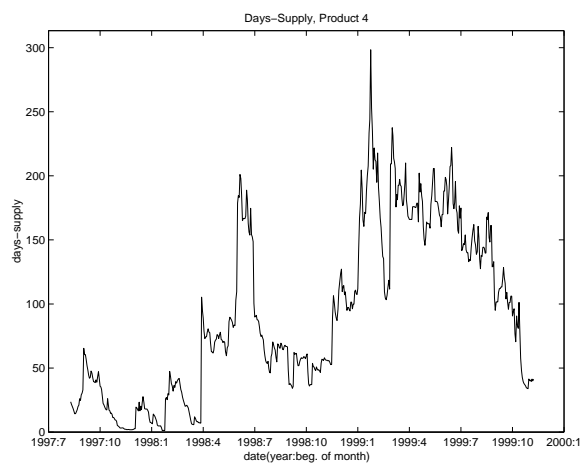


Figure 10: Days-supply of inventory for product 4 (in business days).




Figure 11: Times series plot of the inventory for product 13.

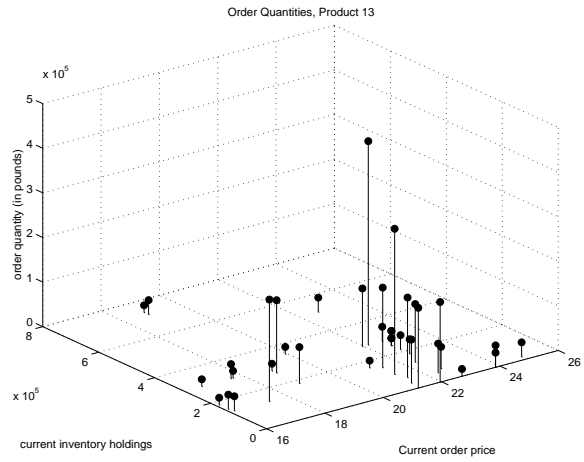


Figure 12: Size of purchases for product 13 as a function current inventory holdings and the buy price.




Figure 13: 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.

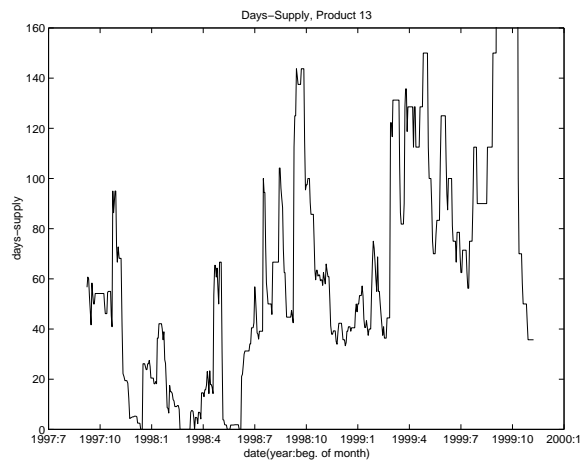


Figure 14: Days-supply of inventory for product 13 (in business days).

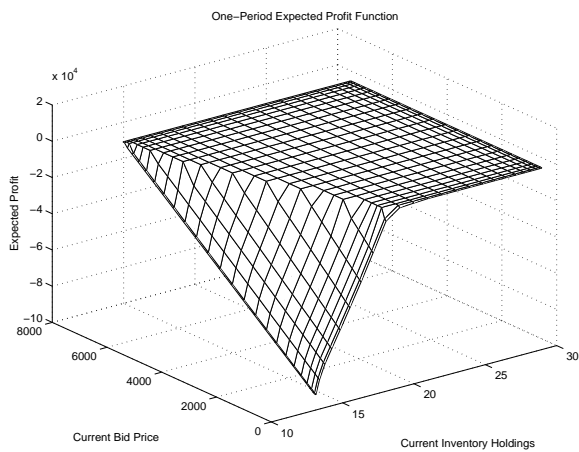


Figure 15: Expected one-period profit function for the calibrated example.

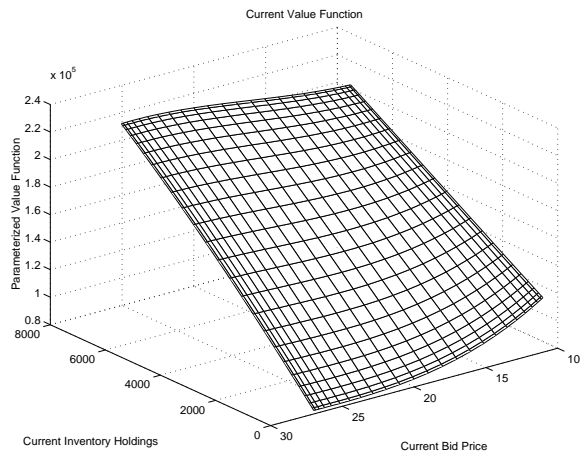


Figure 16: The value function, $V(q, p)$ for the calibrated example.

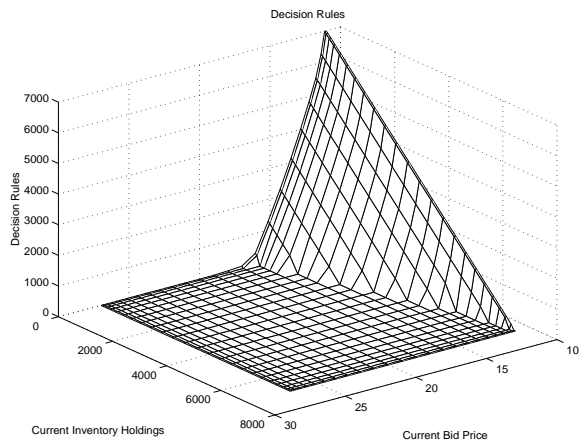


Figure 17: Decision rule, $q^o(q, p)$, for the calibrated example.

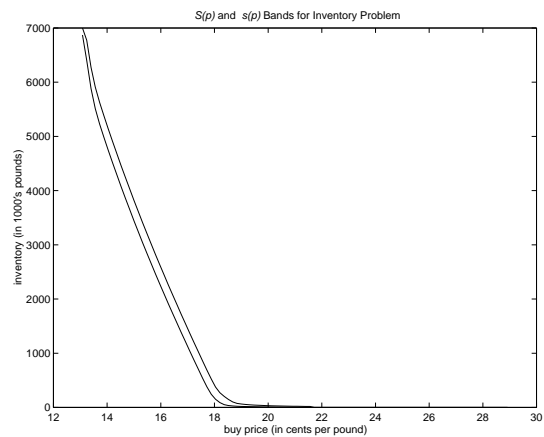


Figure 18: $S(p)$ and $s(p)$ for the calibrated example.

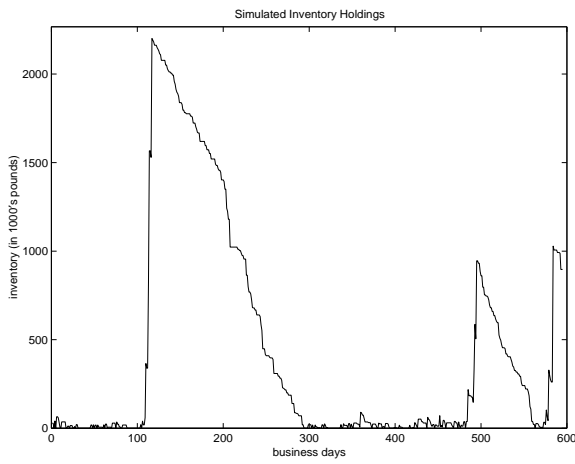


Figure 19: Simulated inventory holdings

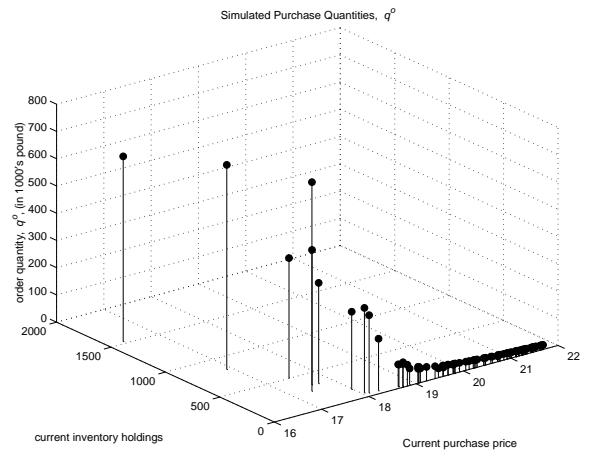


Figure 20: Simulated orders as a function current inventory holdings and buy price.

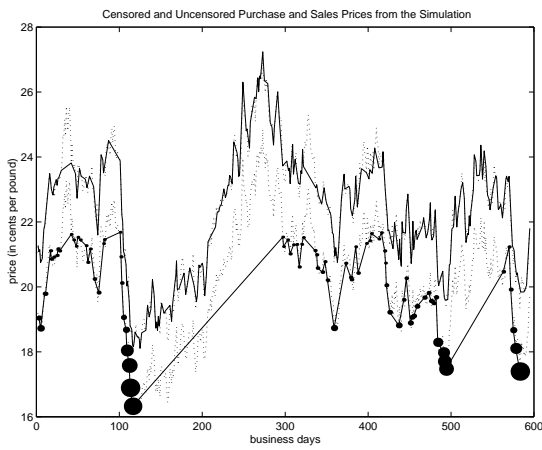


Figure 21: Censored (solid line) and Uncensored (dotted line) order and sales prices from the simulation.

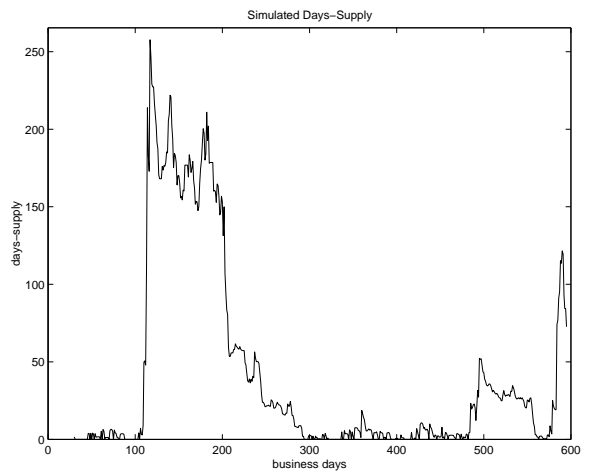


Figure 22: Simulated days-supply of inventory (in business days).

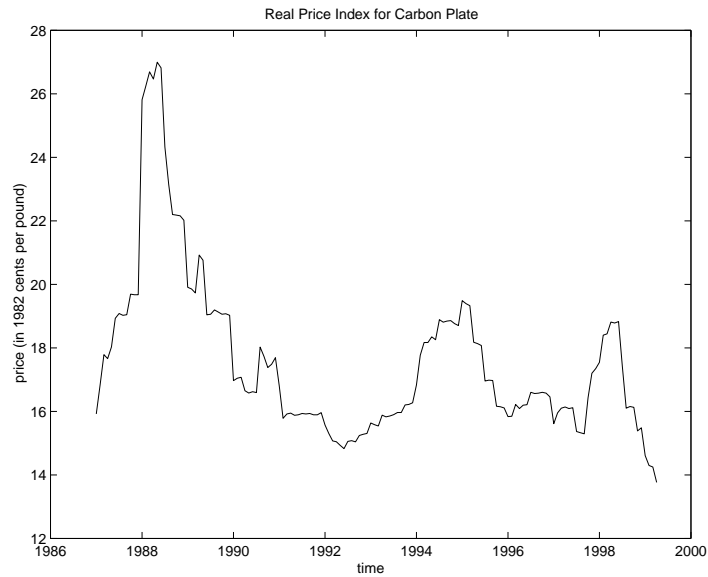


Figure 23: Price index of carbon plate steel from *Purchasing Magazine* deflated by the PPI.




Figure 24: 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.