Simulated Minimum Distance Estimation of a Model of Optimal Commodity Price Speculation with Endogenously Sampled Prices

George Hall, *Yale University*[*]
John Rust, *University of Maryland*

July 2003

Abstract: This paper provides structural parameter estimates of a model of optimal price speculation by a middleman in the steel market. The middleman profits by purchasing large quantities of steel from producers and other middlemen in the wholesale market and subsequently reselling smaller quantities to individual customers in the retail market. The econometric challenge is to make inferences about the law of motion governing wholesale prices given that these prices are *endogenously sampled,* i.e. we only observe market prices on the days when the middleman purchases steel. Since the essence of profitable speculation is to "buy low and sell high", the middleman purchases steel relatively infrequently at prices that are presumably lower than those on other days. To avoid misleading inferences, a standard approach to this "selectivity bias problem" is to use a conditional likelihood function that "integrates out" prices on the days when no purchases are made. However since the middleman purchases steel on only 200 of the 1500 business days in our data set, the conditional likelihood approach is cumbersome since it requires multivariate integration over 1300 dimensions. We propose a less efficient but simpler and computationally tractable *simulated minimum distance* (SMD) estimator to estimate the structural parameters in the presence of endogenous sampling. The SMD estimator minimizes a quadratic form specifying the distance between sample moments and predicted moments from simulations of our model. We show, via a limited Monte Carlo study, that the SMD estimator yields accurate and unbiased estimates of the structural parameters (including the law of motion of the wholesale price process). We use the SMD estimator to estimate the unknown parameters of the middleman's profit function and the wholesale price processes governing two specific types of steel plate. While our commodity price speculation model is formally rejected by goodness of fit tests, the model provides a good approximation to the middleman's trading behavior. Via in-sample comparisons, we find that the optimal trading rule results in only 10 to 20 percent higher profits than those actually earned by the middleman.

Keywords: endogenous sampling, Markov processes, simulation estimation, (S,s) inventory theory, steel, middlemen
JEL classification: C1, C6, L2

[*]Corresponding author. Department of Economics, Yale University, New Haven, CT 06520-8268, phone: (203) 432-3566, fax: (203) 432-5779 e-mail: `george.hall@yale.edu`. An earlier version of this paper, "Econometric Methods for Endogenously Sampled Time Series: The Case of Commodity Price Speculation" was presented at the *JAE* Conference on Inference and Decision Making in Rotterdam in 2000 and at the Cowles Foundation Econometrics Conference at Yale University in October 1999. We thank our discussant, Halbert White, and the conference participants including Herman van Dijk, Robert Engle and George Tauchen for helpful comments. We also received helpful feedback from Michael Keane, Jean-Marc Robin, Kenneth Wolpin and other seminar attendees at subsequent presentations of this paper at various universities and conferences. We are grateful for financial support from National Science Foundation grant SES-9905145.

1 Introduction

This paper estimates a structural model of optimal commodity price speculation using six years of daily microdata from a particular trader in the steel market. The trader who provided us with these data, John,[1] can be described as a middleman since his firm does only minimal production processing. Most of the firm's profits come from frequent sales of relatively small quantities of steel to retail customers at a markup over the wholesale prices it pays to acquire steel inventory in less frequent but larger batch purchases from producers and other middlemen. In previous work (Hall and Rust, 2001) we showed that under fairly general conditions the optimal speculative trading strategy for such a firm takes the form of a *generalized* (S,s) *rule* where the (S,s) bands are functions of state variables, particularly the wholesale price of steel, p. When quantity on hand q falls below the lower threshold $s(p)$ the optimal strategy is to order enough steel to restore inventories to the target threshold $S(p)$. The (S,s) bands are decreasing functions of the wholesale price, capturing the intuitive notion of speculation as "buying low and selling high." Hall and Rust (1999) demonstrated that an informally calibrated version of this model accounts for the stylized facts of the trading behavior that we observe in our dataset. The purpose of this paper is to introduce an econometric estimator that allows us to formally test whether John's trading behavior is governed by a generalized (S,s) rule.

However we face a significant econometric complication: the wholesale price data we observe are *endogenously sampled.* That is, we only observe wholesale prices on days that John's firm purchased steel. The endogenous sampling problem arises from the simple fact that John's firm records prices when it makes a transaction, but it does not record potential purchase or sales prices on days that it did not buy or sell. A potential solution to this problem would be to augment our daily microdata from John with daily market level transaction price data, enabling us to "fill in" the unobserved wholesale prices. Unfortunately, no such data exist since there is no central exchange or marketplace where steel transaction prices are recorded. Instead, the steel market operates as a decentralized "search and bargaining" market where the typical individual transaction is completed privately over the telephone.[2]

[1]Unfortunately we cannot reveal his last name or the name of his company due to a confidentiality agreement we signed in order to gain access to the data.

[2]Surprisingly, while centralized exchanges such as the Chicago Board of Trade exist for commodities such as wheat and pork bellies, there is no central exchange where steel products are traded. Rust and Hall (2003) developed a theory of intermediation to explain why some commodities are traded on exchanges and others are traded via decentralized search and bargaining as in the steel market. In their theory, the microstructure of trade in a commodity or asset is endogenously determined. The relative share of trade handled by "middlemen" and the exchanges ("market makers") depends on a number of parameters including the transactions costs of search and the relative costs of intermediation by middlemen and market makers. There are a range of

John's trading behavior appears to be consistent with the essential feature of profitable speculation, namely the attempt to buy low and sell high. In particular, John tends to purchase large quantities of steel relatively infrequently at prices that are presumably lower than the prices available on days when no purchases are made. The better John is at speculation (i.e. strategically timing purchases at the "troughs" of the wholesale price series), the more serious are the potential biases resulting from the endogenous sampling of wholesale prices. In particular, special econometric corrections are required in order to avoid misleading inferences about the mean, variance and serial correlation properties of the wholesale price series. A standard approach to this "selectivity bias problem" is to use a conditional likelihood function that "integrates out" purchase prices on the days when no purchases are made. However since John purchases steel on roughly 200 of the 1500 business days in our data set, the conditional likelihood approach is intractable since it requires computation integrals over 1300 dimensions. Further the region over which these prices are integrated depends on knowledge of John's trading strategy, i.e. a specification of a time- and state-dependent threshold above which John is unwilling to make new purchases.

We propose a less efficient but simpler and computationally tractable *simulated minimum distance* (SMD) estimator to estimate the structural parameters. The SMD estimator can be viewed as a simulated moments estimator (SME) (Lee and Ingram, 1991 and Duffie and Singleton, 1993), applied to a situation where the data are endogenously sampled. It minimizes a quadratic form specifying the distance between sample moments and predicted moments from simulations of our model of optimal commodity price speculation. We show, via a limited Monte Carlo study, that the SMD estimator yields relatively accurate and unbiased estimates of the structural parameters (including the law of motion of the wholesale price process). We use SMD to estimate unknown parameters in John's profit function and the wholesale price processes governing two specific types of steel plate. While our commodity price speculation model is formally rejected by goodness of fit tests, the generalized (S,s) rule implied by our model does capture several basic features of John's trading behavior.

We conduct several detailed comparisons between the optimal trading rule implied by our estimated model and John's actual trading behavior. We develop a technique for predicting the expected profits that the optimal trading rule would have earned if it had been confronted with the same *ex post* sequence of retail sales and purchase prices that John's firm encountered over the six year sample period. We find

parameters for which all trade occurs via the market maker on a centralized exchange, and other ranges where all transactions transactions occur in a decentralized search and bargaining process via middlemen, much as we observe in the steel market. This theory could explain the variety of different trading institutions that we see in different markets, including the nonexistence of centralized exchanges for steel.

that the conditional expectation of profits from optimal trading rule are 10 to 20 percent higher than the profits John actually earned. Further, under the optimal trading rule a significantly higher share of profits are due to speculation (i.e. the firm's success in buying low and selling high) than from markups. While these results might suggest that the reason for the rejection of our model is that John is not following a profit maximizing strategy, we believe that the most likely source of misspecification is our stationary log-normal AR(1) model of wholesale steel prices. In particular our specification is unable to capture long periods of declining wholesale prices that we observe in the data. A stationary AR(1) price process results in very different predictions of the intertemporal pattern of steel purchases when there is a sustained period of decreasing wholesale prices (an event that is highly unlikely under the stationary AR specification) than would be predicted by a more complicated specification of price dynamics that includes "regime shift" variables that could capture more accurately the sustained trends ("drifts") in wholesale prices. Such a model might also more accurately capture the shift in John's expectations about steel prices and retail sales over the sample period, something we had the luxury to observe through our personal conversations with him.

In order to describe our model and the nature of the endogenous sampling problem more precisely, we need a bit more notation. In our model, we assume that a speculator takes the realized spot wholesale price p_t each day as given.[3] Our model also includes a vector of other state variables x_t that affect the firm's beliefs about future prices, sales, or other factors such as interest rates, holding costs, and so forth that affect the firm's optimal purchasing and inventory holding decisions. At the start of each business day the firm decides how much new inventory to order at the current spot price. The firm also sets the retail price p_t^r for its customers. Hall and Rust (2001) proved that the profit maximizing trading strategy takes the form of a *generalized* (S,s) *rule* in which S and s are functions of p and x. The function $s(p,x)$ is the *purchase threshold* and the function $S(p,x)$ is the *target inventory level* satisfying $S(p,x) \geq s(p,x)$. Under an (S,s) rule, the optimal purchase size is zero whenever the current inventory level, q, exceeds $s(p,x)$. However when q falls below $s(p,x)$ the firm makes a purchase of size $S(p,x) - q$, restoring inventory levels to the target level $S(p,x)$. Both $s(p,x)$ and $S(p,x)$ are decreasing functions of p, capturing the basic feature of successful price speculation, namely, that the firm should buy large quantities when prices are low and sell this inventory when prices are high.

[3]Since no centralized wholesale or retail markets for steel exist, there is a distribution of prices quoted by different producers and traders, and search effort is required to find the best available price at any given point in time. Our model abstracts from the details of this search problem, and we treat p_t as the *lowest* wholesale price that could be obtained via a reasonable amount of search effort on any given business day.

The order threshold function $s(p,x)$ is the source of the endogenous sampling problem since the firm only records the wholesale price p_t on those days when a purchase occurs. Therefore the endogenous sampling rule can be formalized as:

$$p_t \textit{ is observed if and only if } q_t < s(p_t, x_t). \tag{1}$$

Conditional on a purchase occurring, we observe an order of size q_t^o given by

$$q_t^o = S(p_t, x_t) - q_t, \tag{2}$$

and $q_t^o = 0$ otherwise. Using the generalized (S,s) rule as our model of the endogenous determination of sampling dates, our SMD estimator is able to consistently estimate the unknown parameters of the $\{p_t\}$ process even though we only have incomplete information on $\{p_t\}$.

In an earlier version of this paper (Hall and Rust, 2002) we derived a parametric partial-information maximum likelihood (PIML) estimator that solves the endogenous sampling problem and efficiently estimates the unknown parameters of the Markov law of motion for $\{p_t\}$ together with the structural parameters that determine the optimal trading rule. While the PIML estimator is consistent and efficient, it requires high dimensional numerical integrations that can only be feasibly done via recursive quadrature, or by Monte Carlo or quasi-Monte Carlo methods. The SMD estimator we employ is less efficient but computationally simpler. It relies only on the ability to simulate realizations of the optimal trading model. These simulations are then censored in exactly the same way as the observed data are censored, an approach that is similar in many respects to the strategy of "data augmentation" used in Bayesian inference for latent variable models. The idea behind the SMD estimator is to choose parameter values that result in simulated moments that match the observed moments as closely as possible, where both the real and simulated data are censored according to the same sampling rule; namely the one given in equations (1) and (2). Even though the moments entering the SMD criterion are biased and inconsistent estimates of the corresponding long term or "ergodic" moments due to the endogenous sampling problem, the fact that we can censor the simulations in the same way as the actual data are censored implies that the SMD estimator itself *is* consistent. Although the SMD estimator we employ is specialized to our particular steel example, it should be straightforward to generalize this method to other types of endogenous sampling problems that arise in a variety of other contexts.

Examples of where endogenous sampling issues arise include financial applications where transaction prices are observed at randomly spaced intervals (see Aït-Sahalia and Mykland, 2003, Engle and Russell, 1999, and Russell and Engle, 1998), and in marketing applications where the prices of goods that a household purchases are generally only recorded for the items the household purchased and on the dates it purchased them (see Allenby, McCulloch and Rossi 1996, Erdem and Keane, 1996, and Boizot, Robin, and Visser, 2001). The most directly related work is the literature on likelihood-based methods for correcting for endogenous sampling in cross-sectional and panel contexts (Heckman, 1981, Manski and McFadden, 1981, and McFadden, 1997).

Section 2 describes our data set and introduces the steel speculation and inventory problem that motivates this research. Section 3 presents the simulated minimum distance estimator and derives its asymptotic distribution. Section 4 present our estimates of the parameters of the price process and the parameters affecting the firm's cost of purchasing and holding inventory along with some Monte Carlo evidence on the performance of the SMD estimator. We evaluate how well our generalized (S,s) trading strategy fits the data, and use our results to infer the fraction of the firm's discounted profits are due to the markups it charges its retail customers, and the fraction that is due to pure commodity price speculation.

2 Description of the Data and the Model of Price Speculation

In this section we introduce the data and describe a generalized version of a model of commodity price speculation introduced by Hall and Rust (1999, 2001) that allows for additional covariates and unobserved state variables. This model provides the framework for inference and provides the key insights to pose and solve the endogenous sampling problem.

2.1 The Data

Through our contact with John, the general manager of a large U.S. steel wholesaler, we acquired a new high frequency micro database on transactions in the steel market. John's firm has provided us with daily data on all of its 2300+ individual steel products. Our empirical results are based on all transactions made between July 1, 1997 to June 2, 2003 (1500 business days) for two of its highest volume steel 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.

Although this is an exceptionally clean and rich dataset, we only observe prices on the days the firm actually made transactions: the firm does not record any price information on days that it does not transact (either as a buyer or seller of steel). This shortcoming of our dataset is much more important for steel purchases than steel sales, since John's firm purchases new steel inventory in the wholesale market much less frequently than it sells steel to its retail customers. Indeed, even for its highest volume products, it makes purchases only about once every two weeks. It is possible to get weekly and monthly survey data on prices for certain classes of steel products through trade publications such as *Purchasing Magazine* and *American Metal Market.* However, since there are no public exchange markets for steel products, nearly all transactions in the steel market are consummated via private negotiations. Hence the price surveys rely on truthful and accurate recollections of these transaction prices by participants. Also, the average transaction price is often very different from the best transaction price that John can obtain: in our conversations he reported that he could typically obtain significantly lower prices than the average price reported in the survey data.

We illustrate our data by plotting the time series of inventories and prices of one of the firm's products in figures 1 and 2. This product, which we call product 4, is one of the highest volume products sold by this firm. It is also a benchmark product within the industry since the prices of several other steel products are often computed as a function of this product's price. In our plot of wholesale transaction prices in figure 2 (the lower curve with the large black circles), we used straight line interpolations between observed purchase prices at successive purchase dates. The black circle at each purchase date is proportional to the size of the firm's purchase in pounds. This provides a visual indication of the endogenous sampling problem. First, even though we have 1500 daily observations, we observe purchases in the wholesale market on only 212 days. Second, the patterns of the black dots suggests that the firm is more likely to purchase large quantities of steel when wholesale prices are low, although other economic factors seem to be influencing the firm's purchase decisions as well. One key factor is the level of inventory: the firm tends to make large purchases when its inventory is low. We also see that even though wholesale prices declined in 2000 and 2001, the firm's largest purchases of steel occurred during the "turning point" in prices in early 1998. The firm also avoided making large purchases in late 2000 and 2001 due to economic uncertainties resulting from the "dot com crash" and the 9/11/2001 terrorist attack.

Overall, our interpolated plot of steel wholesale prices in figure 2 suggests that we should be wary of using the relatively small number of irregularly spaced observations to make inferences about the underlying law of motion for $\{p_t\}$. The observed purchase prices are unlikely to be representative of the unconditional mean level of prices in the wholesale market (especially if the firm is successful at buying low and selling high), and the estimated serial correlation coefficient for these irregularly spaced transactions is unlikely to be a good estimate of the serial correlation coefficient between daily wholesale prices.

Figure 2 also plots the interpolated sequence of daily retail sales prices. Retail sales occur on about two out every three business days, so the amount of interpolation in the retail price series is modest.

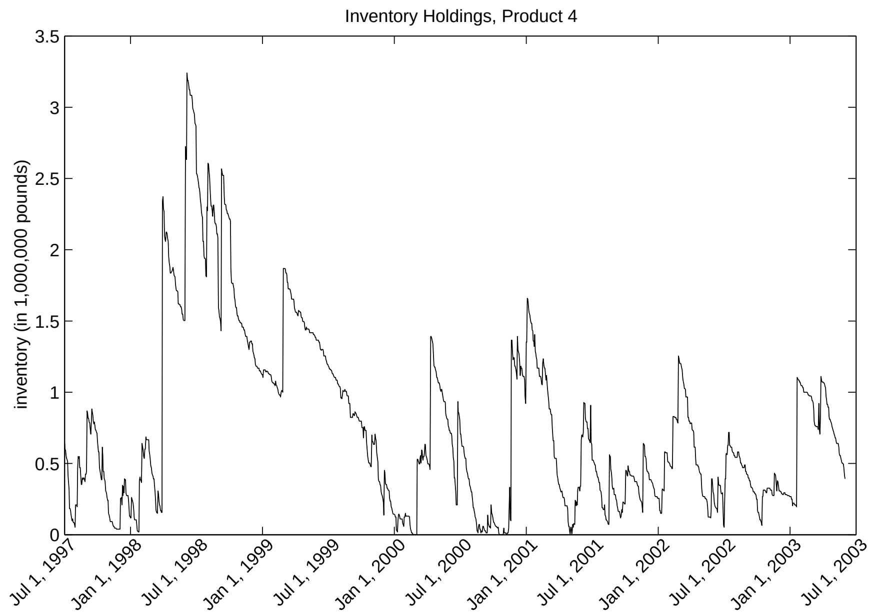

Figure 1: Times series plot of the inventory for product 4.

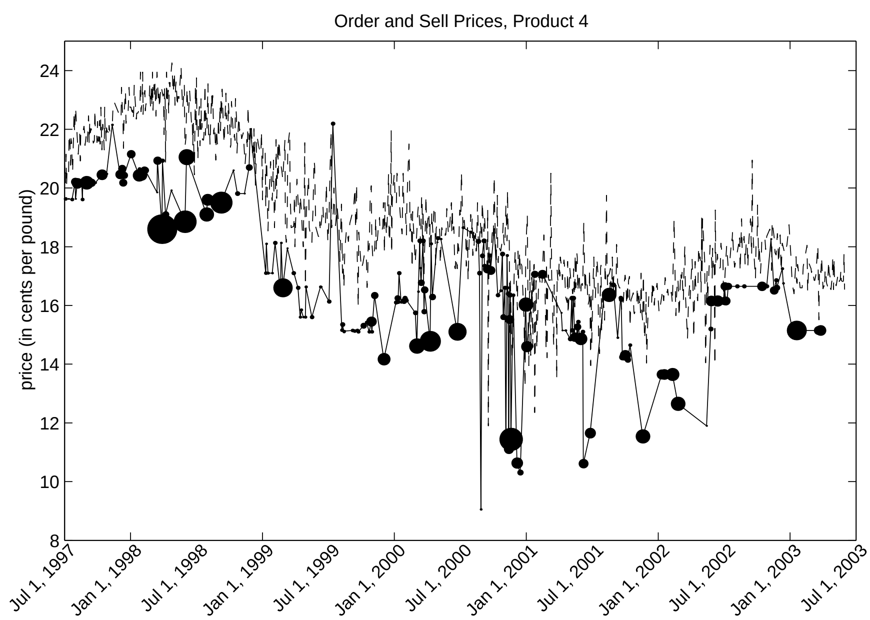

Figure 2: Purchase prices (solid line) and retail prices (dashed line) for product 4. For the purchase price series, the size of the marker is proportional to the size of the purchase.

The wholesale and retail prices move in a roughly parallel way, although there is considerable day-to-day variation in retail prices. Retail prices are quoted net of transportation costs, but still much of the the high frequency variation is due to observable factors. Athreya (2002) finds that roughly 65% of the high frequency variation in retail prices can be explained by observable customer characteristics such as geographical location and past volume of purchases. The remaining 35% appears to be due either to high frequency fluctuations in wholesale prices or to some sort of "informational price discrimination" in the retail market. Using the limited number of days on which both wholesale and retail prices are available, Chan (2002) finds that at most 50% of the variation in retail prices can be explained by variations in the wholesale price of steel. This conclusion is possible due to the fact that on many days there are multiple retail sales to different customers. These findings suggest that a large share of the high frequency variation in retail prices can be ascribed to price discrimination, i.e. the firm charges higher prices to more impatient or poorly informed retail customers.[4] Even though retail sales occur much more frequently than wholesale purchases, the fact that retail prices involve a number of other different considerations (including price discrimination based on observable and unobservable characteristics of the customer) suggest that the retail price is at best a very noisy and (upward) biased signal of the underlying wholesale price.

Figure 1 plots the evolution of inventories over the same period. Purchases of steel are easily recognizable as the discontinuous upward jumps in the inventory trajectories. As is evident from the saw-tooth pattern of the inventory holdings, the firm purchases the product much less frequently than it sells it. The firm's opportunistic purchasing behavior is very clear for this product. As can be seen in figures 1 and 2, during the first ten months of the sample, from July, 1997 until March, 1998, the firm held relatively low levels of inventories at a time when the average price the firm paid for steel was about 20.5 cents per pound. However as the Asian financial crisis deepened, foreign steel producers began cutting their prices and aggressively increasing their exports. We see this clearly in our data, where in April 1998, wholesale prices dropped to 18.5 cents per pound. At that time the firm made a large purchase. As the price of steel continued to fall to historical lows during the remainder of 1998 the firm made a succession of large purchases that lead it to hold historically unprecedented high levels of inventories. We view this as evidence that the firm is attempting to profit from a buy low, sell high strategy.

With 20/20 hindsight, it may appear that John's firm made too many large purchases in 1998 at prices between 18 and 20 cents per pound, given that in the subsequent period from 1999 to 2001 the firm was

[4]See Chan, Hall, and Rust (2003) for a more detailed analysis of bargaining, price setting, and price discrimination in the retail market for steel.

able to acquire inventories at prices as low as 12 cent per pound. Of course John did not have the luxury of hindsight when he made the *ex ante* decisions to purchase inventories in 1998. To provide some context, wholesale prices had been significantly higher than 20 cents per pound prior to 1997, before the start of our sample period. As the Asian crisis reverberated and foreign steel producers began to export large quantities of low price steel into the American market, John began to see lower wholesale prices than he had seen in recent memory. Many traders in the steel market believed that anti-dumping penalties would be quickly enacted and that the low prices that existed during 1998 and 1999 would soon be reversed. However the duties President Clinton put into effect in early 2000 did not seem to shut off the flood of foreign imports and the flow of cheap plate steel from efficient American “mini mills” such as Nucor continued to drive prices downward. Also the stock market crash in 2000 and the terrorist attack in 2001 resulted in cutbacks in demand for steel by end consumers, which combined with the large amount of steel already stockpiled collectively by U.S. steel service centers between 1998 and 2001 is probably responsible for the continued downward trend in the retail price of steel.

Our conversations with John during this period revealed that the “macro shocks” the U.S. experienced in 2000 and 2001 and the recession during this period lead him to take a more cautionary approach to acquiring steel, out of fear that a collapse in steel demand could lead to excessive steel inventories that would take a long time to liquidate, probably at a loss. As it happened, we can see that many of the firm’s retail sales during 1999 to 2001 occurred at prices that were *lower* than the wholesale prices that the firm paid for inventories during 1998. Thus, it is clear that John’s firm faces the downside risk inherent in speculation: while one desires to buy low and sell high, *ex post* there is always the risk that one might inadvertently have bought high and be forced to sell low, incurring potentially substantial losses. It appears that by adopting a more conservative approach to trading between 2000 and 2003, John’s firm was able to avoid incurring further losses.

2.2 The Model

Our model is an extension of previous work by Hall and Rust (2001), who showed that in a broad class of commodity price speculation problems, the optimal trading rule is a generalized version of the classic (S,s) rule from inventory theory. This work links contributions by Arrow *et. al.* (1951) and Scarf (1959) who first proved the optimality of (S,s) policies in inventory investment problems to more recent work by Williams and Wright (1991), Deaton and Laroque (1992) and Miranda and Rui (1997) on the rational expectations commodity storage model. The fixed (S,s) thresholds derived by Scarf under the assumption that the price (cost) of procuring (producing) inventories is constant are clearly suboptimal in a speculative trading environment, since the stochastic fluctuations in the price of steel affects the firm's perception of the optimal level of inventory S, and the threshold for purchasing new inventory s. Hall and Rust showed that the firm's optimal speculative trading strategy is a *generalized the* (S,s) *rule* where S and s are functions of certain underlying state variables including the wholesale price of steel p.[5] Before we describe how the generalized (S,s) rule allows us to formulate and solve the problem of endogenous sampling of steel wholesale prices, we describe the notation and key assumptions underlying the model of commodity price speculation. Then we formally define the (S,s) trading strategy.

We assume that John can purchase unlimited quantities of steel at a time-varying wholesale price p_t that evolves according to a Markov transition density. John subsequently sells this steel to retail customers at a retail price p_t^r. On each business day t the following sequence of actions occurs:

1. At the start of day t John knows his inventory level q_t, the current wholesale price p_t, and the values of the other state variables x_t.

2. Given (q_t, p_t, x_t) John procures additional inventory q_t^o for immediate delivery. Purchases must be non-negative:

$$q_t^o \geq 0 \tag{3}$$

3. Given (q_t, q_t^o, p_t, x_t) John sets a retail price p_t^r that is modeled as a random draw from a density $\gamma(p_t^r | q_t + q_t^o, p_t, x_t)$.

4. Given $(q_t, q_t^o, p_t, p_t^r, x_t)$ John observes a realized retail demand for his steel, q_t^r, modeled as a draw from a distribution $H(q_t^r | p_t, p_t^r, x_t)$ with a point mass at $q_t^r = 0$.

[5] This analysis extends previous results in the operations research literature such as Fabian *et. al.* (1959), Kingman (1969), Kalymon (1971), Golabi (1985), Song and Zipkin (1993), and Moinzadeh (1997) that prove the optimality of generalized versions of the (S,s) rule when the cost (price) of producing (procuring) new inventory fluctuates stochastically. While Hall and Rust (2001) are not the first to prove the optimality of generalized versions of the (S,s) rule, they build on the OR literature by making the connection between models of optimal inventory policies and models of storage and commodity prices. Moreover in the current paper we computationally solve and estimate our model. Thus we can formally compare the model's optimal policies to the inventory policies we see in the data. Besides the work noted above, the most closely related recent work that we are aware of is the ambitious paper by Aguirregabiria (1999) that models price and inventory decisions by a supermarket chain. A supermarket is similar to our steel wholesaler in that both types of firms hold inventories of a substantial number of different products, purchasing them in the wholesale market and selling their inventories at a markup to retail customers. The key difference is that prices in supermarkets are almost always posted so there is no direct price discrimination and there is presumably a larger "menu cost" to changing prices on a day by day basis. Aguirregabiria also did not directly address the endogenous sampling issue, using monthly price averages as proxies for underlying daily prices. For this reason we are unable to directly employ his innovative and ambitious approach to estimation.

5. John cannot sell more steel than he has on hand, so the actual quantity sold satisfies

$$q_t^s = \min[q_t + q_t^o, q_t^r]. \tag{4}$$

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

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

7. New values of (p_{t+1}, x_{t+1}) are drawn from a Markov transition density $g(p_{t+1}, x_{t+1}|p_t, x_t)$.

We abstract from delivery lags and assume that John's firm cannot backlog unfilled orders. Thus, whenever demand exceeds quantity on hand, the residual unfilled demand is lost. Thus, in addition to the censoring of the purchase and retail prices (p_t, p_t^r), we only observe a truncated measure of the firm's retail demand, i.e., we only observe the *minimum* of q_t^r and $q_t + q_t^o$ as given in equation (4). Since the quantity demanded has support on the $[0, \infty)$ interval, equation (4) implies that there is always a positive probability of a *stockout* given by:

$$\delta(q, p, p^r, x) = 1 - H(q|p^r, p, x). \tag{6}$$

Since retail sales occur much more frequently that purchases of new inventory, the retail sales price p_t^r provides an important source of information about the wholesale price p_t. Presumably for most transactions we should have $p_t^r \geq p_t$, reflecting nonnegative markups over the current wholesale price of steel. However as noted above markups vary considerably from day to day, so at best p_t^r is a biased and noisy indicator of the wholesale price p_t. We bypass some of the difficult issues associated with modeling endogenous price setting and price discrimination by adopting a "reduced-form" model of price setting. We model the daily average retail price as a draw from a conditional density $\gamma(p_t^r|q_t + q_t^o, p_t, x_t)$. This way of modeling prices is sufficiently flexible to be consistent with a variety of theories of bargaining and price discrimination by the firm.[6]

[6] Chan, Hall, and Rust (2003) solve a version of the model in which the firm chooses both q_t^o and p_t^r. In this case, the value function is no longer guaranteed to be K-concave, and the solution to the inventory problem may no longer be of the generalized (S, s) form. Solving this model takes considerably longer than the model presented here for two reasons. First, the Chan, Hall, and Rust model requires a two-dimensional optimization instead of an one-dimensional optimization at each iteration of the Bellman equation. Second, in models with endogenous price setting, the generalized (S, s) rule is not always guaranteed to be an optimal trading strategy. As a result we cannot restrict our search to the subclass of generalized (S, s) policies as we can when we solve the model presented here. This greatly increases the computational time required to solve models that incorporate either uniform endogenous price setting or in models of bargaining and price discrimination.

The firm's expected sales revenue function, $ES(p,q,x)$ is the conditional expectation of realized sales revenue $p^r q^r$ given the current wholesale price p, quantity on hand q, and the observed information variables x. The firm's retail sales on date t is a random draw q_t^r from a conditional distribution $H(q_t^r|p_t^r,p_t,x_t)$ that depends on the retail price quote p_t^r, the current wholesale price p_t, and the values of the other observed state variables x_t. We assume that there is a positive probability $\eta(p^r,p,x) = H(0|p^r,p,x)$ that the firm will not make any retail sales on a particular day, so H can be represented by

$$H(q^r|p^r,p,x) = \eta(p^r,p,x) + [1-\eta(p^r,p,x)]\int_0^{q^r} h(q|p^r,p,x)dq, \tag{7}$$

where h is a continuous strictly positive probability density function over the interval $[0,\infty)$. Given this stochastic "demand function", the firm's expected sales revenue $ES(p,q,x)$ is:

$$\begin{aligned} ES(p,q,x) &= E\{\tilde{p}^r\tilde{q}^s|p,q,x\} \\ &= E\{\tilde{p}^r E\{\min[q,\tilde{q}^r]|p^r,p,q,x\}|p,q,x\} \\ &= \int_0^\infty p^r[1-\eta(p^r,p,x)]\left[\int_0^q q^r h(q^r|p^r,p,x)dq^r + \delta(q,p^r,p,x)q\right]\gamma(p^r|q,p,x)dp^r. \end{aligned} \tag{8}$$

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

$$c^o(q^o,p) = \begin{cases} pq^o + K & \text{if} \quad q^o > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

The firm's remaining costs are summarized by the *holding cost function* $c^h(q,p,x)$. These costs include physical storage costs, and "goodwill costs" representing the present value of lost future business from customers whose orders cannot be filled due to a stockout. Goodwill costs can be viewed as the inverse of the "convenience yield" discussed in the commodity storage literature (Kaldor, 1939, Williams and Wright, 1991). In this case a convenience yield emerges from a desire to hold a buffer stock or precautionary level of inventories in order to minimize goodwill costs from stockouts. This allows the model to capture other reasons besides pure price speculation for holding inventories.[7] The firm's single-period profits π equals its sales revenues, less the cost of new orders for inventory $c^o(q^o,p)$ and inventory holding costs $c^h(q+q^o,p,x)$:

$$\pi(p,p^r,q^r,q+q^o,x) = p^r q^s - c^o(q^o,p) - c^h(q+q^o,p,x). \tag{10}$$

[7]The firm obtains much of its steel from foreign sources. In the model orders occur instantaneously with certainty. In practice, however, delivery lags can be several months and the steel delivered can often be of lower quality than agreed on. The firm does have the option of refusing to take delivery if the steel is not of the quality promised. Having a buffer stock of inventories on hand reduces the cost to firm of exercising this option. Also foreign producers of steel do from time to time renege on previously negotiated deals, failing to deliver the amount of steel originally promised.

where $q^s = min[q^r, q + q^o]$. Each period the firm chooses investment q^o_t given $\{p_t, q_t, x_t\}$ to maximize the discounted present value of profits:

$$V(p_t, q_t, x_t) = \max_{q^o} E\left\{\sum_{j=t}^{\infty} \rho^{(j-t)} \pi(p_j, p^r_j, q^r_j, q^o_j + q_j, x_j) \Big| p_t, q_t, x_t\right\}, \tag{11}$$

where $\rho = 1/(1+r)$ and r is the firm's discount rate. The value function $V(p,q,x)$ is given by the unique solution to Bellman's equation:

$$V(p,q,x) = \max_{0 \leq q^o \leq \overline{q} - q} \Big[W(p, q+q^o, x) - c^o(q^o, p)\Big], \tag{12}$$

where $\overline{q}$ is the firm's maximum storage capacity and

$$W(p,q,x) \equiv \Big[ES(p,q,x) - c^h(q,p,x) + \rho EV(p,q,x)\Big], \tag{13}$$

and EV denotes the conditional expectation of V given by:

$$\begin{aligned} EV(p,q,x) &= E\{V(\tilde{p}, \max[0, q - \tilde{q}^r], \tilde{x}) | p, q, x) \\ &= \lambda_1(p,q,x) \int_{p'} \int_{x'} V(p', q, x') g(p', x' | p, x) dp' dx' \\ &+ \lambda_2(p,q,x) \int_{p'} \int_{x'} V(p', 0, x') g(p', x' | p, x) dp' dx' \\ &+ \lambda_3(p,q,x) \int_{p'} \int_{x'} \int_0^q V(p', q - q', x') h(q' | p, q, x) g(p', x' | p, x) dq' dp' dx', \end{aligned} \tag{14}$$

where

$$\begin{aligned} \lambda_1(p,q,x) &= \int_{p^r} \eta(p^r, p, x) \gamma(p^r | p, q, x) dp^r \\ \lambda_2(p,q,x) &= \int_{p^r} [1 - \eta(p^r, p, x)] \delta(p^r, p, q, x) \gamma(p^r | p, q, x) dp^r \\ \lambda_3(p,q,x) &= \int_{p^r} [1 - \eta(p^r, p, x)] \gamma(p^r | p, q, x) dp^r \\ h(q' | p, q, x) &= \int_{p^r} h(q' | p^r, p, q, x) \gamma(p^r | p, q, x) dp^r. \end{aligned} \tag{15}$$

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

$$q^o(p,q,x) = \inf \underset{0 \leq q^o \leq \overline{q} - q}{argmax} \Big[W(p, q + q^o, x) - c^o(q^o, p)\Big]. \tag{16}$$

We invoke the inf operator in the definition of the optimal decision rule in equation (16) 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* of the optimizing values of q^o.

In this model the variables q and q^o do not enter as separate arguments in the value function W given in (13): rather they enter as the sum $q+q^o$ as shown in equation (16). This 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 inventory on hand after new orders q^o have arrived. Thus if the firm is holding less than its desired level of inventories $S(p_t,x_t)$ at the start of day t, it will only have to order the amount $q^o(p,q,x) = S(p,x) - q$ in order to achieve its target inventory level $S(p,x)$. 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:

$$\frac{\partial W}{\partial q^o}(p, q+q^o, x) = p. \tag{17}$$

If W were strictly concave in q, there would be a unique value of $q+q^o$ that solves equation (17) for any value of p. Call this solution $S(p,x)$:

$$\frac{\partial W}{\partial q^o}(p, S(p,x), x) = p. \tag{18}$$

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

In turns out that if $K > 0$ the function $W(p,q,x)$ will not be strictly concave. However under fairly general conditions W is *K-concave* as a function of q for each fixed p.[8] Using the K-concavity property we can prove that whenever $q \geq s(p,x)$, it is not optimal to order: $q^o(p,q,x) = 0$. When $q < s(p,x)$ the symmetry property implies that $q^o(p,q,x) = S(p,x) - q$ as discussed above. In particular Hall and Rust (2001) proved:

Theorem 1: *Consider the function $W(p,q+q^o,x)$ defined in equation (13), where W is defined in terms of the unique solution V to Bellman's equation (12). Under appropriate regularity conditions given in Hall and Rust (2001), the optimal speculative trading strategy $q^o(p,q,x)$ takes the form of an (S,s) rule. That is, there exist a pair of functions (S,s) satisfying $S(p,x) \geq s(p,x)$ where $S(p,x)$ is the desired or target inventory level and $s(p,x)$ is the inventory order threshold, i.e.*

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

where $S(p,x)$ is given by:

$$S(p,x) = argmax_{0 \leq q^o \leq \overline{q} - q} \Big[W(p, q^o, x) - c^o(q^o, p) \Big] \tag{20}$$

[8]A function $W(p,q) : [\underline{p}, \overline{p}] \times [0, \overline{q}] \to R$ is *K-concave* in its second argument q if and only if $-W(p,q)$ is *K-convex* in its second argument. More directly, $W(p,q)$ is *K-concave* in q iff $\exists K \geq 0$ such that for every $p \in [\underline{p}, \overline{p}]$, and for all $z \geq 0$ and $b \geq 0$ such that $q+z \leq \overline{q}$ and $q-b \geq 0$ we have $W(p,q+z) - K \leq W(p,q) + z[W(p,q) - W(p,q-b)]/b$.

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

$$s(p,x) = \inf_{q \geq 0} \{q | W(p,q,x) - pq \geq W(p,S(p,x),x) - pS(p,x) - K\}. \tag{21}$$

By a simple substitution of the generalized (S,s) rule in equation (19) into the definition of V in equation (12) we obtain the following corollaries:

Corollary 1: *The value function V is linear with slope p on the interval $[0,s(p,x)]$:*

$$V(p,q,x) = \begin{cases} W(p,S(p,x),x) - p[S(p,x) - q] - K & \text{if} \quad q \in [0,s(p,x)] \\ W(p,q,x) & \text{if} \quad q \in (s(p,x),\overline{q}]. \end{cases} \tag{22}$$

Corollary 2: *The $S(p,x)$ and $s(p,x)$ functions are non-increasing in p and are strictly decreasing in p in the set $\{p|0 < S(p,x) < \overline{q}\}$.*

Corollary 3: *If fixed costs of ordering is zero, $K = 0$, then the minimum order size is 0 and*

$$S(p,x) = s(p,x). \tag{23}$$

Theorem 1 does yield a first order condition that could possibly provide a basis for a generalized methods of moments (GMM) strategy for estimating the unknown parameters of the model.

$$\frac{\partial W}{\partial q}(p,S(p,x),x) - p = 0. \tag{24}$$

However, the existence of the frequently binding inequality constraint on inventory investment, equation (3), complicates the use of standard Euler equation methods to estimate the unknown parameters of the model via GMM.

If we assume that there is additive measurement error ε in the wholesale price p, or assume that ε represents other unobserved (per unit) components of the cost of ordering new inventory, then it is tempting to treat equation (24) as an "Euler equation" and use GMM to estimate parameters of the model. However there are several big obstacles to this approach. First, we do not have a convenient analytical formula for the partial derivative of the value function, $\partial W/\partial q$. Second, even if the unconditional mean of ε is zero, the conditional mean of ε over those values of (p,ε) for which it is optimal to purchase (i.e. for which $q < s(p,x)$), is generally nonzero. Finally, there is the issue of endogenous sampling, and the fact that we observe purchases only an a relatively small subset of business days in our overall sample. These problems motivate the use an alternative estimation approach that is capable of incorporating other information such as retail sales prices in order to improve our ability to make inferences about the $\{p_t\}$ process.

3 The Simulated Minimum Distance Estimator

This section introduces a *simulated minimum distance estimator* (SMD) that may be less efficient than maximum likelihood, but which does not require the high dimensional integration and is much easier to compute.[9] Similar estimators have been proposed in other contexts by Lee and Ingram (1991) and Duffie and Singleton (1993). The idea behind the SMD estimator is quite straightforward, and is similar in spirit to the method of "calibration". The main difference is that the SMD estimator is based on an explicit statistical criterion function that enables us to compute asymptotic distributions for the parameter estimator, evaluate the fit of alternative specifications, and to conduct goodness of fit tests.

The SMD estimator is simply the parameter value that minimizes the distance between a set of simulated and sample moments using the observed censored observations. First we calculate sample moments using the censored observations in the data, i.e. with $p_t = 0$ when $q_t^o = 0$. Then we generate one or more simulated realizations of the (S,s) model for a given trial value θ of the unknown parameter vector. We define $\hat{\theta}_{smd}$ as the value of θ that minimizes a quadratic form in the difference between the sample moments for the actual data and the sample moments of the simulated data, where the simulated data has been censored in exactly the same fashion as the actual data, i.e. we set $p_t = 0$ whenever the simulated value of $q_t^o = 0$. Thus even though various moments based on censored data may be biased and inconsistent estimators of the corresponding moments of the ergodic process in the absence of censoring, this does not prevent us from deriving a consistent SMD estimator for θ^*. In fact the SMD estimator is consistent even if we use only a single simulated realization of the (S,s) model.

The asymptotic variance of the SMD estimator is multiplied by a factor $(1 + 1/S)$ where S is the number of simulations. Consequently, there is an efficiency gain to running additional simulations since it reduces the variance of the estimator. However the "penalty" to forming an SMD estimator based on only a single realization appears relatively small: the asymptotic variance is only twice as large as the variance of an estimator that eliminates all simulation noise by letting $S \to \infty$. This increase in variance seems small in comparison to the substantial reduction in computational burden from using only a single simulation of the model. Estimation still requires a nested fixed point algorithm to solve for the optimal (S,s) policy and a re-simulation of the model using a fixed set of uniform random "seed" variates (see below) each time the parameter θ is updated, so the SMD estimator is still fairly computationally demanding. Its other

[9]See Hall and Rust (2002) for a derivation of the "partial information" maximum likelihood estimator and a discussion of its asymptotic properties in the case where only time series information is available for a single commodity trader.

drawback is that it requires the analyst to determine an appropriate set of moments to represent the relevant metric for assessing the distance between the predictions of the model and the data. In principle an infinite number of different moment conditions could be specified, but only a finite number can be used in practice.

Let the conditional density of next period's inventory q_{t+1} given $(p_t, p_t^r, x_t, q_t, q_t^o)$ be denoted by μ. From our discussion of the model in section 2, it is easy to see that μ is a mixed discrete/continuous density with three classes of outcomes for q_{t+1}: 1) with probability $\eta(p^r, p, x)$ the firm will not make any sales and $q_{t+1} = q_t + q_t^o$; 2) with probability $(1-\eta(p_t^r, p_t, x_t))\delta(p_t^r, p_t, q_t + q_t^o, x_t)$ the firm will have a stockout and $q_{t+1} = 0$; 3) otherwise q_{t+1} is distributed continuously over the interval $(0, q_t + q_t^o)$ with density given by $(1-\eta(p_t^r, p_t, x_t))h(q_t + q_t^o - q_{t+1}|p_t^r, p_t, x_t)$ where h is the density of retail sales and $q_t^r = q_t + q_t^o - q_{t+1}$ is the implied value of retail sales given (q_{t+1}, q_t, q_t^o). We summarize this as:

Theorem 2: *The (mixed discrete/continuous) density of next period inventory q' given (p, p^r, q, q^o, x) is given by:*

$$\mu(q'|p, p^r, q, q^o, x) = \begin{cases} (1-\eta(p^r, p, x))\delta(p^r, p, q+q^o, x) & \text{if} \quad q' = 0 \\ \eta(p^r, p, x) & \text{if} \quad q' = q + q^o \\ (1-\eta(p^r, p, x))h(q+q^o - q'|p^r, p, x) & \text{otherwise} \end{cases} \tag{25}$$

Under our setup, we can show that the observables $\{p_t, p_t^r, q_t, q_t^o, x_t\}$ evolve as a joint Markov process which also has a discrete/continuous transition probability density λ. We state this as Theorem 3:

Theorem 3: *The joint process $\{p_t, p_t^r, q_t, q_t^o, x_t\}$ is Markov with a degenerate discrete/continuous transition density λ given by:*

$$\begin{aligned} \lambda(p_{t+1}, p_{t+1}^r, q_{t+1}, q_{t+}^o, x_{t+1}|p_t, p_t^r, q_t, q_t^o, x_t) &= g(p_{t+1}, x_{t+1}|p_t, x_t) \\ &\times \mu(q_{t+1}|p_t, p_t^r, q_t, q_t^o, x_t) \\ &\times \gamma(p_{t+1}^r|p_{t+1}, q_{t+1} + q_{t+1}^o, x_{t+1}) \\ &\times I\{q_{t+1}^o q^o(p_{t+1}, q_{t+1}, x_{t+1})\}. \end{aligned} \tag{26}$$

Note that the degeneracy in the transition probability λ is a result that conditional on the (stochastic) realizations of $(p_{t+1}, q_{t+1}, x_{t+1})$, the optimal order quantity q_{t+1}^o is a deterministic function of these variables, given by the optimal decision rule given in equation (19) in Theorem 1. Although it is possible to introduce an "unobserved state variable" ε_t into the problem so that q_{t+1}^o, as a function of ε_t, is nondegenerate (i.e. random) from the standpoint of the econometrician, we do not attempt to do this here. The deterministic prediction of optimal order quantities from the (S, s) rule will provide an extremely strong test of the validity of the model.[10] The degeneracy has implications for the properties of the estimator, which we discuss further below. For maximum likelihood estimation, the degeneracy creates severe problems due to a "zero likelihood problem" for any observation (q_t^o, q_t, p_t, x_t) that does not lie on the graph of the optimal decision rule (an event that is almost certain to occur in the absence of unobserved state variables ε_t). The SMD estimator is more tolerant of the degeneracy caused by failure to include unobserved state variables ε_t in the decision rule. We will show below that this leads to a situation where discontinuities can occur in SMD objective function, but asymptotically these discontinuities disappear and the SMD objective converges to a smooth function of the parameters θ.

Assumption 1: *The Markov process $\{\xi_t\} = \{p_t, p_t^r, q_t, q_t^o, x_t)\}$ is ergodic (i.e. it possesses a unique stationary distribution).*

Assumption 1 can be proven from more primitive conditions on the structural objects in the speculation model, and it can be shown to hold for the specification we estimate in this paper. The key to proving ergodicity is to establish that the transition density λ satisfies a stochastic equicontinuity condition, a generalization of the "Doeblin condition" in the theory of Markov chains (see, e.g. Futia, 1982). In our case, these conditions can be established from an assumption that 1) wholesale and retail prices are truncated and remain in a compact interval and do not have any "absorbing states", and 2) speculation is a sufficiently profitable activity that the firm would never want to exit the market (or hold zero inventory perpetually). However delving into the details of such a proof would distract from the main issues in this paper and we simply assume that ergodicity holds, since it is quite easy to verify from simulations of the model when conditions are such that the ergodicity condition fails.

Let θ denote the $L \times 1$ vector of parameters to be estimated. The SMD estimator is based on finding a parameter value that best fits a $J \times 1$ vector of moments of the observed process:

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

where $J \geq K$ and h is a known (smooth) function of (ξ_t, ξ_{t-1}) that determines the moments we wish to match. We include ξ_t and its lag ξ_{t-1} as arguments of h in order to handle situations were we are trying to fit moments such as means and covariances of the components of ξ_t. It is straightforward to allow moments that involve more than one lag: we only include a single lagged value of ξ_t in our presentation below for notational simplicity.

[10]See Hall and Rust (2002) for an example of how to include unobserved state variables to result in a nondegenerate transition density λ.

By Assumption 1, the process $\{\xi_t\}$ is ergodic so that, with probability 1, h_T converges to a limit $E\{h(\xi',\xi)\}$ where the expectation is taken with respect to the ergodic distribution of (ξ',ξ) (i.e. the limiting distribution of (ξ_{t+1},ξ_t) as $t \to \infty$). Under suitable additional regularity conditions, a central limit theorem will hold for h_T, i.e. we have

$$\sqrt{T}\,[h_T - E\{h\}] \Longrightarrow N(0,\Omega(h)), \tag{28}$$

where

$$\Omega(h) = E\left\{(h(\xi',\xi) - E\{h\})(h(\xi',\xi) - E\{h\})'\right\}, \tag{29}$$

where the expectation in (29) is taken with respect to the ergodic distribution of (ξ',ξ).

Now assume it is possible to generate simulated realizations of the $\{\xi_t\}$ process for any candidate value of θ, and that this process is censored in exactly the same way as the observed $\{\xi_t\}$ process is censored, i.e., with $p_t = 0$ when $q_t^o = 0$. These simulations depend on a $T \times 1$ vector, u, of *IID* $U(0,1)$ random variables that are drawn once at the start of the estimation process and held fixed thereafter in order for the estimator to satisfy stochastic equicontinuity conditions necessary to establish asymptotic normality of the SMD estimator. We will consider simulated processes of the form

$$\{\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0)\}, \quad t = 2,\ldots,T \tag{30}$$

where for each $t > 1$, $\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0)$ is a function of θ. The notation $\{u_s\}_{s\leq t}$ reflects the fact that the simulated process is *adapted* to the realization of the $\{u_t\}$ process, i.e. the first t realized values of $\{\xi_t(\{u_s\}_{s\leq t},\theta)\}$ depend only on the first t realized values of $\{u_s\}$ and not on subsequent realized values of u_s for $s > t$. Note that we allow the simulated process to depend on the first value ξ_0 of the observed data as an initializing condition.

To show that it is possible to construct *smooth* (i.e. continuously differentiable) simulators, consider the unidimensional case where $\xi_t \in R^1$ for all t. Let $\lambda(\xi_{t+1}|\xi_t,\theta)$ denote its transition density and $P(\xi_{t+1}|\xi_t,\theta)$ be the corresponding conditional CDF. The first value of the simulated process is simply set to the observed value ξ_0. Using the probability integral transform, we can define $\xi_1(u_1,\theta,\xi_0)$ by:

$$\xi_1(u_1,\theta,\xi_0) = P^{-1}(u_1|\xi_0,\theta). \tag{31}$$

Clearly $\xi_1(u_1,\theta,\xi_0)$ will be a continuously differentiable function of θ if $P^{-1}(u_1|\xi_0,\theta)$ is a continuously differentiable function of θ. Now define recursively for $t = 2,4,\ldots$

$$\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0) = P^{-1}(u_t|\xi_{t-1}(\{u_s\}_{s\leq t-1},\theta,\xi_0),\theta). \tag{32}$$

We can see recursively that $\xi_t(\{u_s\}_{s\le t},\theta,\xi_0)$ will be a continuously differentiable function of θ provided that $P^{-1}(u|\xi,\theta)$ is a continuously differentiable function of ξ and θ.

In the case where $\{\xi_t\}$ is the multidimensional process with $\xi_t=(p_t,p_t^r,q_t,q_t^o,x_t)$, we can do a similar simulation as in the univariate case described above, using a factorization of the transition density of $\{x_t\}$ into a product of univariate conditional densities such as given in Theorem 3. For example, if ξ_t has two components, $\xi_t=(\xi_{1,t},\xi_{2,t})$, suppose that its transition density λ can be factored as

$$\lambda(\xi_{t+1}|\xi_t,\theta)=\lambda_2(\xi_{2,t+1}|\xi_{1,t+1},\xi_t,\theta)\lambda_1(\xi_{1,t+1}|\xi_t,\theta), \tag{33}$$

with corresponding conditional CDFs denoted by P_1 and P_2. Now we can generate simulations of $\{\xi_t\}$ that will be smooth function of θ just as in the univariate case, except that in the two-dimensional case we need to generate two random $U(0,1)$ variables $u_t=(u_{1,t},u_{2,t})$ for each time period simulated. For example to generate a simulated value of $\xi_1=(\xi_{1,1},\xi_{2,1})$ we compute

$$\begin{aligned}\xi_{1,1} &= P_1^{-1}(u_{1,1}|\xi_0,\theta)\\ \xi_{2,1} &= P_2^{-1}(u_{2,1}|\xi_{1,1},\xi_0,\theta).\end{aligned} \tag{34}$$

Clearly, the resulting realization for ξ_1 is of the form $\xi_1(u_1,\xi_0,\theta)$ and will be a smooth function of θ provided that P_1 and P_2 are smooth functions of (ξ,θ). Continuing recursively we have:

$$\begin{aligned}\xi_{1,t+1} &= P_1^{-1}(u_{1,t+1}|\xi_t,\theta)\\ \xi_{2,t+1} &= P_2^{-1}(u_{2,t+1}|\xi_{1,t+1},\xi_t,\theta).\end{aligned} \tag{35}$$

The resulting simulations take the form $\{\xi_t(\{u_s\}_{s\le t},\theta,\xi_0)\}$ and will be smooth functions of θ provided that P_1 and P_2 are smooth functions of their conditioning arguments (ξ,θ).[11]

[11] Note that the degeneracy of the transition density λ given in Theorem 3 implies that the simulated trajectories will not necessarily be smooth. The reason is that the (S,s) policy involves an inherent discontinuity that will be reflected in the simulations: for q_t just above the $s(p_t)$ order threshold the optimal order is zero, but as q_t moves just below the order threshold, the optimal order jumps discontinuously from 0 to $S(p_t)-s(p_t)$. As we will see subsequently, the gap between the $S(p)$ and $s(p)$ bands for reasonable parameter values is not large. Also as the sample size T tends to infinity, the impact of such discontinuities in individual realizations on the simulated moments (which are time averages) tends to zero. In the limit the expectation of the simulated moment is a smooth function of the parameters, a consequence of the fact that the only discontinuity in the model is for (p,q) points that lie on $q=s(p)$ band, but as long as prices and quantities have continuous distributions, this set of discontinuities has measure zero, and the Lebesgue dominated convergence theorem can be used to prove that the expected moments $E\{h|\theta\}$ (defined in equation (41) below) are smooth functions of θ even though individual simulated realizations of the process $\{\xi_t(\{u_s\}_{s\le t},\theta,\xi_0)\}$ can be discontinuous in θ. Further, we will show via the Monte Carlo study in section 4 that even in small samples, the discontinuities in the simulated realizations do not seem to affect the asymptotic properties of the SMD estimator. The reasons for this are well understood and we will avoid repeating the details here. See for example Pakes and Pollard (1989) for details on the machinery to prove asymptotic normality of simulated method of moment estimators when there are discontinuities in the objective function in finite samples.

Now consider using a single simulated realization of $\{\xi_t(\{u_s\}_{s\le t},\theta,\xi_0)\}$ to form a simulated sample moment $h_T(\{u_s\}_{s\le T},\xi_0,\theta)$ given by

$$h_T(\{u_s\}_{s\le T},\xi_0,\theta)=\frac{1}{T}\sum_{t=1}^{T}h\left(\xi_t(\{u_s\}_{s\le t},\theta,\xi_0),\xi_{t-1}(\{u_s\}_{s\le t-1},\theta,\xi_0)\right). \tag{36}$$

Let $(\{u_s^1\}_{s\le T},\ldots,\{u_s^S\}_{s\le T})$ denote S *IID* $T\times 1$ sequences of $U(0,1)$ random vectors used to generate the S independent realizations of the endogenously sampled process $\{\xi_t(\{u_s^i\}_{s\le t},\theta,\xi_0)\}$, $i=1,\ldots,S$. Define $h_{S,T}(\theta)$ as the average of S independent time averages $h_T(\{u_s^i\}_{s\le T},\xi_0,\theta)$

$$h_{S,T}(\theta)=\frac{1}{S}\sum_{i=1}^{S}h_T(\{u_s^i\}_{s\le T},\xi_0,\theta). \tag{37}$$

Definition 1: *The* **simulated minimum distance estimator** $\hat{\theta}_T$ *is defined by:*

$$\hat{\theta}_T=\underset{\theta\in\Theta}{argmin}(h_{S,T}(\theta)-h_T)'W_T(h_{S,T}(\theta)-h_T), \tag{38}$$

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

In order to simplify the asymptotic analysis, we initially assume that we have a correct parametric specification of the endogenous sampling problem. That is we make

Assumption 2: *The parametric model introduced in section 2 is correctly specified, i.e., there is a $\theta^*\in\Theta$ such that:*

$$\{\xi_t(\{u_s\}_{s\le t},\theta^*,\xi_0)\}\sim\{\xi_t\} \tag{39}$$

that is, when $\theta=\theta^*$, the simulated sequence initialized from the observed value ξ_0 has the same probability distribution as the observed sequence $\{\xi_t\}$.

We believe that it is possible to relax assumption 2 to allow the parametric model to be misspecified, following an analysis similar to that of Hall and Inoue (2002) who characterized the asymptotic properties of the GMM estimator in the misspecified case. We conjecture that their analysis will also apply to the case of SMD estimation and that the asymptotic properties of the SMD estimator that we derive for the correctly specified case will still hold, except that now θ^* is interpreted as the value of θ the minimizes the distance between the moments of the true data generating process and the parametric simulated process, where the expectation is taken in the limit as both $S\to\infty$ and $T\to\infty$.[12]

[12]When there is misspecification, the standard formula for the asymptotic covariance matrix when the model is correctly specified will generally not be consistent when the model is misspecified. However similar to the case of maximum likelihood estimation of misspecified models (White, 1982), there are alternative estimators of the asymptotic covariance matrix that are consistent when the model is misspecified.

We now sketch the derivation of the asymptotic distribution of the SMD estimator, listing the key assumptions and showing how its asymptotic variance depends on the number of simulations S.

Assumption 3: *For any* $\theta \in \Theta$ *the process* $\{\xi_t(\{u_s\}_{s \le t}, \theta, \xi_0)\}$ *is ergodic with unique invariant density* $\psi(\xi|\theta)$ *given by:*

$$\psi(\xi'|\theta) = \int \lambda(\xi'|\xi,\theta) d\psi(\xi|\theta). \tag{40}$$

Define the functions $E\{h|\theta\}$, $\nabla E\{h|\theta\}$, and $\nabla h_{S,T}(\theta)$ by:

$$\begin{aligned} E\{h|\theta\} &= \int h(\xi',\xi) d\lambda(\xi'|\xi,\theta) d\psi(\xi|\theta) \\ \nabla E\{h|\theta\} &= \frac{\partial}{\partial\theta} E\{h|\theta\} \\ \nabla h_{S,T}(\theta) &= \frac{\partial}{\partial\theta} h_{S,T}(\theta). \end{aligned} \tag{41}$$

Assumption 4: θ^* *is* **identified***; that is, if* $\theta \ne \theta^*$*, then* $E\{h|\theta\} \ne E\{h|\theta^*\} = E\{h\}$*. Furthermore,* $rank(\nabla E\{h|\theta\}) = L$ *and* $\lim_{T\to\infty} W_T = W$ *with probability 1 where* W *is a* $J \times J$ *positive definite matrix.*

The consistency of the SMD estimator can be established by providing appropriate regularity conditions under which the simulated process is uniformly ergodic, i.e., under which with probability 1 we have

$$\lim_{T\to\infty} \sup_{\theta\in\Theta} \left| (h_{S,T}(\theta) - h_T)' W_T (h_{S,T}(\theta) - h_T) - (E\{h|\theta\} - E\{h|\theta^*\})' W (E\{h|\theta\} - E\{h|\theta^*\}) \right| = 0. \tag{42}$$

Assumption 3 guarantees that the unique minimizer of $(E\{h|\theta\} - E\{h|\theta^*\})' W (E\{h|\theta\} - E\{h|\theta^*\})$ is θ^*, and this combined with the uniform consistency result implies the consistency of $\hat\theta_T$. The asymptotic normality of $\hat\theta_T$ can be established by a Taylor series expansion of the first order condition

$$(h_{S,T}(\hat\theta_T) - h_T)' W_T \nabla h_{S,T}(\hat\theta_T) = 0. \tag{43}$$

Expanding $h_{S,T}(\hat\theta_T)$ about $\theta = \theta^*$ we have

$$h_{S,T}(\hat\theta_T) = h_{S,T}(\theta^*) + \nabla h_{S,T}(\tilde\theta_T)(\hat\theta_T - \theta^*), \tag{44}$$

where $\tilde\theta_T$ denotes a vector that is (elementwise) on the line segment between $\hat\theta_T$ and θ^*. Substituting (44) into the first order condition for $\hat\theta_T$ in equation (43) and solving for $(\hat\theta_T - \theta^*)$ we obtain

$$(\hat\theta_T - \theta^*) = -\left[\nabla h_{S,T}(\tilde\theta_T)' W_T \nabla h_{S,T}(\hat\theta_T)\right]^{-1} \nabla h_{S,T}(\hat\theta_T)' W_T \left[h_{S,T}(\theta^*) - h_T\right], \tag{45}$$

where we assume that $\left[\nabla h_{S,T}(\tilde{\theta}_T)' W_T \nabla h_{S,T}(\hat{\theta}_T)\right]$ is invertible, which will be the case with probability 1 for sufficiently large T due to assumptions 3 and 4. Now multiply both sides of equation (45) by $\sqrt{T}$ and apply a Central Limit theorem to the difference $\sqrt{T}[h_{S,T}(\theta^*) - h_T]$ to obtain

$$\sqrt{T}\,[h_{S,T}(\theta^*) - h_T] \Longrightarrow N\left(0, (1+1/S)\Omega(h,\theta^*)\right). \tag{46}$$

To understand this result, note that $h_{S,T}(\theta^*)$ is an average of S independent realizations of $\{\xi_t(\{u_s\}_{s\leq t}, \theta, \xi_0)\}$, which by assumption 2 has the same distribution as $\{\xi_t\}$. As a result each of the terms entering $h_{S,T}(\theta^*)$, $h_T(\{u_s^i\}_{s\leq T}, \theta^*)$, has the same probability distribution as h_T and are distributed independently of h_T. The Central Limit Theorem applied to h_T yields

$$\sqrt{T}\,[h_T - E\{h|\theta^*\}] \Longrightarrow N\left(0, \Omega(h,\theta^*)\right). \tag{47}$$

Similarly, for each $i = 1, \ldots, S$ we have

$$\sqrt{T}\left[h_T(\{u_s^i\}_{s\leq T}, \theta^*) - E\{h|\theta^*\}\right] \Longrightarrow N\left(0, \Omega(h,\theta^*)\right). \tag{48}$$

Note that

$$[h_{S,T}(\theta^*) - h_T] = \left[\frac{1}{S}\sum_{i=1}^{S}[h_T(\{u_s^i\}_{s\leq T}, \theta^*) - E\{h|\theta^*\}] + E\{h|\theta^*\} - h_T\right], \tag{49}$$

so that we have

$$\sqrt{T}\,[h_{S,T}(\theta^*) - h_T] \Longrightarrow \left[\frac{1}{S}\sum_{i=1}^{S}\tilde{X}_i + \tilde{X}_0\right], \tag{50}$$

where $(\tilde{X}_0, \tilde{X}_1, \ldots, X_S)$ are *IID* $N(0, \Omega(h,\theta^*))$ random vectors. It follows immediately that the asymptotic distribution of $\sqrt{T}\,[h_{S,T}(\theta^*) - h_T]$ is $N(0, (1+1/S)\Omega(h,\theta^*))$. Using this result and equation (45) we have

$$\sqrt{T}\left[\hat{\theta}_T - \theta^*\right] \Longrightarrow N(0, (1+1/S)\Lambda_1^{-1}\Lambda_2\Lambda_1^{-1}), \tag{51}$$

where

$$\begin{aligned}
\Lambda_1 &= [\nabla E\{h|\theta^*\}]' W \nabla [E\{h|\theta^*\}] \\
\Lambda_2 &= [\nabla E\{h|\theta^*\}]' W \Omega(h,\theta^*) W [\nabla E\{h|\theta^*\}].
\end{aligned} \tag{52}$$

Borrowing from the literature on generalized method of moments estimation, the optimal weight matrix $W = [\Omega(h,\theta^*)]^{-1}$ results in an SMD estimator with minimal variance. In this case the asymptotic distribution of $\hat{\theta}_T$ simplifies to:

Theorem 4: *Consider the SMD estimator $\hat{\theta}_T$ formed using a weighting matrix W_T equal to the inverse of a consistent estimator of $\Omega(h,\theta^*)$. Then we have:*

$$\sqrt{T}[\hat{\theta}_T - \theta^*] \Longrightarrow N(0,(1+1/S)\Lambda^{-1}) \tag{53}$$

where:

$$\Lambda = \left[\nabla E\{h|\theta^*\}'[\Omega(h,\theta^*)]^{-1}\nabla E\{h|\theta^*\}\right]. \tag{54}$$

The most important point to note about this result is that the penalty to forming an SMD estimator using only a single realization $S = 1$ of the endogenously sampled process $\{\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0)\}$ is fairly small. The variance of the resulting estimator is only twice as large as an estimator that computes the expectation of $h_T(\{u_s\}_{s\leq T},\theta)$ exactly, such as would be done via Monte Carlo integration when $S \to \infty$.

The SMD estimator can be implemented in practice by solving

$$\hat{\theta}_T = \underset{\theta\in\Theta}{argmin}(h_{S,T}(\theta) - h_T)'\left[\hat{\Omega}(h,\theta)\right]^{-1}(h_{S,T}(\theta) - h_T), \tag{55}$$

where

$$\hat{\Omega}(h,\theta) = \frac{1}{T}\sum_{t=1}^{T}\varepsilon_t(\theta)\varepsilon_t(\theta)' \tag{56}$$

where

$$\varepsilon_t(\theta) = h(\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0),\xi_{t-1}(\{u_s\}_{s\leq t-1},\theta,\xi_0)) - h_T(\{u_s\}_{s\leq T},\xi_0,\theta). \tag{57}$$

Thus, an estimate of the optimal weighting matrix $\Omega(h,\theta)$ is recomputed each time the parameter θ is updated.

4 Model Solution and Estimation

We now consider a special case of the model in which there are no additional state variables, x. In this case, the (S,s) bands are only functions of the current wholesale price. We first present Monte Carlo evidence on the small sample properties of the SMD estimator. Second, we estimate the model using actual data for two different products. Finally, we decompose the firm's profits by product into four components. We use this decomposition to infer the share of the firm's profits that are due to markups paid by retail customers and the share due to price speculation. We also use this decomposition to compare John's purchasing decisions to the model's trading rules.

4.1 A special case of the model

Consider a version of the model in which John solves the following problem:

$$\max_{\{q_t^o\}} E \sum_{t=0}^{\infty} \rho^t \left\{ p_t^r q_t^s - c^o(q_t^o, p_t) - c^h(q_t + q_t^o, p_t) \right\} \tag{58}$$

subject to (3), (4), and (5), and where $c^o(q_t^o, p_t)$ is given by (9) and $c^h(q_t + q_t^o, p_t) = \phi\sqrt{q_t + q_t^o}$. As before, John takes the wholesale price p_t and quantity demanded q_t^r as given. He knows p_t before deciding q_t^o. He then draws q_t^r. The holding cost function is a square root so the marginal convenience yield is decreasing in the level of inventories.

We assume the wholesale price evolves according to a truncated lognormal $AR(1)$ process:

$$log(p_{t+1}) = \mu_p + \lambda_p log(p_t) + w_t^p \tag{59}$$

where w_t^p is an IID $N(0, \sigma_p^2)$ sequence. If we let $\overline{\mu}_p$ and $\overline{\sigma}_p$ denote the uncensored mean and standard deviation of the wholesale price distribution, we can compute

$$\tilde{\sigma}_p = \sqrt{log(\overline{\sigma}_p^2 + \overline{\mu}_p^2) - 2log(\overline{\mu}_p)}. \tag{60}$$

Using $\tilde{\sigma}_p$ we can compute μ_p and σ_p by:

$$\mu_p = (1-\lambda_p)(log(\overline{\mu}_p) - \tilde{\sigma}_p^2/2) \text{ and } \sigma_p = \tilde{\sigma}_p \times \sqrt{1-\lambda_p^2}. \tag{61}$$

We assume John sets the retail price using a fixed linear markup rule over the wholesale price:

$$p_t^r = \alpha_0 + \alpha_1 p_t. \tag{62}$$

John draws a quantity demanded q_t^r each period from a mixed truncated lognormal distribution conditional on p_t That is, with probability η, $q_t^d = 0$, and with probability $1-\eta$, q_t^d is drawn from a truncated normal distribution with location parameter $\mu_q(p) = \mu_p - \varsigma log(p_t)$. Both ς, the price elasticity of demand, and η are fixed, time-invariant constants.

Let $\overline{\mu}_q$ and $\overline{\sigma}_q$ denote the unconditional mean and standard deviation of the quantity demanded distribution. We can compute

$$\tilde{\mu}_q = log(\overline{\mu}_q) - \tilde{\sigma}_q^2/2 \text{ and } \tilde{\sigma}_q = \sqrt{log(\overline{\sigma}_q^2 + \overline{\mu}_q^2) - 2log(\overline{\mu}_q)}.$$

Then the mean and standard deviation of quantity demanded conditioned on p_t and a sales occurring, μ_q and σ_q, are computed by:

$$\mu_q = \tilde{\mu}_q + \varsigma \times \mu_p/(1-\lambda_p) \text{ and } \sigma_q = \sqrt{\tilde{\sigma}_q^2 - \varsigma^2 \times \tilde{\sigma}_p^2/(1-\lambda_p^2)}.$$

Finally θ denotes the $(L \times 1)$ parameter vector to be estimated: $\theta = \{K, \alpha_0, \alpha_1, \lambda_p, \overline{\mu}_p, \overline{\sigma}_p, \overline{\mu}_q, \overline{\sigma}_q, \varsigma, \phi\}$.

4.2 Computation

The SMD estimation procedure requires us to solve for the optimal investment rule each time we evaluate the criterion for a new parameter vector. We solve the model by the method of parameterized policy iteration (PPI) introduced by Benitez-Silva *et. al.* (2001). The PPI algorithm involves approximating the value function $V(p,q)$ given in equation (12) as a linear combination of N basis functions, $\{\varphi_1(p,q),\varphi_2(p,q),...,\varphi_N(p,q)\}$:

$$V(p,q) \approx \sum_{n=1}^{N} \vartheta_n \varphi_n(p,q). \tag{63}$$

We discretize the state space into M pairs (p,q), and we denote the m^{th} pair by (p_m,q_m). Thus we transform the value function into a system of M linear equations with N unknowns $\{\vartheta_1,\vartheta_2,....,\vartheta_N\}$:

$$\begin{aligned} \sum_{n=1}^{N} \vartheta_n \varphi_n(p_m,q_m) &= \max_{0\le q^o \le \overline{q}-q_m} \Bigg[ES(p_m,q_m) - c^o(q^o,p_m) - c^h(q_m,p_m) + \\ & \rho E\left\{ \sum_{n=1}^{N} \vartheta_n \varphi_n(p', max[0,q_m-q^s+q^o])|p_m,q_m \right\}\Bigg] \text{ for } m=1,...,M. \end{aligned} \tag{64}$$

As the name suggests, PPI employs an iterative strategy to find the N coefficients on the basis functions that solve the system of equations in (64). Given an initial guess of the coefficient vector, ϑ, we solve the two-period problem on the right-hand side of (64) for each discretized pair (p,q). This yields an $(M\times 1)$ vector containing the current estimate of the optimal decision rule $q^o(p,q)$ at each grid point (p,q). Note that although we discretized the state variables, q^o is a continuous variable subject to the frequently binding constraint: $0 \le q_i^o \le \overline{q} - q$.

Using the decision rule vector, we construct two $(M\times N)$ matrices, P and EP, with elements $P_{m,n}$ and $EP_{m,n}$ given by:

$$\begin{aligned} P_{m,n} &= \varphi_n(p_m,q_m) \\ EP_{m,n} &= E\left\{\varphi_n(p',q_m-q^s+q^o(p_m,q_m))|p_m,q_m\right\}. \end{aligned}$$

Define the $(M\times 1)$ vector y with the m^{th} element given by

$$y_m = ES(p_m,q_m) - c^o(q^o(p_m,q_m),p_m) - c^h(q_m,p_m),$$

and let the $(M\times N)$ matrix X be given by $X=(P-\rho EP)$. Then the system of equations (64) can written in matrix form as $y = X\vartheta$. If $M=N$ and X is invertible, the solution for ϑ is simply $\hat{\vartheta} = y/X$. If $M>N$, we form an approximate solution using ordinary least squares estimation, i,e. $\hat{\vartheta} = (X'X)^{-1}X'y$. Using $\hat{\vartheta}$ as our updated coefficient vector, we iterate on this procedure until the coefficient vector converges to a fixed point.

We approximated the value function by a complete set of Chebychev polynomials of degree 3 in p and q (so $N = 10$). We discretized the state space into 225 (p,q) pairs choosing 15 discrete values for p and 15 discrete values for q. The grid points are fixed at the Chebychev zeros, so they are more heavily weighted toward the boundaries of the state space. This parameterization of the value function does not guarantee concavity of the value function; nevertheless, for the problem at hand we found PPI to be relatively accurate, robust, and fast compared to alternative solution methods. See Benitez-Silva *et. al.* (2001) for detailed comparisons of the PPI algorithm with other solution techniques for a variety of different models.

4.3 Finite Sample Performance of SMD: A Monte Carlo Experiment

We have considerable freedom in our choice of moments functions, the h vector, to use in the criterion. Of course, the most efficient moment functions are the score of the likelihood function. Such an estimator could attain the Cramer-Rao efficiency bound, but it would involve a ratio of integrals, and it is not clear that these integrals can be replaced by simulation estimates and still obtain a consistent SMD estimator. If accurate numerical integrals are required, the computational advantage of the SMD estimator is lost and it may be less computationally burdensome to compute a maximum likelihood estimator directly. We instead match the means and histograms (four of the five quintile bins) of the p, p^r, q^o, q^s, and q processes for a total of 25 moment conditions. We set the number of simulations, S, to 100.

Computing histogram bins requires the use of indicator functions; but indicator functions would create discontinuities into the criterion function, so we used logistic transforms of the indicator functions, approximating $I\{x \leq y\}$ by the logistic function $\exp\{(x-y)/\sigma)\}/(1+\exp\{(x-y)/\sigma\})$ for a small positive number σ. This logistic transform help smooth the criterion function. However, in our simulations we did not allow for unobserved *IID* components ε_t to the wholesale order price p_t as described in section 2. Without the smoothing provided by the ε's, the estimation criterion is no longer guaranteed to be continuously differentiable. The reason is that even though the $s(p)$ function is a continuously differentiable function of θ, small shifts in $s(p)$ can have a discontinuous impact on simulated orders. For example a small change in θ that shifts a given point (p,q) from being above the $s(p)$ band to below it would result in a discontinuous shift in simulated purchases from 0 to $S(p)-q$. In fact, we did find regions in the parameter space in which concentrated "slices" of the criterion function had "steps" and "cliffs." Nevertheless, as one can see from figure 5, there are relatively few points that are near $s(p)$ at low prices where the gap between $S(p)$ and $s(p)$ is large. Most simulated points are close to $s(p)$ only at high prices where $S(p)$ is very close to $s(p)$ and thus the potential discontinuity caused by shifts in $s(p)$ is small. With the additional help from the averaging that occurs in formulating the simulated moments, we observed that the estimation criterion appeared to be smooth for most parameter values. To guard against possible discontinuities or local minima, we employed MATLAB's constrained minimization routine `fmincon.m`, and we visually inspected concentrated slices of the criterion function after each estimation. However we acknowledge that even though plots of the objective appear to be smooth, there may be "microscopic" discontinuities in the slopes in the criterion that may be responsible for the unusually small estimated standard errors that we discuss below.

As presented in equations (56) and (57) of the previous section, the inverse of the optimal weighting matrix, $\hat{\Omega}(h,\theta)$ is the variance-covariance of the residuals from the simulation sequence. However if the model is correctly specified, then when $\theta = \theta^*$, the simulated sequence will have the probability distribution as the observed sequence; therefore we use inverse of the variance-covariance matrix of the residuals of the observed sequence as our weighting matrix, W, where the residuals are given by $\varepsilon_t = h(\xi_t, \xi_{t-1}) - h_T$ where h_t is the sample mean given in formula (27). Since this weighting matrix is just a function of the sample moments, it remains fixed throughout the estimation.

parameter	truth	mean point estimate	std point estimates	mean standard error
K	50.0	53.6	17.4	11.5
α_0	1.00	0.97	0.25	0.41
α_1	1.10	1.10	0.04	0.02
λ_p	0.975	0.976	0.003	0.002
$\overline{\mu}_p$	20.00	20.18	1.22	0.61
$\overline{\sigma}_p$	5.00	5.10	0.87	0.32
$\overline{\mu}_q$	250.0	250.1	28.5	12.5
$\overline{\sigma}_q$	500.0	504.2	94.4	49.2
ς	1.50	1.68	0.54	0.30
ϕ	-4.00	-4.95	2.30	1.12
η	0.33	0.35	0.02	
r	0.075/365	0.075/365		

Table 1: Estimation results on 125 simulated datasets generated by the model.
Two parameters were fixed prior to estimation. The daily interest rate, r, was set to 0.075/365, and the fraction of days in which quantity demanded is zero, η, was set to $1-(\sum I(q_t^s > 0))/T$.

In our Monte Carlo exercise, there are two sets of simulations: first, we fixed the parameter values in the model to those in second column of table 1; we solved the model and created 125 simulated data sets of 1500 periods from the model; second using 125 simulated data sets, we estimated the model using our SMD estimator. The mean and standard deviation of the point estimates as well as the mean of the asymptotic standard errors for each of the ten parameters are reported in table 1. Prior to estimation, we set the interest rate equal to its' true value and η equal to the fraction of days in which no sale occurred.

The quantity data are in hundred-weight (i.e. in 100's of pounds) so the price parameters are in dollars-per-hundredweight (or cents per pound). The fixed cost, K, is set to \$50 per order. The parameter choices for $\overline{\mu}_p$ and $\overline{\sigma}_p$ imply the uncensored price process has a mean of \$17.60 per hundred-weight or 17.6 cents per pound and a standard deviation of \$3.70 dollars per hundred-weight. The parameter values of $\overline{\mu}_q$, $\overline{\sigma}_q$ and ς imply the average sale is 107 hundred-weight or 1,070 pounds. The interest rate r is set to 7.5 percent per annum. The storage cost net of convenience yield, ϕ is set -4.00 dollars per squared-root hundred-weight, so the convenience yield dominates the physical storage cost.

For all ten parameters, the mean point estimates are quite close to their true values. The average point estimates across the 125 simulated datasets are all well within a single standard error of their true values. The SMD estimator recovers the correct parameter values. However, most of the asymptotic standard errors reported in the far right column of table 1 are roughly half the standard deviation of the point estimates reported in the third column. This difference between the standard deviation of the point estimates and the average standard errors can be seen in figure 3. In these ten graphs we plot the kernel density of the ten point estimates along with the density implied from the true parameter values and the estimated standard errors. These plots illustrate that (with the exception of α_0) the density of the point estimates has fatter tails than the density implied by the asymptotic standard errors.

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

$$\frac{T}{(1+1/S)^2}(h_{S,T}(\hat{\theta}) - h_T)' \left[\hat{\Omega}(h)\right]^{-1} (h_{S,T}(\hat{\theta}) - h_T) \rightarrow \chi^2(J-L) \tag{65}$$

In figure 4 we plot a smooth histogram of the J-statistics from each of the 125 simulated datasets. We multiplied each J-statistic by one-eighth. Using a dashed line, we also plot the χ^2 distribution with 15 degrees of freedom. The good news is that the distribution of the J-statistics appear to be distributed χ^2; the bad news is that the J-statistics appear to be about 8 times too big.

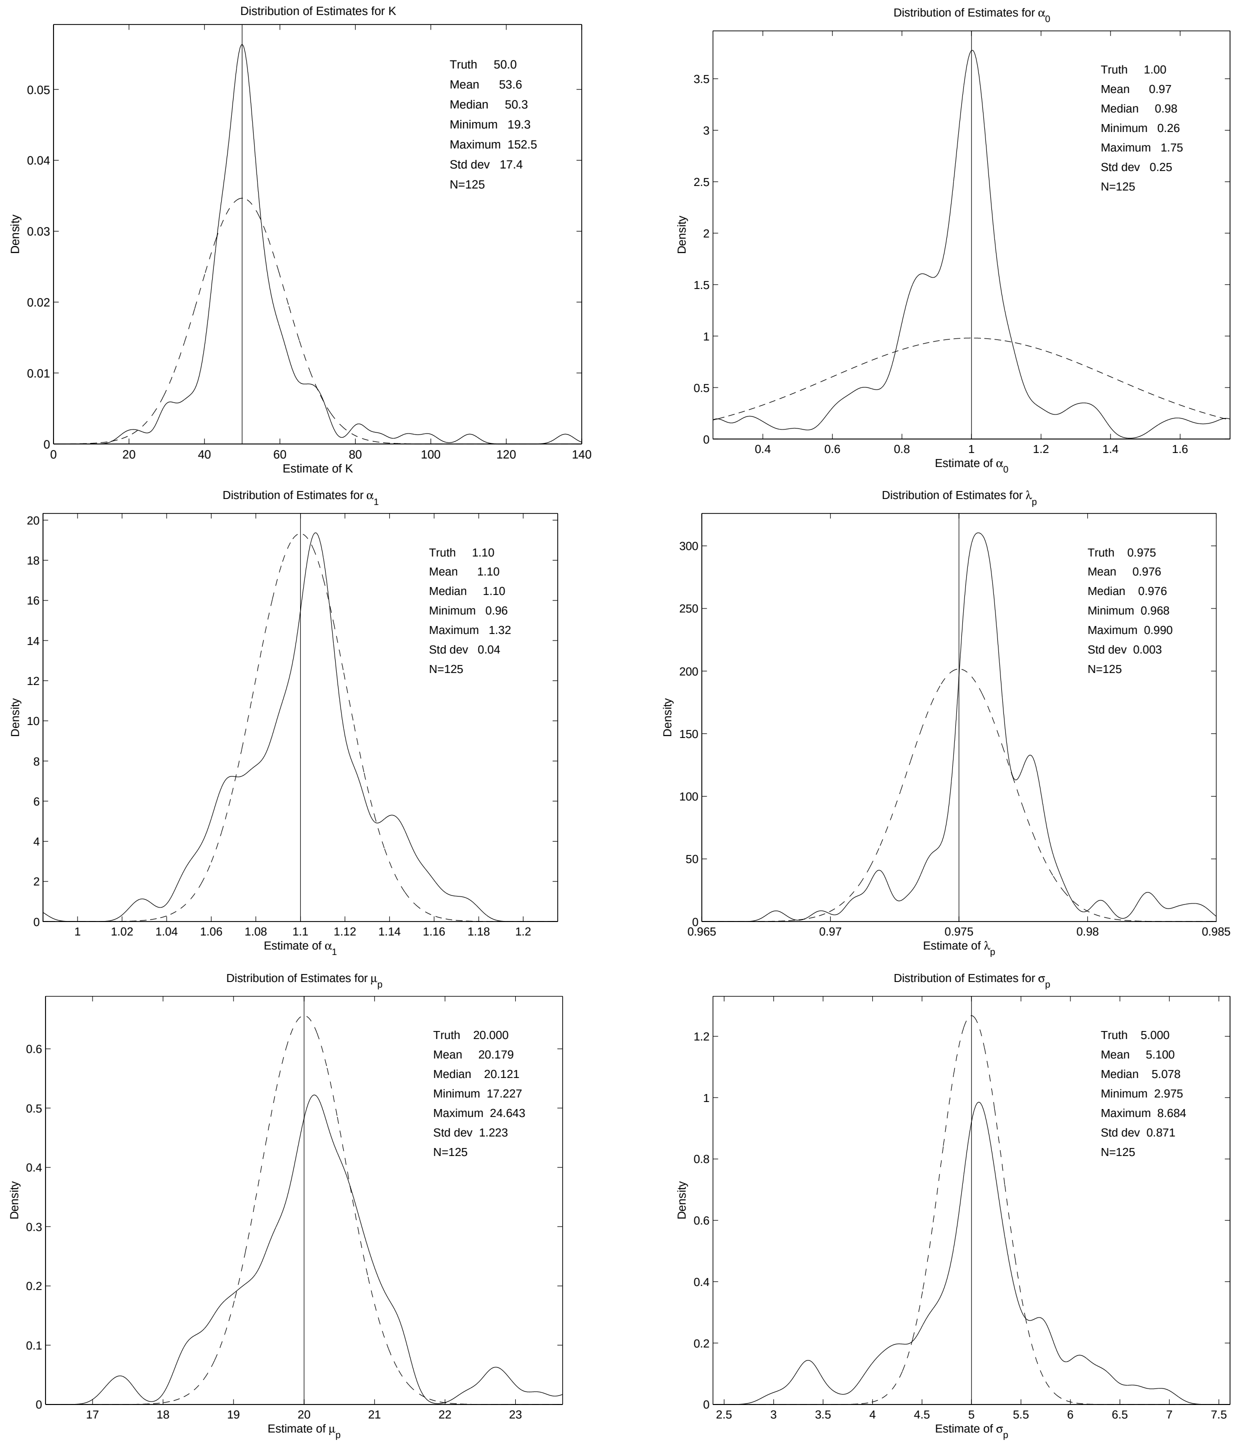

Figure 3: Plot of the kernel density of the simulated results for each parameter (solid line) and the density implied by the true value of each parameter and the mean asymptotic standard error for each parameter (dashed line). The solid vertical line marks the true parameter value.

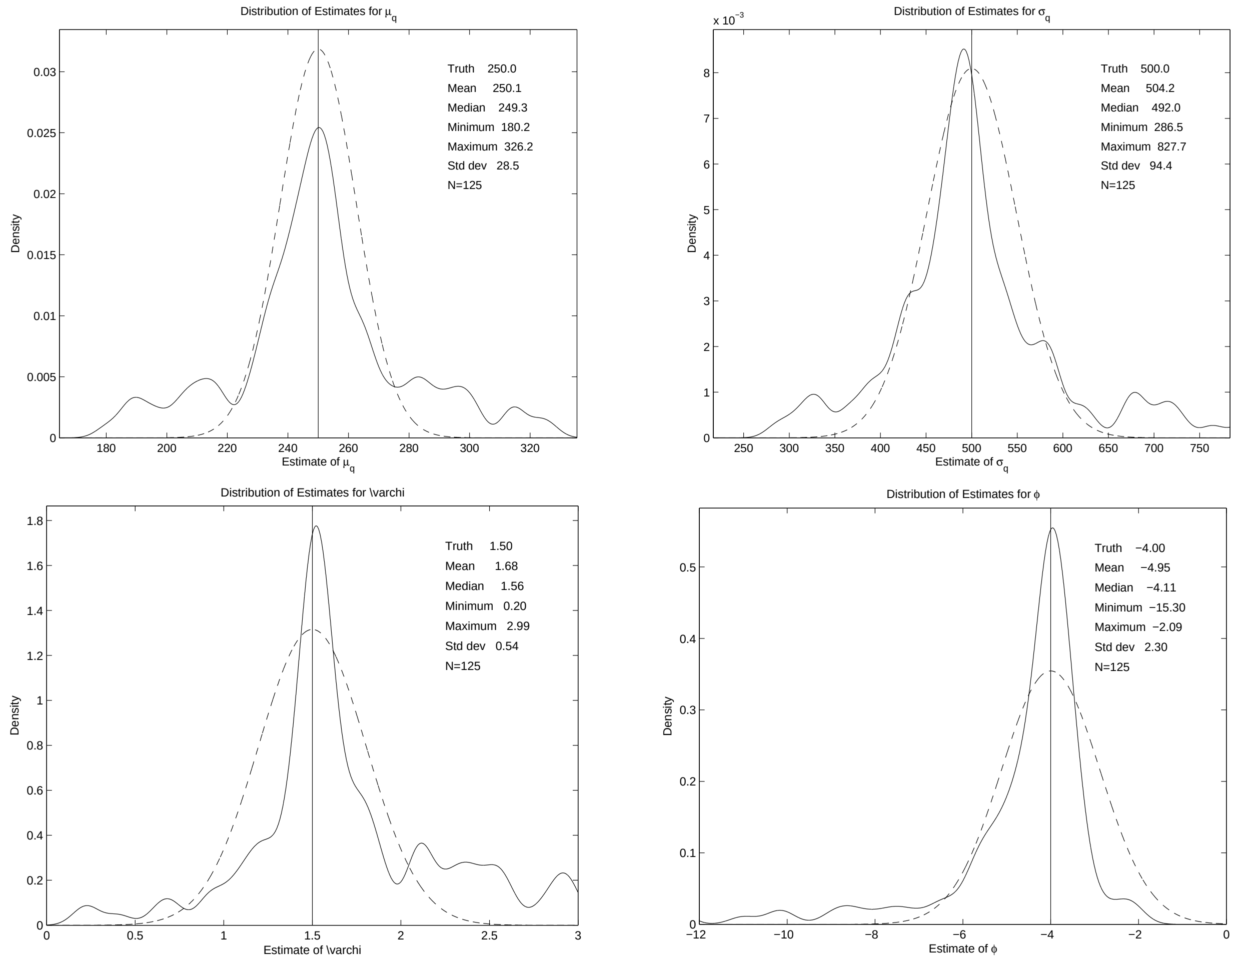

Figure 3: (continued) Plot of the kernel density of the simulated results for each parameter (solid line) and the density implied by the true value of each parameter and the mean asymptotic standard error for each parameter (dashed line). The solid vertical line darks the true parameter value.

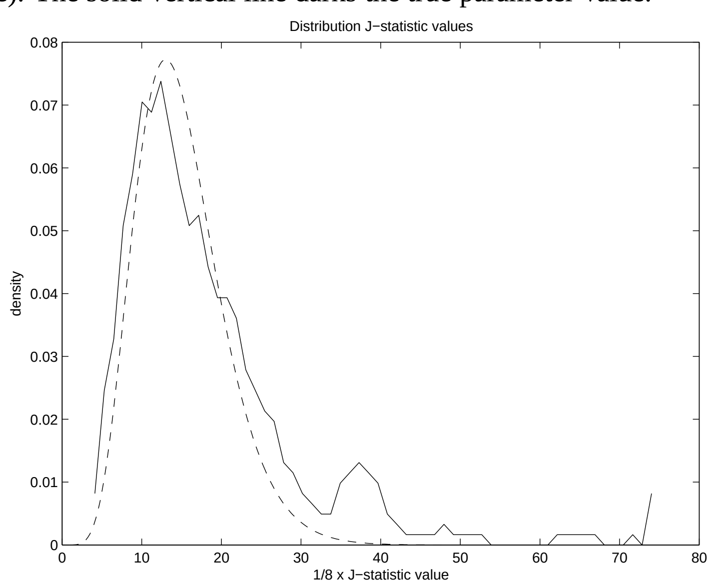

Figure 4: The solid line plots the distribution of J-statistics from the 125 estimations. Each J-statistic is divided by 8. The dashed line plotd the density of a χ^2 distribution with 15 degrees of freedom.

The small asymptotic standard errors and the large χ^2 statistics may be due to small discontinuities in the estimation criterion, a result of our failure to account for unobservable components ε_t of purchase prices p_t.[13] These results suggest that although the consistency of the SMD estimator is not jeopardized by small discontinuities, the estimated covariance matrix and standard errors may be much more sensitive to small discontinuities in the simulated moments. In future work we plan to investigate how discontinuities could affect the asymptotic properties of the SMD estimator, but this investigation is beyond the scope of this paper. In the absence of an asymptotic theory that accounts for discontinuities in the estimation criterion, it may be important to include unobservables such as ε_t in the simulations in order to smooth out these discontinuities in order to obtain consistent estimates of the asymptotic covariance matrix.

4.4 Empirical Results

We now estimate the model for two products independently. In table 2 we report the point estimates and standard errors for the parameters of the model for products we call product 2 and 4. As before, the interest rate r, and η are fixed prior to estimation: r is set to 0.075/365, and η is set to the the fraction of days no sale occurred. John would not provide us specific data on the firm's borrowing and lending (many sales involve trade credit), but told us that one and three-quarter points over a short-term LIBOR rate was a good estimate of the interest rate the firm faced. The average 3-month LIBOR rate over the period studied is about 5.75, which implies an average annual borrowing rate for the firm of about 7.5%.

Although we estimated the parameters for each of these products independently, it is reassuring that several of the point estimates are similar across the two products. It is reasonable to expect that the parameters, K, α_0, α_1, λ_p, ς, and ϕ to be quite similar, if not identical, across products.[14] In general this is case. Prior to estimating the model, we asked John for his estimate of the fixed cost of placing an order (i.e. the parameter K). He stated $50, while our point estimate is about $120. He thought the main fixed cost to ordering is the value of his and his administrative assistant's time it takes to complete the paperwork. Our larger estimate may be picking up some the search costs John incurs prior to making a purchase.

The marginal cost of storage parameter, ϕ, is negative for both products so the marginal convenience yield dominates the physical costs of storage. This result is consistent with the observation in the commodity storage literature that negative storage costs are a key determinate of the autocorrelation in commodity

[13] We initially suspected a bug in our computer code. But a coding error that generates standard errors one half the correct size would have generated J-statistics that are 4 (not 8) times too big.

[14] We could have estimated the model jointly across the two products, constraining these value to be equal across products.

	Product 2		Product 4	
parameter	point estimate	standard error	point estimate	standard error
K	118.7	8.3	123.5	18.6
α_0	1.75	0.11	2.14	0.40
α_1	1.02	0.01	1.00	0.02
λ_p	0.9824	0.0003	0.9819	0.0005
$\overline{\mu}_p$	22.31	0.18	20.89	0.15
$\overline{\sigma}_p$	6.20	0 09	6.00	0.06
$\overline{\mu}_q$	191.7	5.5	218.4	8.7
$\overline{\sigma}_q$	252.6	18.8	265.8	14.9
ς	3.16	0.13	3.34	0.09
ϕ	-3.25	1.14	-3.35	0.18
r	0.075/365		0.075/365	
η	0.35		0.34	
$\chi^2(15)$	1586		948	

Table 2: Estimation Results using data for product 2 and product 4.
Two parameters were fixed prior to estimation. For both products, the daily interest rate, r, was set to 0.075/365; for each product individually, the fraction of days in which quantity demanded is zero, η, was set to $1-(\sum I(q_t^s > 0))/T$.

prices. We experimented with various function forms for the holding cost function and/or a stock-out penalty function. If the marginal value of holding inventories is small when inventories are close to zero (i.e. when the wholesale price is high), the optimal strategy is for the firm to effectively shut down by holding no inventories until the wholesale price falls. In other words, the $s(p)$ band equals zero for p greater than some threshold. While we do observe near-zero levels inventories in the data from time to time, these near stockout levels do not persist for more than a few days. If the marginal value of holding inventories is “too large” even when the firm is holding large levels of inventories, the model implies the firm should (counterfactually) always hold inventories near its capacity constraint. Hence having some convexity in the holding cost helps match the mean and spread of inventories holdings we see in the data.

The endogenous sampling problem is illustrated in figure 5. We plot the $S(p)$ and $s(p)$ bands derived from the model’s optimal decisions rules using the estimated parameter vector for product 4. Due to the fixed costs of ordering, 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 minimum order size is a decreasing function of the price. In figure 5 we also scatterplot a set of simulated state space pairs (p_t, q_t). According to the firm’s optimal trading rule, the firm only makes purchases when the (p_t, q_t) pair is below the $s(p)$ band (in the southwest corner of the graph). In the simulation presented, this occurs less than 16 percent of time.

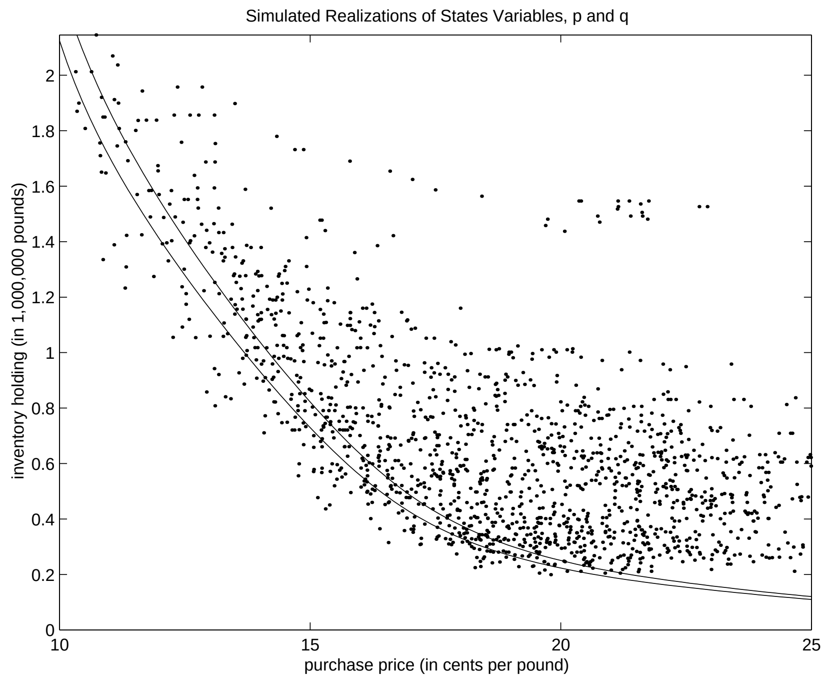

Figure 5: Scatterplot of 1500 purchase price and inventory holding pairs from a simulation for product 4. The solid lines are the $S(p)$ and $s(p)$ bands from the model.

4.5 Goodness of Fit

In table 2 we also report the minimized SMD estimation criterion. Both models are decisively rejected; even if we divide the J-statistics by 8 (as suggested by the Monte Carlo results) the marginal significance level is less than .0001. Figure 6 illustrates the reason for this rejection. In this figure, the downward sloping curves are the (S,s) bands from the model for product 4. We superimpose a scatterplot of the purchase price and inventory level pairs we observe on 212 days John purchased product 4. We then draw vertical lines from each one of these (p,q) pairs to the point $(p,q+q^o)$, so the length of each vertical line represents the size of the purchase q^o. If our model was true, each of the observed (p,q) pairs would be to the southwest of the $s(p)$ band. Further, the top of each of the vertical lines would just touch the $S(p)$ band. Clearly this is not what we see.

If we divide the sample into four periods, we can see evidence of shifting (S,s) bands. For the initial period from July 1, 1997 to March 27, 1998, we mark the (p,q) pairs in figure 6 with circles. During this period prices were around 20 cents per pound, and the large purchase on March 27, 1998 at 18.5 cents per pound discussed in section 2.1 looks consistent with sharply downward-sloping (S,s) bands. In the second period, from March 28, 1998 to December 31, 1998, we denote the (p,q) pairs with diamonds. Prices during this period return range from 19 to 22 cents, but the (S,s) bands have shifted upward. Why the shift? In figure 7 we plot a thirty-day moving average of sales for product 4. During this second period, sales were unusually high, averaging between 40,000 and 50,000 pounds per day. It appears there was an outward shift in the demand schedule for steel during this period that our model completely misses. But then with the large increase in imported steel in 1999, the demand schedule appears to shift in dramatically. During this third period which we date from January 1, 1999 to October 8, 1999 and mark with squares, average daily sales fell to 10,000 pounds per day – just 1/4 of what they were the previous year. It appear the (S,s) bands shifted in as John reduced his inventory in response to this demand shock despite falling prices. The fourth regime from October 9, 1999 to the end of the sample is denoted with stars. During this period, sales are relatively stable and larger purchases are associated with lower prices.

The main shortcoming of the estimation is our inability to match the downward trend of the price process that we see in almost all of the firm's products. As illustrated in figure 2 the wholesale price for product 4 fell from 20 cents per pound in 1997 to about 12 cents per pound in 2002. No such trend is evident in simulations such as the one presented in figures 8 and 9. In our model, prices are stationary though highly persistent. Consequently, as can been seen in the (S,s) bands plotted in figures 5 and 6 the optimal decision rules imply counterfactually that the firm should make only small purchases and hold low levels of inventories whenever the wholesale price is above 16 cents per pound. From figures 1 and 2 we see that, for product 4, the firm made large purchases around 18.5 cents per pound in April 1998, and around 10.5 cents per pound in December 2000.

An often suggested solution to this trend problem is that we assume that the log of steel prices follow a random walk. For product 4, if we concentrate out all the other parameters except λ_p, the criterion surface is a steeply sloped and smooth cup centered around 0.982 so the small standard error associated with the AR(1) coefficient is not surprising. But the concentrated criterion surface actually turns down slightly between .995 and 1.01. (The model still solves numerically for values of λ_p slightly greater than one.) The global minimum is still located at 0.982, but there appears to be a local minimum just above 1.00. However if we assume the log price process follows a (truncated) random walk, the optimal decision rules implies frequent small- to medium-size orders such that the inventory level fluctuates closely around a fixed target level. A version of the model which assumes $\log(p_t)$ follows a random walk will not imply the large variation in inventory holdings that we see in the data. A second potential solution is to detrend the data. However when we first started working on this project, no one we talked to expected steel prices

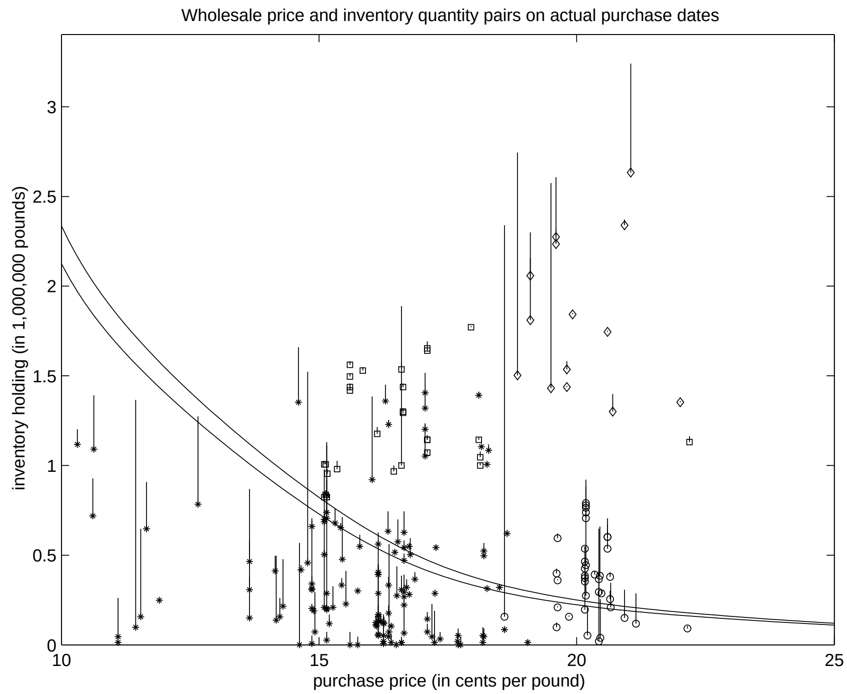

Figure 6: Scatterplot of purchase price and inventory holding pairs on the 212 actual purchase dates for product 4. The length of each vertical line measures the size of each purchase. The downward-sloping curves are the $S(p)$ and $s(p)$ bands from the model.

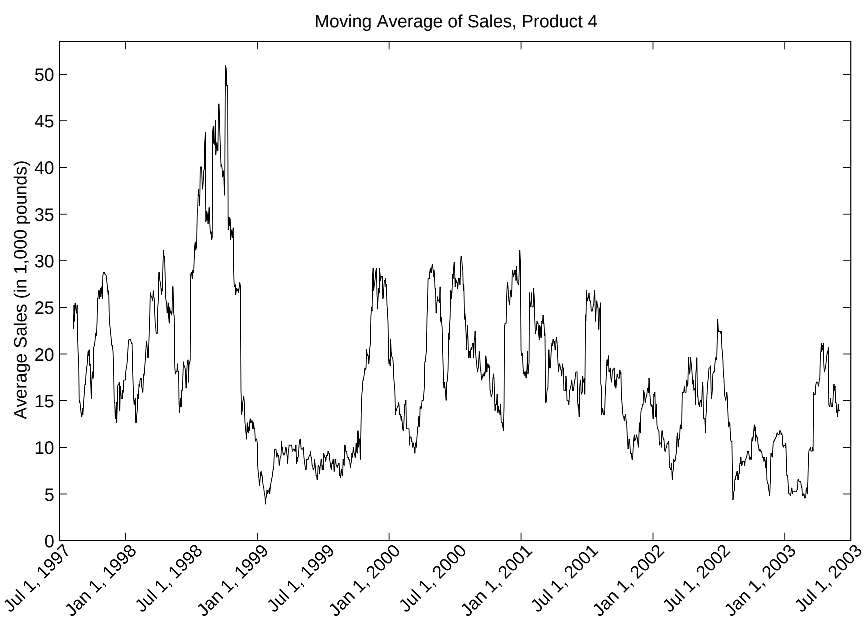

Figure 7: Thirty-day moving average of sales for product 4.

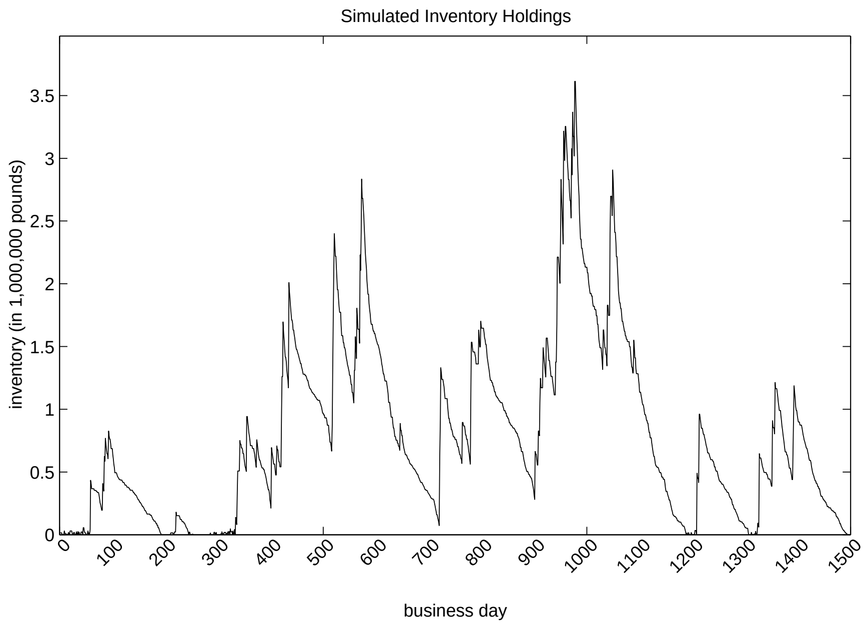

Figure 8: Simulated inventory data from the estimated model for product 4.

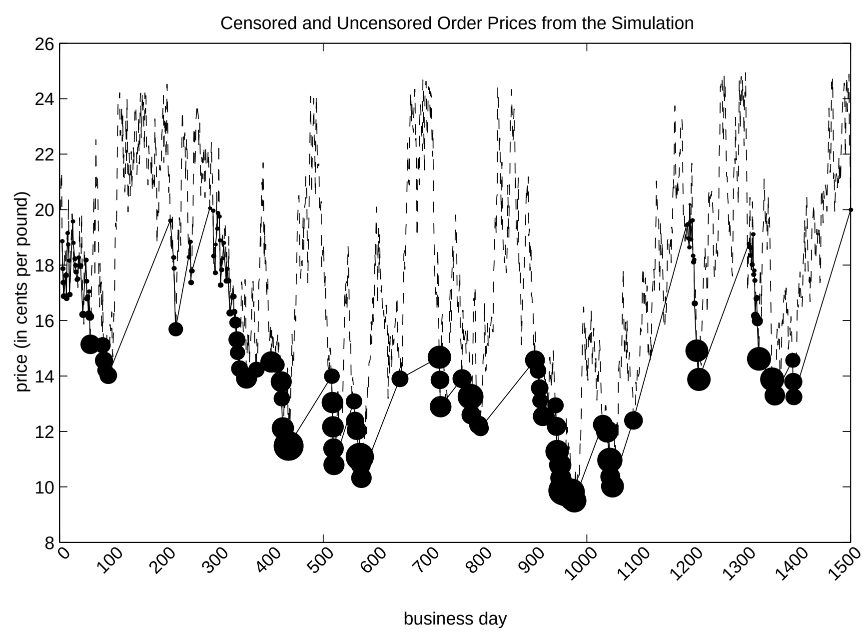

Figure 9: Censored (solid line) and uncensored (dotted line) purchase prices, p_t from simulation for product 4. For the censored series, the size of the marker is proportional to the size of the purchase.

to decline 40% in the first four years of the sample. To some extent we are just working with too short a sample period. A third candidate solution is to add an additional macroeconomic state variable. Such a variable could allow for ‘high demand”, “medium demand” and “low demand” regimes. As we discuss below, we view this third solution as the most promising.

However even with a macro shock, the model would still need a transitory shock such an unobservable component to the purchase price to fully reconcile the model to the data. It is clear from figure 6 that even within each of the four regimes, the (S,s) bands are not stable. Augmenting the model with a regime shock and a transitory unobserved price shock is left for future work.

While figure 6 highlights some limits of the models, the models at the estimated parameter values capture several of the salient features of the inventory and price data. Figures 5, 8, and 9 highlight some of the strengths of the model. First, in the data purchases are made infrequently. Figure 9 presents the censored and uncensored purchase price series, p_t. The solid line is the analog of what we observe in the data: we linearly interpolated between the prices at which transactions took place; the dotted line includes the unobserved prices at which no transactions occurred. During periods of low prices (e.g. days 400-600 and 900-1100) the firm aggressively made purchases to build up large levels of inventories. The large levels of inventories were slowly drawn down as prices inevitably rose. After exploiting a low price opportunity, the firm may subsequently make no new purchases for many days. Second, we observe both small and large purchases in the data. Again this can been in all three graphs. In figure 5 when the (p_t, q_t) pair (dot) is below the $s(p)$ band, the size of the order is the vertical distance between the $S(p)$ band and the (p_t, q_t) pair (dot). When the purchase price is less than 15 cent per pound, we observe both large and small orders. When the purchase price is above 18 cents per pound we only observe small orders. In figure 9, the size of the marker is proportional to the size of the purchase. Again once can see that the model predicts relatively large purchases when the price is low and relatively small purchases when the price is high. Third, in the data we observe periods of with high levels of inventories and periods with low levels of inventories. From the scatterplot in figure 5 and the time path of inventories plotted in figure 8 we can see that the model predicts that inventory levels will vary over the sample between almost zero and 3.5 million pounds. Compare this to figure 1.

Finally, we use simulations of the estimated model to deduce the relative importance of capital gains versus markups for the overall profitability of the firm. This also provides another diagnostic to illustrate where the model succeeds and fails in describing John’s investment behavior. By substituting the law of motion for inventories (5) into the firm’s objective function, (58), the discounted present value of the firm’s profits can be expressed as:

$$\sum_{t=0}^{T}\rho^t\pi(p_t,p_t^r,q_t^r,q_t+q_t^o) = \sum_{t=0}^{T}\rho^t(p_t^r-p_t)q_t^s+q_0p_0+\sum_{t=1}^{T}\rho^t(p_t-(1+r)p_{t-1})q_t - \sum_{t=0}^{T}\rho^t I(q_t^o)K-\sum_{t=0}^{T}\rho^t c^h(q_t+q_t^o,p_t). \quad (66)$$

The first term on the right hand side of equation (66) can be interpreted as the discounted present value of the markup paid by the firm's retail customers over the current wholesale price while the third term can be interpreted as the discounted present value of the capital gains or loss from holding the steel from period $t-1$ into period t. The fourth, and fifth terms are the discounted present values of the order costs and the holding costs incurred by the firm over the sample period.

Since this decomposition depends on the wholesale price path between purchases, we simulate between purchase dates via importance sampling. That is, for each interval between successive purchase dates, we simulate wholesale price paths that are consistent with the estimated law of motion (59) and the observed purchase prices at the beginning and end of the interval. Since our theory implies that the firm places an order anytime the quantity falls below the order threshold, $s(p)$, we truncate the simulated price process by rejecting any paths such that $q_t < s(p_t)$ for any draw within the simulated paths. We discuss our simulation method in more detail in the appendix.

We first employ this decomposition to evaluate John's actual performance over the nearly-six year sample period for products 2 and 4. For a given interpolated price series, we decomposed the firm's profits using the actual data for q_t, q_t^s, and q_t^o, our fixed value for the interest rate, r, and our point estimates for K, and ϕ. In table 3 we report the average decomposition from 100 simulated wholesale price paths. As discussed in the introduction, the price of steel fell dramatically in the second and third years of the sample. Nevertheless, by our accounting, the firm made \$430,000 (product 2) and almost \$500,000 (product 4) from the markup and capital gains on each of these two product over the nearly-six year period.[15] Ignoring the fixed order cost and the returns from the convenience yield, about 65 percent (product 2) and 71 percent (product 4) of these profits came from the markup, while the remaining 35 and 29 percent came from capital gains. We find it remarkable and evidence of John's acumen in steel trading that the firm made positive capital gains over this period despite the price of steel falling about 40 percent. While John's success in price speculating is good for the firm's profits, it increases the potential biases from failing to account for the endogeneity of the sampling process.

[15] Profits are discounted back to July 1, 1997.

	Product 2						Product 4					
	G.M.'s actual Performance			Model's Policy Prescription			G.M.'s actual Performance			Model's Policy Prescription		
markup	$281,499	(34,226)	65%	$211,543	(27.761)	41%	$352,997	(35,798)	71%	$267,730	(29,881)	41%
capital gain	148,865	(36,360)	35%	306,661	(30,732)	59%	144,440	(38,248)	29%	455,410	(46,683)	59%
holding cost	299,200	(0)		285,266	(13,746)		354,847	(0)		331,108	(28,898)	
order costs	-21,695	(0)		-15,754	(1,480)		-22,829	(0)		-31,668	(1,728)	
total profits	707,869	(4,364)		787,715	(55,167)		829,455	(3,825)		1,022,581	(85,810)	

Table 3: Profit Decomposition For Product 2 and 4 Using Equation (66)

Both the actual and the counter-factual profits cover the 1500 days studied and are discounted back to the start of the sample period, July 1, 1997. The profit numbers reported are the average across 100 simulations. The numbers in parentheses are the standard deviations from the 100 simulations. Total profits are the sum of the first four rows.

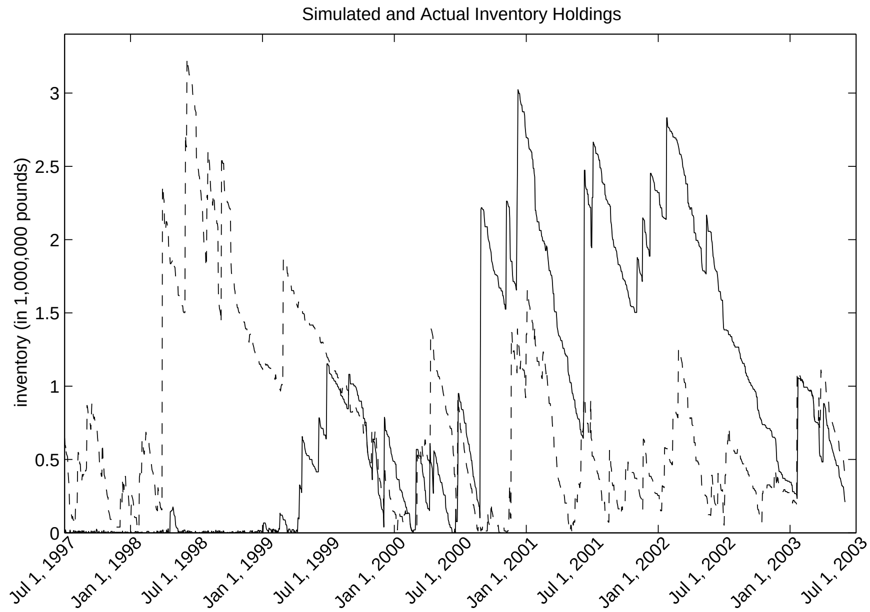

Figure 10: Actual (dashed line) and counter-factual (solid line) inventory holdings for product 4.

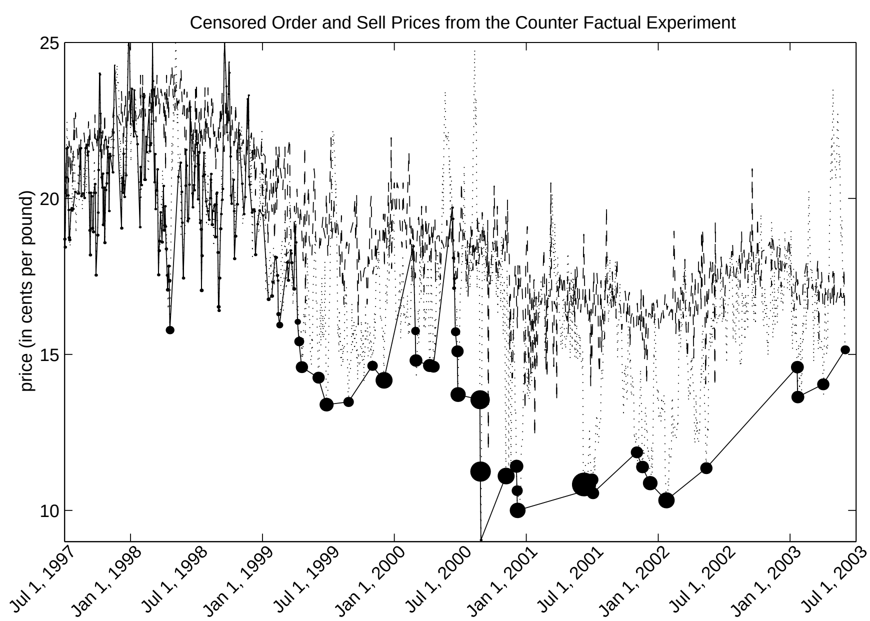

Figure 11: Counter-factual uncensored purchase prices (dotted line), censored purchase prices (solid line), and retail prices (dashed line) for product 4. For the censored purchase price series, the size of the marker is proportional to the size of the purchase.

As a diagnostic of our model, we compare John's performance to the model's predictions. In this exercise we take as given the 100 interpolated wholesale price series, the firm's quantity demanded series, and the firm's initial level of inventories for each product. But in this case, we let the model's optimal decision rule dictate when and how much to order.[16] Inventories follow the accumulation identity given by equation (5). As reported in table 3, had John counter-factually followed the optimal order strategy implied by our model, his discounted profits from the markup would have been roughly $70,000 less for product 2, and $85,000 less for product 4. However, his capital gains would have been considerably larger: $160,000 more for product 2; $300,000 more for product 4.

This counter-factual provides another measure of fit for the model and guidance for where to improve the model in the future. The model suggests that John should price speculate much more aggressively than we observe. In figures 10 and 11 we plot the prices and inventory holdings for one simulation of the model. In figure 10 we plot both the actual inventory holdings along with the implied holdings under the model's decision rules. In figure 11 we plot the corresponding retail and wholesale price paths. The model's counter-factual inventory path differs considerably from the firm's actual inventory path. In the beginning of the sample, years 1997 and 1998, when prices were high, the model implies John should have made frequent small purchases and held relatively low levels of inventories. As was discussed in the introduction, in April 1998 when the wholesale price of steel dropped from 20 cents per pound to 18.5 cents per pound, John built up his inventory of product 4 substantially. In contrast the model does not view 18.5 cents as a particularly good price; as can be seen in the (S, s) bands plotted in figure 5, the target inventory level at 18.5 cents is around 300,000 pounds. In April 1998, John's inventory of product 4 exceeded 2,000,000 pounds. It is not until July 1999 when simulated prices fell below 14 cents a pound that the model recommends holding more than 1,000,000 pounds of inventory.

During the second half of 1999 both the model and John let inventories fall to almost zero. However by the second half of 2000 the counter-factual inventory policy diverges again from John's observed behavior. In August 2000 the model recommends making a set of large purchases to build inventory levels to above 2,000,000 pounds while John kept inventory levels near zero. In general, during the second half of the sample period, John held relatively low levels of inventories, whereas the model's inventory was often in excess of 2,000,000 pounds. Basically, the model recommends that John's purchasing strategy should have

[16] We placed one ad hoc restriction on our decision rule. In mid-December 2000, John had an opportunity to buy a limited quantity of products 2 and 4 for a little over 10 cents per pound. John bought as much as he could at these prices. Our model dictated that he should have purchased large quantities at these prices. For the counter-factual experiment we constrained the model purchase no more steel than we actually observe on these dates.

been the opposite of what he did: John should have held low inventory levels in 1997, 1998 and 1999, and high inventory levels in 2001 and 2002 .

This comparison between the model and the data is “rigged” in the model’s favor in one dimension and “rigged” against the model in another. Since we used the entire sample period to both estimate the model and evaluate the model’s performance, the model “knows” the mean and the standard deviations of prices and quantity demanded for the entire period. The model knows, whereas John did not know, that a price of 18.5 cents per pound in the Spring of 1998 was an above-average price for the 1997-2003 period. In this way the model has an advantage over the manager. However the model is constrained to sell at most the quantity of steel John actually sold. The model does not get the opportunity make any sales John might have had the option to make but decided to turn down.

While we do not report an out-of-sample comparison between our model and John, if we had estimated the model through the Fall of 2001, and then used our model to dictate purchases for the firm for the Winter and Spring of 2002, our model would still have outperformed John. In the Fall of 2001, the firm was purchasing steel around 10 to 12 cents per pound. We told John at that time that our model recommended building up inventories at these prices. He did not follow this advice since he anticipated further price declines. He argued (and to be honest, we did not disagree) that our model did not take into account the potential slowdown in the economy in the wake of the terrorist attack of September 11, 2001 that he expected to reduce demand for steel. He also expected new production capacity from the Nucor Corporation to put additional downward pressure on prices. However, with the bankruptcy of Bethlehem Steel in October 2001 as well as both the anticipation of an increase and the actual increase in steel tariffs imposed by President Bush in March 2002, steel price increased about 20 percent in the Spring of 2002 to the 12 to 14 cent range. In the Spring of 2002, we reminded John that in the fall our model recommended he build up inventories. He sighed, “I wish I had.”

In this case, our model “got it right” but perhaps not for the right reasons. Our model was predicting an increase in prices since our model always expects prices to return to the sample mean. Our model does not use information on where the economy is going as a covariate for predicting steel prices or steel sales. For example, there is no way for our current model to update expectations of steel prices in response to news, such as President Bush’s decision to impose steel tariffs. To obtain a more realistic model that might be able to rationalize John’s apparently more cautious speculative strategy, we would need to add macroeconomic state variables x. Then we could use our model jointly with a macroeconomic forecasting model to provide conditional inventory level recommendations to the firm such as “If you expect the economy to remain strong, the model recommends holding inventories in a range from X to Y; if you expect the economy to weaken, ...” These additional state variables would enable us to capture apparently non-stationary features of steel prices (such as persistently increasing or decreasing price trajectories over relatively long periods of time), and would serve as additional covariates that shifting the (S, s) bands up and down in response to news of persistent macroeconomic shocks, helping the model to better fit the observed purchase and inventory data.

5 Conclusion

In this paper we estimate a dynamic programming model of a single steel trader. Since we only observe prices of days transactions occurred, we have an irregularly spaced, non-random sample of the price process. To address this sampling issue, we present and employ a simulated minimum distance estimator for estimating an endogenously-sampled Markov process. This SMD estimator is computationally simpler than maximum likelihood, but it requires solving the dynamic programming problem at each trial value of the unknown parameter vector for the endogenous sampling rule. Using this sampling rule, the SMD estimator is able to consistently estimate the unknown parameters of the Markov process even though the econometrician has incomplete information on the process.

While this research was motivated by a new dataset from a single steel intermediary, most datasets in which agents have the choice of whether and when to participate in a market activity will be endogenously sampled. In most markets, the only prices recorded are the transaction prices – econometricians almost never get to observe prices offered but not transacted on. For example, econometricians rarely get to observe the wages unemployed job seekers are offered but refuse. It should be straightforward to apply the SMD estimator to other types of endogenous sampling problems that arise in time series contexts.

References

[1] Aït-Sahalia, Y. and P. A. Mykland (2003) "The Effects of Random and Discrete Sampling When Estimating Continuous-Time Diffusions" *Econometrica* 71, 483-549.

[2] Aguirregabiria, V. (1999) "The Dynamics of Markups and Inventories in Retailing Firms" *Review of Economic Studies* 66, 275–308.

[3] Allenby, G., McCulloch, R. and P. Rossi (1996) "On the Value of Household Information in Target Marketing" *Marketing Science* 15, 321–340.

[4] Arrow, K.J., Harris, T. and J. Marschak (1951) "Optimal Inventory Policy" *Econometrica* 19-3, 250–272.

[5] Athreya, R. (2002) "Price Dispersion in the Wholesale Market for Steel" manuscript, Yale University.

[6] Benitez-Silva, H., G. Hall, G. Hitsch, G. Pauletto and J. Rust (2001) "A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems" manuscript, State University of New York at Stony Brook.

[7] Boizot, C., J.M. Robin, and M. Visser (2001) "The Demand for Food Products: An Analysis of Interpurchase Times and Purchased Quantities" *Economic Journal* 111, 391-419.

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

[9] Chan, H., G. Hall and J. Rust (2003) "An Empirical Model of Optimal Intertemporal Price Discrimination" manuscript, Yale University.

[10] Deaton, A. and G. Laroque (1992) "On the Behavior of Commodity Prices" *Review of Economic Studies* 59, 1–23.

[11] Duffie, D. and K.J. Singleton (1993) "Simulated Moments Estimation of Markov Models of Asset Prices" *Econometrica* 61-4, 929–952.

[12] Engle, R. and J. Russell (1999) "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Data" forthcoming, *Econometrica.*

[13] Erdem, T. and M. Keane (1996) "Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets" *Marketing Science* 15–1, 1–20.

[14] Fabian, T., J.L. Fisher, M.W. Sasieni, and A. Yardeni (1959) "Purchasing Raw Material on a Fluctuating Market" *Operations Research* 7, 107-122.

[15] Futia, C. (1982) "Invariant Distributions and the Limiting Behavior of Markovian Economic Models" *Econometrica* 50-1, 377-408.

[16] Golabi, K. (1985) "Optimal Inventory Policies when Ordering Prices are Random" *Operations Research* 33–3, 575-588.

[17] Hall, A. and A. Inoue (2002) "The Large Sample Behavior of the Generalized Method of Moments Estimator in Misspecified Models" manuscript NC State University.

[18] Hall, G. and J. Rust (1999) "An Empirical Model of Inventory Investment by Durable Commodity Intermediaries" *Carnegie-Rochester Conference Series on Public Policy*, 52.

[19] Hall, G. and J. Rust (2001) "The (S,s) Rule is an Optimal Trading Strategy in a Class of Commodity Price Speculation Problems" manuscript, Yale University.

[20] Hall, G. and J. Rust (2002) "Econometric Methods for Endogenously Sample Time Series: The Case of Commodity Price Speculation" NBER Technical Working Paper 278.

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

[22] Heckman, J.(1981) "Statistical Models for Discrete Panel Data" in C. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data* MIT Press.

[23] Kaldor, N. (1939) "Speculation and Economic Stability" *Review of Economic Studies* 7, 1-27.

[24] Kalymon, B. (1971) "Stochastic Prices in a Single-Item Inventory Purchasing Problem" *Operation Research* 19, 1434-1458.

[25] Kingman, B.G. (1969) "Commodity Purchasing" *Operational Research Quarterly* 20-1, 59-79.

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

[27] Manski, C. and D. McFadden (1981) "Alternative Estimators and Sample Designs for Discrete Choice Analysis" in C. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data* MIT Press.

[28] McFadden, D. (1997) "On the Analysis of Endogenously Recruited Panels" manuscript, University of California at Berkeley.

[29] 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*

[30] Moinzadeh, K. (1997) "Replenishment and Stocking Policies for Inventory Systems with Random Deal Offerings" *Management Science* 43–3, 334-342.

[31] Pakes, A. and D. Pollard (1989) "Simulation and Asymptotics of Optimization Estimators" *Econometrica* 57-5, 1027-1057.

[32] Russell, J. and R.F. Engle (1998) "Econometric Analysis of Discrete-Valued Irregularly-Spaced Financial Transactions Data Using a New Autoregressive Conditional Multinomial Model" UCSD Discussion Paper 98-10 `ftp://weber.ucsd.edu/pub/econlib/dpapers/ucsd9810.pdf`

[33] Rust, J. and G. Hall (2003) "Middlemen versus Market Makers: A Theory of Competitive Exchange" *Journal of Political Economy* 111-2, 353-403.

[34] Scarf, H. (1959) "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.

[35] Song, J. and P. Zipkin (1993) "Inventory Control in a Fluctuating Demand Environment" *Operations Research* 41-2, 351-370.

[36] White, H. (1982) "Maximum Likelihood Estimation of Misspecified Models" *Econometrica* 50–1, 1–26.

[37] Williams, J.C. and B. Wright (1991) *Storage and Commodity Markets* Cambridge University Press, New York.

Appendix: Simulating Price Paths with Fixed Starting and Ending Points

We thank Michael Keane for suggesting and explaining this procedure to us.

We assume that wholesale prices follow the AR(1) process given in equation (59) of the paper. To simplify the presentation in this appendix, let p_t denote the $log(p_t)$. Assume we observe p_{t_1} and p_{t_2} on dates t_1 and t_2, but we do not observe any prices on dates in between. We want to simulate realizations of $\{p_{t_1+1}, p_{t_1+2}, ..., p_{t_2-1}\}$ that are consistent with both p_{t_1} and p_{t_2} and the law of motion (59). Let $\tau = t_2 - t_1$, be the recurrence time.

We write the price system using state-space notation using a nonstandard ordering of the state vector:

$$\begin{bmatrix} 1 \\ p_{t_1} \\ p_{t_2} \\ p_{t_1+1} \\ p_{t_1+2} \\ \vdots \\ p_{t_2-1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & \dots & 0 & 0 \\ \mu_p & 0 & \lambda_p & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & & \ddots & & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ p_{t_1-1} \\ p_{t_2-1} \\ p_{t_1} \\ p_{t_1+1} \\ \vdots \\ p_{t_2-2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \sigma_p \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} w^p_{t_2}. \tag{67}$$

We rewrite this equation using more compact notation as:

$$\mathbf{p}' = A\mathbf{p} + Cw^{p'} \tag{68}$$

where the **p** denotes the vector of logged prices and the prime denotes the next period's values.

We then compute the variance-covariance matrix of the price vector:

$$\Omega = \sum_{j=0}^{\tau+1} A^j CC' A'^j.$$

We then compute the Cholesky decomposition of the (2:τ+2,2:τ+2) elements of $\Omega = \Upsilon\Upsilon$. This allows us to write $\mathbf{p}' - \mu_p = \Upsilon\eta$ where η is a vector of shocks drawn from a standard normal distribution. Writing in more expansive notation yields

$$\begin{bmatrix} p_{t_1} - \mu_p \\ p_{t_2} - \mu_p \\ p_{t_1+1} - \mu_p \\ \vdots \\ p_{t_2-1} - \mu_p \end{bmatrix} = \begin{bmatrix} \upsilon_{11} & 0 & 0 & \dots & 0 \\ \upsilon_{21} & \upsilon_{22} & 0 & \dots & 0 \\ \upsilon_{31} & \upsilon_{32} & \upsilon_{33} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \upsilon_{\tau+11} & \upsilon_{\tau+12} & \upsilon_{\tau+13} & \dots & \upsilon_{\tau+1\tau+1} \end{bmatrix} \begin{bmatrix} \eta_{t_1} \\ \eta_{t_2} \\ \eta_{t_1+1} \\ \vdots \\ \eta_{t_2-1} \end{bmatrix}. \tag{69}$$

Since we know p_{t_1} and p_{t_2} we can solve for η_{t_1} and η_{t_2} directly from

$$\begin{aligned} (p_{t_1} - \mu_p) &= \upsilon_{11}\eta_{t_1} \\ (p_{t_2} - \mu_p) &= \upsilon_{21}\eta_{t_1} + \upsilon_{22}\eta_{t_2}. \end{aligned}$$

$\{\eta_{t_1+1}, \eta_{t_1+2}, \ldots, \eta_{t_2-1}\}$ are random draws from a standard normal distribution. Once the η vector is constructed, we use equation (69) to compute the simulated price vector $\mathbf{p}' = \Upsilon\eta + \mu_p$. Note that each of the simulated prices is a function of η_{t_1} and η_{t_2}.

To construct a single simulation for the entire time period we repeated this procedure for each interval between successive purchase dates. For each interval, we then applied an acceptance/rejection criterion. Since our model implies that the firm makes a purchase whenever current inventories fall below the order threshold $s(\exp(p))$, we rejected paths such that $exp(p_t) < s^{-1}(q_t)$ for any $t_1 < t < t_2$. For each interval, we repeated the procedure described above until we found a path that did not violate the order threshold constraint. For both products there are intervals in the price series in which we could not find any acceptable paths. In these cases, we accepted one of the rejected price paths.