Econometric Methods for Endogenously Sampled Time Series: The Case of Commodity Price Speculation in the Steel Market

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

May 24, 2002

Abstract: This paper studies the econometric problems associated with estimation of a stochastic process that is *endogenously sampled.* Our interest is to infer the law of motion of a discrete-time stochastic process $\{p_t\}$ that is observed only at a subset of times $\{t_1,\dots,t_n\}$ that depend on the outcome of a probabilistic sampling rule that depends on the history of the process as well as other observed covariates x_t. We focus on a particular example where p_t denotes the daily wholesale price of a standardized steel product. There is no centralized spot market for steel, which is best described as a "dealer market" where individual transactions result from private bilateral negotiations between buyers and sellers, typically intermediated by middlemen known as *steel service centers.* Although there is no central record of daily transactions prices in the steel market, we do observe transaction prices for a particular trader — a middleman (i.e. a steel service center) that purchases steel in the wholesale market for subsequent resale in the retail market. The endogenous sampling problem arises from the fact that we only observe p_t on the days that the trader decides to make purchases. We present a parametric analysis of this problem under the assumption that the timing of steel purchases is part of an optimal trading strategy that maximizes the intermediary's expected discounted trading profits. We derive a parametric partial information maximum likelihood (PIML) estimator that solves the endogenous sampling problem and efficiently estimates the unknown parameters of a Markov transition probability that determines the law of motion for the underlying $\{p_t\}$ process. The estimator also yields estimates of the structural parameters underlying that determine the optimal trading rule. We also introduce an alternative consistent, less efficient, but computationally simpler *simulated minimum distance* (SMD) estimator that avoids high dimensional numerical integrations required by the PIML estimator. Using the SMD estimator, we provide estimates of a truncated lognormal AR(1) model of the wholesale price processes for particular types of steel plate. We use this to infer the fraction of the intermediary's discounted profits that are due to the markups it charges its retail customers, and what fraction is due to pure commodity price speculation, i.e. its success in timing purchases of steel in order to profit from "buying low and selling high."

Keywords: endogenous sampling, Markov processes, maximum likelihood, simulation estimation
JEL classification: C13-15

[*]Corresponding author. Department of Economics, University of Maryland, College Park, MD 20742, phone: (301) 405-3489, fax: (301) 405-3542 e-mail: `jrust@gemini.econ.umd.edu`, web page: `http://gemini.econ.umd.edu/jrust`. This paper was originally prepared for the Cowles Foundation Econometrics Conference, October 23-24, 1999. We thank our discussant, Halbert White, and conference participants including Robert Engle and George Tauchen for helpful comments. We also received helpful feedback from Kenneth Wolpin and other seminar attendees at a subsequent presentations of this paper at the University of Pennsylvania, Boston University, John's Hopkins, Rice, Penn State, the Universities of Maryland, Pittsburgh, Virginia and Florida, the 2001 Conference of the Society for Computational Economics, and the 2001 Midwest Econometrics Group Annual Meeting. We are grateful for financial support from National Science Foundation grant SES-9905145. John Rust is grateful to the Alfred P. Sloan Foundation for research support via its funding of the proposal, "Does Computational Complexity Limit Knowledge in Economics?"

1 Introduction

This paper studies the econometric problems associated with estimation of a stochastic process that is *endogenously sampled.* Our interest is to infer the law of motion of a discrete-time stochastic process $\{p_t\}$ that is observed only at a subset of times $\{t_1,\ldots,t_n\}$ that depend on the outcome of a probabilistic sampling rule that depends on the history of the process as well as other observed covariates x_t. We focus on a particular example where p_t denotes the daily wholesale price of a standardized steel product. There is no centralized spot market for steel, which is better described as a "dealer market" where individual transactions result from private bilateral negotiations between buyers and sellers, typically intermediated by middlemen known as *steel service centers.*[1] Although there is no central record of daily transactions prices in the steel market, we do observe transaction prices for a particular trader — a middleman (i.e. a steel service center) that purchases steel in the wholesale market for subsequent resale in the retail market. *The endogenous sampling problem arises from the fact that we only observe p_t on the days that the trader decides to make purchases.*

We present a parametric analysis of this problem under the assumption that the timing of steel purchases is part of an optimal trading strategy that maximizes the intermediary's expected discounted trading profits. We derive 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. We also introduce an alternative consistent, less efficient, but computationally simpler simulated minimum distance (SMD) estimator that avoids high dimensional numerical integrations required by the PIML estimator. The SMD estimator is essentially an an application of the simulated moments estimator (SME) introduced by Lee and Ingram (1991) and Duffie and Singleton (1993) to the case of endogenously sampled data. Using the SMD estimator, we provide estimates of a truncated lognormal AR(1) model of the wholesale price processes for particular types of steel plate. We use this to infer the fraction of the intermediary's discounted profits that are due to the markups it charges its retail customers, and what fraction is due to pure commodity price speculation, i.e. its success in timing purchases of steel in order to profit from "buying

[1] It is a puzzle why centralized exchanges exist for some commodities such as pork bellies, but not for steel. Rust and Hall (2003) develop a theory of intermediation in which the microstructure of trade in a commodity or asset is endogenously determined. Depending on the parameters of this model there are equilibria consistent with all trade occuring via a *market maker* on a centralized exchange, or all trade occurring via decentralized transactions with *middlemen*, or some mixture of middlemen and market makers. This theory could explain the variety of different trading institutions that we see in different markets.

low and selling high."

This paper was motivated by previous work (Hall and Rust, 1999, 2000 and 2001) on modeling the speculative trading and inventory investment decisions of a particular steel wholesaler. This firm does minimal production processing: its main activity is to stockpile quantities of various types of steel via bulk purchases at wholesale prices from producers and other large intermediaries in order to profit from subsequent resale to retail customers at a mark-up. This firm has provided us with a unique new data set with daily observations on purchases and sales of the more than 2,300 products it carries. While these data are unique in their level of detail and quality, the firm does not record any prices in its computerized data base unless a purchase, sale, or adjustment occurs. *The essence of the endogenous sampling problem is that we only observe purchase prices on the days that purchases occur.*

Let $\{p_t\}$ denote the stochastic process representing the *lowest price* offered by any seller of a particular steel product on day t. We assume that the firm observes p_t at each day t, but it only records p_t when it decides to place an order. Let q_t^o denote the quantity orders (purchased) on day t. The endogenous sampling rule can be stated as follows:

$$p_t \quad \text{is observed} \quad \Longleftrightarrow \quad q_t^o > 0.$$

It is notationally convenient to treat the endogenous sampling problem as a censored sampling problem: i.e., we set p_t to some arbitrary value such as $p_t = 0$ when $q_t^o = 0$, and let p_t equal the observed purchase price when $q_t^o > 0$. Note that we also observe the retail sales prices $\{p_t^r\}$ that the intermediary charges its customers. Since retail sales occur much more frequently than purchases on the wholesale market, retail price data $\{p_t^r\}$ can provide a key source of information for learning about $\{p_t\}$. However on the subset of days where both p_t and p_t^r are observed, we observe that markups $p_t^r - p_t$ are quite volatile, and vary by time, location, and type of the customer. In other words, there is considerable price discrimination in the retail market for steel. As a result the retail price of steel p_t^r is best regarded as a noisy and biased signal of the wholesale price p_t and therefore may not provide information that is directly relevant for estimating the unknown parameters of the wholesale price process.

Our treatment of the wholesale price process $\{p_t\}$ as an exogenously specified "forcing process" that is known up to a finite number of parameters is admittedly only a first approximation to reality. The assumptions that $\{p_t\}$ is observed each day by the intermediary and evolves as an exogenous stochastic process (i.e. its realizations do not depend on actions of the intermediary) are particularly strong restrictions that we intend to relax in future work. As we noted above, prices in the steel market are determined via bilateral negotiations: there is no central market place where the lowest price can be easily observed. Instead, in order to get price quotes, purchasing agents within the firm must communicate with steel producers or other intermediaries via telephone, fax, telex, or recently, the WWW. Thus each price quote involves a small monetary and time cost. However this leads potential endogeneity problems, since the best price the firm is able to negotiate depends on the intensity of its search/bargaining process, and this intensity level could vary depending on the conditions it faces. We defer the difficult issues associated with potential endogeneity in $\{p_t\}$ to future research. However while a more realistic model of speculation would result in a more complicated dynamic programming problem, we believe the general approaches to estimation of the underlying price processes described in this paper will still apply, (although with some modification since when there is no spot market and the "law of one price" does not hold, we would need to estimate a conditional probability distribution representing the trader's beliefs about the distribution of prices in the marketplace given his information).

We ignore the classical "simultaneous equations" endogeneity problems in $\{p_t\}$ since the intermediary's purchases in the wholesale market constitute a negligible fraction of total steel transactions in any given day. Consequently, the assumption that the firm has no influence on wholesale prices is a good first approximation. We also believe our assumption that the firm is continuously aware of the best wholesale price for each of its products is a good first approximation. Indeed, if this firm is to be successful it must have a comparative advantage in "shopping the market". This can be obtained by making investments in communications technologies and personal contacts to reduce the marginal cost of searching for prices as low as possible. Once these investments are in place, it is reasonable to presume that on any given business day, the firm is aware of the best wholesale price p_t for any one of its products. If the firm's actions do not have a significant impact on the best available purchase price, we assume its beliefs about tomorrow's best wholesale price p_{t+1} can be represented by a transition density $g(p_{t+1}, z_{t+1}|p_t, z_t)$, where z_t is a vector of additional observable state variables that help the firm predict future wholesale prices and retail demand for its products.

Our initial interest is to summarize the law of motion for wholesale prices using a flexibly parameterized specification for $g(p_{t+1}, z_{t+1}|p_t, z_t)$. Once we have learned more about the nature of wholesale steel prices and the behavior of traders within this market, our ultimate interest is to analyze the endogenous determination of $\{p_t\}$ in the context of a model of equilibrium price dispersion that allows for positive transactions costs, costs of search, and other frictions that prevent the "Law of One Price" from holding. If it is possible to develop a tractable model that allowed for search, bargaining, and the endogenous determination of equilibrium prices in the steel market, we could use this model and other available aggregate and market-level information to obtain more efficient estimators of the parameters of the wholesale price process $\{p_t\}$. However an estimator that requires nested solutions of an equilibrium model of price dispersion is likely to be a very computationally intensive. In this paper we demonstrate that it is possible to obtain consistent but less efficient estimators that do not require a full blown model of how p_t is determined, and ignore the market and firm-level endogeneity issues associated with the endogenous determination of p_t. Even so, the estimators we propose are still computationally demanding because they deal with a different set of endogeneity issues, namely those associated with the way $\{p_t\}$ is *sampled.*

The method we propose requires repeatedly solving a rather complicated dynamic programming problem to determine the intermediary's optimal trading strategy for each trial value for the unknown parameter vector θ. Extending a seminal result by Scarf (1959) for a simpler class of inventory investment problems, Hall and Rust (2001) proved that the optimal speculative investment strategy for a fairly general class of commodity price speculation problems takes the form of a *generalized* (S,s) *rule* where p is the current spot price of steel and x represents a vector of other state variables (including interest rates, demand shifters, and variables representing unobservable cost shocks) affecting the firm's investment decisions. The functions $S(p,x)$ and $s(p,x)$ satisfy $s(p,x) \leq S(p,x)$. The lower band $s(p,x)$ is the firm's *order threshold:* it is optimal for the firm to place an order whenever its current inventory level q falls below $s(p,x)$. The upper band $S(p,x)$ is the firm's *target inventory level:* whenever the firm places an order to replenish its inventory, it orders an amount sufficient to insure that inventory on hand (the sum of the current inventory plus new orders) equals $S(p,x)$.

The order threshold function $s(p,x)$ is the source of the endogenous sampling problem since the firm only records the spot price p on those days where a purchase occurs. Therefore we have the following *endogenous sampling rule:*

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

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, we propose tractable estimators that are able to consistently estimate θ despite the absence of complete information on $\{p_t\}$.

We have focused on the issue of endogenous sampling since we view it as a first order problem that impedes our ability to learn about the law of motion of $\{p_t\}$. Our theory implies that a profit-maximizing intermediary is more likely to purchase steel when p_t is low than when it is high, which makes it clear that ignoring the endogeneity problem and treating the observations of $\{p_t\}$ as occurring on a randomly selected subset of dates could lead to misleading inferences about $\{p_t\}$. For example, we show via simulations that the sample average of p_t at observed purchase dates is an inconsistent and substantially downward biased estimator of the unconditional long run mean of wholesale prices. By developing econometric methods that successfully deal with the endogeneity of the sampling process, we believe we can overcome the most important endogeneity problem limiting our ability to learn about price determination in the steel market. In subsequent work we plan to revisit the other endogeneity problems connected with search and bargaining and equilibrium price determination that we have outlined above. In the remainder of this paper we proceed under the assumption that our failure to model equilibrium in the steel market will affect our ability to learn about price determination in this market only to the extent that the equilibrium law of motion for $\{p_t\}$ cannot be well approximated by a flexibly parameterized specification for the transition density $g(p_{t+1}, z_{t+1}|p_t, z_t)$.

The problem of endogenous sampling arises in a variety of contexts, including financial applications where transaction prices are sampled at irregularly spaced intervals (see Engle and Russell, 1999, and Russell and Engle, 1998), and in marketing applications where prices for household purchases are only observed on dates purchases are made (see Allenby, McCulloch and Rossi 1996, Erdem and Keane, 1998). However we are not aware of any econometric literature that is directly relevant for handling endogenous sampling issues in a time series context. The most directly related work is the literature on likelihood-based methods for correcting for endogenous sampling in cross-sectional and panel contexts (Manski and McFadden, 1981, McFadden, 1997).

The main idea behind the likelihood based approach is to write down a likelihood that reflects a correctly specified probability law for the endogenous sampling scheme. In some cases, consistent, but less efficient quasi maximum likelihood and GMM estimators have been proposed. These estimators work by appropriately re-weighting the observations to adjust for the effects of non-random sampling, similar in some respects to the way the conditional probabilities in the likelihood reflect an appropriate weighting of the outcomes. We follow this general strategy in this paper, and propose a partial information maximum likelihood (PIML) estimator that is consistent and asymptotically normally distributed. However the PIML estimator requires high dimensional numerical integrations that can only be feasibly done via Monte Carlo or quasi Monte Carlo methods.

We introduce an alternative less efficient but computationally simpler simulated minimum distance (SMD) estimator that does not attempt to re-weight the observations in order to insure consistency. Instead, the estimator is similar to the simulated moments estimator (SME) of Lee and Ingram (1991) and Duffie and Singleton (1993). The SMD estimator only relies on the ability to simulate realizations of the optimal trading model. These simulations can then be censored in exactly the same way as the observed data are censored, similar in many respects to the strategy of "data augmentation" used in Bayesian inference of latent variable models. The idea of the SMD estimator is simply 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 equation (1). It should be apparent that although the two estimation methods we present here are specialized to our particular steel example, it should be straightforward to generalize these methods to other types of endogenous sampling problems that arise in a wide variety of other contexts.

Section 2 introduces the steel speculation and inventory problem that motivated this research. Section 3 presents a parametric, full information approach to inference using a generalization of a model of optimal commodity price speculation and inventory investment developed in Hall and Rust (1999,2000,2001). An independent contribution of this section is to provide a tractable specification for unobserved state variables affecting the speculator's trading decisions that accounts for the frequently binding inequality constraints that purchases of steel must be non-negative. The fact that this constraint is strictly binding at $q_t^o = 0$ prevents the use of standard Euler equation methods to uncover the trader's decision rule and the associated endogenous sampling rule for wholesale steel prices. By introducing an unobserved state variable, we are able to derive a nondegenerate conditional probability distribution for q_t^o that allows us to derive a partial information likelihood function for the full set of data that we observe, $\xi_t = (q_t, q_t^o, p_t, p_t^r, z_t)$. We establish the consistency of the PIML estimator by showing that the values of the joint process $\{\xi_{t_i}\}$ on successive purchase dates t_i (when all components of ξ_t are observed) is an *embedded Markov chain.* This allows us to invoke a standard Information Inequality argument to establish the consistency of the PIML estimator. Via a standard Taylor series approximation and an appeal to an appropriate Central Limit Theorem for mixing processes, it is possible to establish the asymptotic normality of the PIML estimator. Section 4 introduces the simulated minimum distance estimator and derives its asymptotic distribution. Section 5 presents some initial Monte Carlo evidence on the performance of the estimators proposed in this paper. Section 6 presents the results of an empirical application to several plate steel products for which wholesale prices are assumed to evolve according to a univariate truncated lognormal AR(1) process. We estimate the unknown parameters of the price process and the unknown parameters affecting the intermediary's cost of purchasing and holding inventory. We then evaluate how well our generalized (S,s) trading strategy fits these data, and use our results to infer the fraction of the intermediary's discounted profits are due to the markups it charges its retail customers, and the fraction that is due to pure commodity price speculation, i.e., its success in timing purchases of steel in order to profit from "buying low and selling high."

We find that a simple version of the model that posits that $\{p_t\}$ is a lognormal AR(1) process does not fit the data well: our model predicts too many small trades relative to what is observed in the data. This suggests that it is necessary to estimate a model that allows for richer dynamics in steel prices, including allowing for "macro shock" variables and overall indices of steel prices that can provide a better approximation of the firm's beliefs about future price movements. However simulations that compare the recommended trades from our optimal trading rule to the actual trades made by the company over the period of our data set reveal that our model results in higher realized trading profits despite the potential misspecification in the AR(1) price process.

2 Endogenous Sampling in a Price Speculation Problem

This section extends an (S,s) model of commodity price speculation developed in Hall and Rust (1999,2000,2001) to allow for additional covariates and unobserved state variables. The extended model provides a framework for conducting inference and analyzing the endogenous sampling problem. Hall and Rust's 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 (1992), Deaton and Laroque (1992) and Miranda and Rui (1997) on the rational expectations commodity storage model. The fixed (S,s) thresholds derived by Scarf are not optimal in the class of problems we are considering due to the presence of serially correlated cost shocks such as the spot price of steel. However Hall and Rust (2001) proved 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 spot price p. In doing so, this work follows a long tradition in the operation research literature on optimal inventory policies for a single item when the procurement price fluctuates.[2] We now show how the generalized (S,s) rule allows us to formulate and solve the problem of endogenous sampling of steel spot prices.

The motivation for this paper a new database that we acquired via personal contacts with a large U.S. steel wholesaler. The firm provided us with an ongoing data feed to on virtually all aspects of its operations, including the purchase and sale prices and quantities and the identities of its customers for all of its 2300+ individual steel products on a daily basis. The empirical results presented in section 6 are based on data on every transaction the firm made between July 1, 1997 to June 30, 2001 for several 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: it does not record any price information on days that it does not transact

[2] In the operations research literature, see Fabian *et. al.* (1959), Kingman (1969), Kalymon (1971), Golabi (1985), Song and Zipken (1993), Moinzadeh (1997), and Ozekici and Parlar (1999). While these papers argue (and several prove) that the firm's optimal decision rule takes the form of a generalized (S,s) policy, we computationally solve and estimate our model and compare the optimal policies to 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.

(either as a buyer or seller of steel). This shortcoming of our dataset is much more important for steel purchases than steel sales, since the firm purchases new steel inventory in the wholesale market much less frequently than it sells steel to its retail customers. Indeed, even for its highest volume products, it makes purchases only about once every two weeks. The (S, s) theory (as well as plain old common sense) suggest that purchases are not made at random. Instead, it is reasonable to presume that the firm is continually watching the wholesale market looking for buying opportunities. This suggests that the purchase prices we observe will tend to be lower on the days that it purchases than on the days that it doesn't purchase.

To illustrate this, we plot the time series of inventories and prices of one of the firm's products in figures 1 and 2. This is one of highest volume products sold by this firm and is considered to be a benchmark product within the industry, reflected by the fact that the prices of several other steel products are computed as function of this product's price. While it is possible to get average transaction prices for certain classes of steel products through trade publications such as the *American Metal Market*, we are not aware of any data that provides a random sample of daily transaction prices for individual steel products. The purchase price trajectories plotted in figure 2 use straight line interpolatations between observed purchase prices at successive purchase dates. The black circle at each purchase date is proportional to the size of the purchase in pounds. This gives us a clear visual confirmation of the prediction that the firm is more likely to purchase large quantities of steel when wholesale prices are low. Above the interpolated sequence of purchase prices is the interpolated series of average daily retail sales prices. We see that the wholesale and retail prices move in a roughly parallel way, although there appears to be considerable day to day variation in retail prices p_t^r that cannot be explained by variation in the wholesale price p_t. Indeed Chan (1999) estimated a components of variance decomposition of retail prices, using daily time dummies to account for variations in wholesale prices (this is possible due to the fact that the firm frequently sells to more than one retail customer on any given business day). Her results suggest that roughly only 50% of the variation in retail prices p_t^r can be accounted for by variation in wholesale prices p_t.

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 wholesaler purchases this product on average once only every ten business days, while it sells this product 3 out every 5 business days. The evidence of opportunistic purchase 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

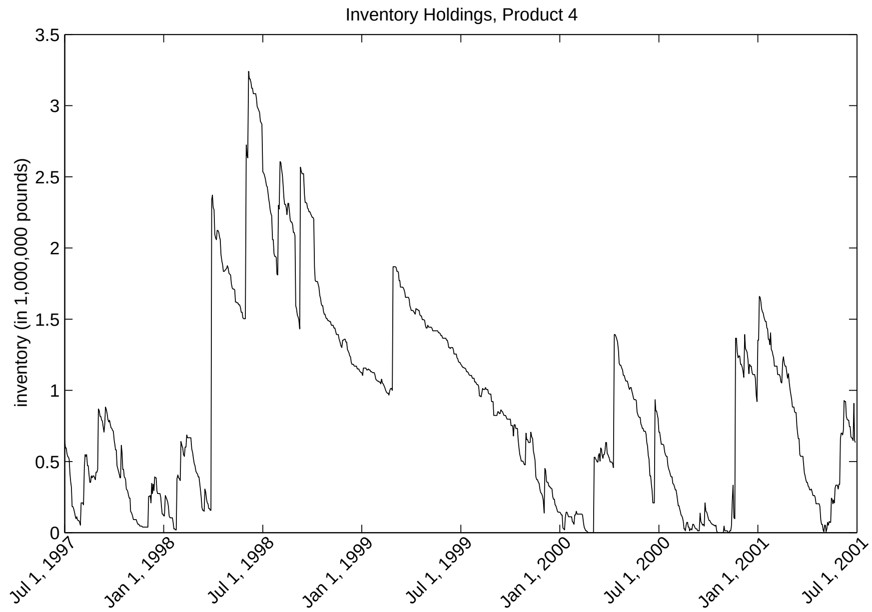

Figure 1: Times series plot of the inventory for product 4.

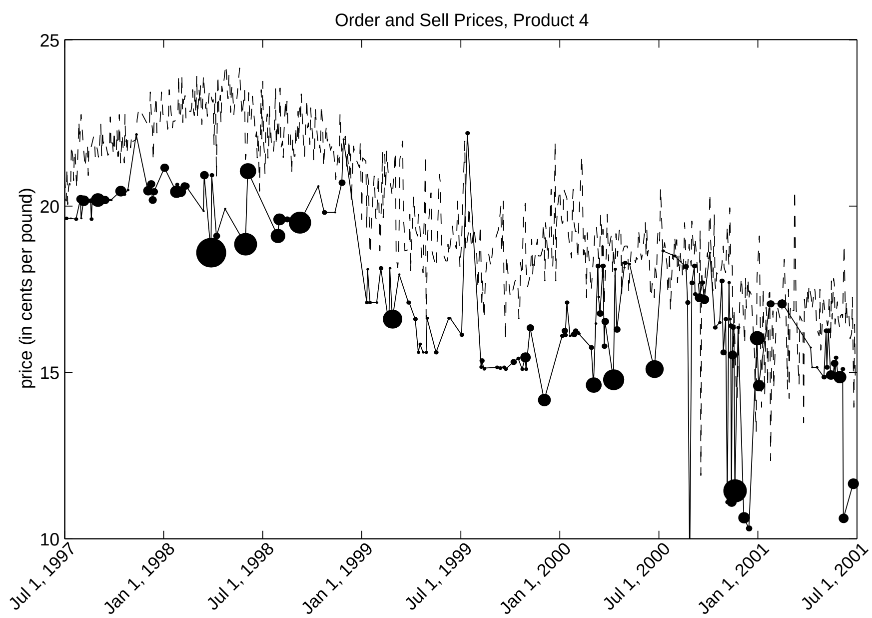

Figure 2: Purchase prices (solid line) and retail prices (dashed line) for product 4. For the purchase price series, the size of the marker is proportional to the size of the purchase.

steel was about 20.5 cents per pound. However the Asian financial crisis started to take a toll on foreign steel producers, which began cutting their prices and increasing their exports in a big way during 1998. 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 clear evidence that the firm is attempting to profit from a "buy low, sell high" strategy.

We now introduce a model of commodity price speculation and describe the (S,s) trading strategy, which in a broad class of models constitutes the optimal strategy for buying low and selling high. We assume that middleman (which we also refer to as the "firm") can purchase unlimited quantities of steel at a time-varying spot price p_t that evolves according to a Markov transition density to be specified in more detail below. We assume that the middleman subsequently sells this steel to retail customers at a retail price p_t^r that includes a randomly varying markup over the current spot price p_t. We observe the firm's quantity on hand q_t and its orders q_t^o on each business day, but we only observe purchase prices and retail sales prices on the days the firm purchases and sells, respectively. Using the convention that $p_t^r = 0$ on days that the firm does not sell any steel to retail customers and $p_t = 0$ on days that it does not purchase any steel, our observations consist of time series on the vector $\{q_t, q_t^o, p_t, p_t^r\}$, and possibly a subset of the components of the vector x_t. We write $x_t = (z_t, \varepsilon_t)$ where z_t are state variables affecting the firm's decisions that we can observe and ε_t denotes variables that we can not observe. We assume that the firm observes all of the variables $\{q_t, q_t^o, p_t, p_t^r, x_t\}$ on each business day.[3]

Let $t = 1, \dots, T$ index T successive business days over which the firm is observed. As discussed above, p_t is observed only when the firm places an order, $q_t^o > 0$, which occurs at a subset of n purchase dates $\{t_1, \dots, t_n\}$ satisfying $1 \leq t_1 \leq t_2 \cdots \leq t_n \leq T$. We now specify the firm's problem in somewhat greater detail in order to derive the likelihood of observing the sequence $\{q_t, q_t^o, p_t, p_t^r, z_t\}$. First, consider the timing of the firm's actions. At each business day t the following sequence of actions occurs:

1. At the start of day t the firm knows its inventory level q_t, the current spot price p_t, and the values of the other state variables x_t.

2. Given (q_t, p_t, x_t) the firm orders additional inventory q_t^o for immediate delivery.

[3] In reality the firm may also not "observe" p_t on certain days where it decides not to undertake search costs to find the best available purchase price for steel, and there is no value for p_t^r on days when it does not receive any retail orders by its customers.

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

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

5. The firm cannot sell more steel than it has on hand, so the actual quantity sold satisfies

$$q_t^s = \min\left[q_t + q_t^o, q_t^r\right]. \tag{3}$$

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:

$$q_{t+1} = q_t + q_t^o - q_t^s. \tag{4}$$

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

8. A new value of the unobservable state variable ε_{t+1} is drawn from the density $\phi(\varepsilon)$.

Note that we abstract from delivery lags and assume that the firm cannot backlog unfilled orders. Thus, whenever demand exceeds quantity on hand, the residual unfilled demand is lost. 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 (3). Since the quantity demanded has support on the $[0, \infty)$ interval, equation (3) implies that there is always a positive probability of a *stockout* given by:

$$\delta(q, p, p^r, z) = 1 - H(q | p^r, p, z). \tag{5}$$

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 spot price p_t. Presumably for most transactions we should have $p_t^r \geq p_t$, reflecting nonnegative markups over the current spot price of steel. However as noted above markups vary in an apparently random fashion from day to day, so at best p_t^r is a biased and noisy indicator of the wholesale price p_t. In this version of the paper 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

$\gamma(p_t^r|q_t+q_t^o,p_t,z_t)$. This way of modeling prices is sufficiently flexible to be consistent with a variety of theories of bargaining and price discrimination by the firm.[4]

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

$$H(q^r|p^r,p,z)=\eta(p^r,p,z)+[1-\eta(p^r,p,z)]\int_0^{q^r}h(q|p^r,p,z)dq, \tag{6}$$

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 $ES(p,q,z)$ is:

$$\begin{aligned} ES(p,q,z) &= E\{\tilde{p}^r\tilde{q}^s|p,q,z\} \\ &= E\{\tilde{p}^r E\{\min[q,\tilde{q}^r]|p^r,p,q,z\}|p,q,z\} \\ &= \int_0^\infty p^r[1-\eta(p^r,p,z)]\left[\int_0^q q^r h(q^r|p^r,p,z)dq^r+\delta(q,p^r,p,z)q\right]\gamma(p^r|q,p,z)dp^r. \end{aligned} \tag{7}$$

We now briefly describe the per period profit function. The firm incurs various costs of holding inventory which we summarize by the *holding cost function* $c^h(q,p,z)$. Following much of the literature on commodity storage (going back to Kaldor, 1939) we interprete the cost of storage as comprising of not only the physical cost of storage, but also a "convenience yield." This marginal convience yield represents the amount the middleman is willing to pay to hold a buffer stock or precautionary level of inventories. This captures any other reasons besides pure price speculation and a stockout motive the firm may have to want to vertically raise the order threshold band s. In particular an increase the order threshold reduces

[4] In Hall and Rust (1998b), we 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 Hall and Rust (1998b) model requires a two-dimensional optimization instead of an one-dimensional optimization at each iteration of the Bellman equation. Second, when solving the Hall and Rust (1998b) we can not restrict our search to just generalized (S,s) policies as we can when we solve the model presented here. Consequently the computational time required for solving the Hall and Rust (1998b) model is such that implementation of our estimation proceedure is not currently feasible.

the firm's vulnerability to interruptions in deliveries.[5] Finally, the firm faces a fixed cost K associated with placing new orders for inventory, which we embody in the *order cost function* $c^o(q^o, p, \varepsilon)$ given by

$$c^o(q^o, p, \varepsilon) = \begin{cases} (p+\varepsilon)+K & \text{if} \quad q^o > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{8}$$

where ε constitutes additional per unit costs associated with placing an order such as transportation costs and unexpected shipping delays.[6] In what follows we will treat ε as an unobserved state variable. It will play a key role in the derivation of the likelihood function in the next section.

Assumption 1: *The intermediary faces an ex post per unit "goodwill cost" of $c^g(p,q,z) \geq 0$ for any unsatisfied demand due to a stockout.*

The speculator's *ex ante* expected total goodwill costs at the beginning of a business day with a spot procurement price (p) and inventory on hand of q is given by:

$$\begin{aligned} EG(p,q) &= E\{\tilde{c}^g(p^r, p, z)\max[(q^r - q), 0]|p, q\} \\ &= \int_0^\infty c^g(p^r, p, z)[1-\eta(p^r, p, z)]\int_q^\infty (q^r - q)h(q^r|p^r, p, x)dq^r\gamma(dp^r|q, p, z). \end{aligned} \tag{9}$$

Assumption 2: *The intermediary has a maximum storage capacity equal to $\overline{q} \leq \infty$.*

Under these assumptions, the intermediary's single-period profits π equals its sales revenues, less the cost of new orders for inventory $c^o(q^o, p, \varepsilon)$ and inventory holding costs $c^h(q+q^o, p, z)$:

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

where $q^s = min[q^r, q+q^o]$. The intermediary's inventory investment behavior is governed by the decision rule:

$$q_t^o = q^o(p_t, q_t, z_t, \varepsilon_t), \tag{11}$$

where the function q^o is the solution to:

$$V(p_t, q_t, x_t) = \max_{q^o} E\left\{\sum_{j=t}^{\infty} \beta^{(j-t)}\pi(p_j, p_j^r, q_j^r, q_j^o + q_j, x_j)\Big| p_t, q_t, x_t\right\}. \tag{12}$$

[5] The firm obtains much of its steel from foreign sources. In the model orders occur instaneously 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 reneg on deals and fail to deliver the steel promised.

[6] Our model assumes that all orders are delivered instantaneously which is obviously not the case in practice. Some purchases, especially those from foreign steel producers in Russia or Japan, can take more than one month to arrive. The random ε_t cost is a crude way of "monetizing" the inconvenience costs of the delivery lag. It is obviously preferable to model delivery lags explicitly, but doing so substantially complicates our model and at the present time we see no tractable way of modeling steel speculation when there are arbitrary stochastic delivery lags. This is another item on our agenda for future research.

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

$$V(p,q,x) = \max_{0 \leq q^o \leq \overline{q}-q} \left[W(p,q+q^o,z) - c^o(q^o,p,\varepsilon)\right], \tag{13}$$

where:

$$W(p,q,z) \equiv \left[ES(p,q,z) - EG(p,q,z) - c^h(q,p,z) + \beta EV(p,q,z)\right], \tag{14}$$

and EV denotes the conditional expectation of V given by:

$$\begin{aligned}
EV(p,q,z) &= E\{V(\tilde{p},\max[0,q-\tilde{q}^d],\tilde{x})|p,q,z) \qquad (15)\\
&= \lambda_1(p,q,z)\int_{p'}\int_{z'}\int_{\varepsilon} V(p',q,z',\varepsilon)\phi(\varepsilon)g(p',z'|p,z)d\varepsilon dp'dz' \\
&+ \lambda_2(p,q,z)\int_{p'}\int_{z'}\int_{\varepsilon} V(p',0,z',\varepsilon)\phi(\varepsilon)g(p',z'|p,z)d\varepsilon dp'dz' \\
&+ \lambda_3(p,q,z)\int_{p'}\int_{z'}\int_0^q\int_{\varepsilon} V(p',q-q',z',\varepsilon)\phi(\varepsilon)h(q'|p,q,z)g(p',z'|p,z)d\varepsilon dq'dp'dz',
\end{aligned}$$

where

$$\begin{aligned}
\lambda_1(p,q,z) &= \int_{p^r} \eta(p^r,p,z)\gamma(p^r|p,q,z)dp^r \qquad (16)\\
\lambda_2(p,q,z) &= \int_{p^r} [1-\eta(p^r,p,z)]\delta(p^r,p,q,z)\gamma(p^r|p,q,z)dp^r \\
\lambda_3(p,q,z) &= \int_{p^r} [1-\eta(p^r,p,z)]\gamma(p^r|p,q,z)dp^r \\
h(q'|p,q,z) &= \int_{p^r} h(q'|p^r,p,q,z)\gamma(p^r|p,q,z)dp^r.
\end{aligned}$$

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

$$q^o(p,q,x) = \inf \underset{0 \leq q^o \leq \overline{q}-q}{argmax} \left[W(p,q+q^o,z) - c^o(q^o,p,\varepsilon)\right]. \tag{17}$$

Note that we invoke the inf operator in the definition of the optimal decision rule in equation (17) 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.

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

$$q^o(p,q,x) = \begin{cases} 0 & \text{if} \quad q \geq s(p,x) \\ S(p,x) - q & \text{otherwise} \end{cases} \tag{18}$$

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

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

$$S(p,x) = q^o(p,0,x). \tag{19}$$

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

$$s(p,x) = inf_{0 \leq q}\{q | q^o(p,q,x) = 0\}. \tag{20}$$

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

In this model the variables q and q^o do not enter as separate arguments in the value function W given in (14): rather they enter as the sum $q+q^o$ as shown in equation (17). 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 sum $q+q^o$, representing inventory on hand at the beginning of the period after new orders q^o have arrived. It follows that if the firm is holding less than its desired level of inventories $S(p_t,x_t)$ 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:

$$\frac{\partial W}{\partial q^o}(p,q+q^o,z) - (p+\varepsilon) = 0. \tag{21}$$

Assuming that W is strictly concave in q, it is clear that there will be a unique value of $q+q^o$ that solves equation (21) for any value of p. Call this solution $S(p,x)$:

$$\frac{\partial W}{\partial q^o}(p,S(p,x),z) = p+\varepsilon. \tag{22}$$

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

In turns out that the function $W(p,q,z)$ may not necessarily be strictly concave but Hall and Rust (2001) show that under fairly general conditions W is K-concave as a function of q for each fixed p.[7] Using the K-concavity property we can prove that whenever $q \geq s(p,x)$, it is not optimal to order: $q^o(p,q,x) = 0$.

[7] A function $W(p,q) : [\underline{p},\overline{p}] \times [0,\overline{q}] \rightarrow R$ 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 $W(p,q+z) - K \leq W(p,q) + z[W(p,q) - W(p,q-b)]/b$.

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 $W(p,q+q^o,z)$ defined in equation (14), where W is defined in terms of the unique solution V to Bellman's equation (13). Under appropriate regularity conditions given in Hall and Rust (2001) (including Assumptions 1 and 2), 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.*

$$q^o(p,q,x) = \begin{cases} 0 & \text{if} \quad q \geq s(p,x) \\ S(p,x) - q & \text{otherwise} \end{cases} \tag{23}$$

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

$$S(p,x) = argmax_{0 \leq q^o \leq \overline{q} - q} \Big[W(p,q^o,z) - c^o(q^o,p,\varepsilon) \Big] \tag{24}$$

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:

$$s(p,x) = \inf_{q \geq 0} \{ q | W(p,q,z) - [p+\varepsilon]q \geq W(p,S(p,x),z) - [p+\varepsilon]S(p,x) - K \}. \tag{25}$$

By a simple substitution of the generalized (S,s) rule in equation (23) into the definition of V in equation (13) we obtain the following corollaries:

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

$$V(p,q,x) = \begin{cases} W(p,S(p,x),z) - [p+\varepsilon][S(p,x) - q] - K & \text{if} \quad q \in [0,s(p,x)] \\ W(p,q,z) & \text{if} \quad q \in (s(p,x),\overline{q}]. \end{cases} \tag{26}$$

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

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

$$S(p,z,\varepsilon) = s(p,z,\varepsilon). \tag{27}$$

3 Maximum Likelihood Estimation

This section derives the likelihood for the commodity price speculation problem presented above. The problem is complicated by the existence of frequently binding inequality constraints on inventory investment, q^o. This implies that it is not possible to use standard Euler equation methods to estimate the unknown parameters of the model via generalized method of moments. Note that Theorem 1 does provide an appropriate first order or "Euler equation" but only on those business days when the firm finds it optimal to purchase:

$$\frac{\partial W}{\partial q}(p, S(p,z,\varepsilon), z) - p = \varepsilon. \tag{28}$$

It is tempting to assume the ε has mean zero and use the Euler equation (28) as a basis for GMM estimation of the 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, as we show in Theorem 2 below, 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 $q < s(p,z)$), 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 a search for an alternative likelihood-based approach. We show how to derive a non-degenerate likelihood function via the inclusion of a single *IID* unobservable state variable ε_t in the firm's optimization problem. The resulting conditional probability distribution function for q^o has a mass point at $q^o = 0$ that respects the frequently binding constraint that inventory investment cannot be negative. This conditional distribution allows us to derive a full-information maximum likelihood estimator that provides a complete solution to the problem of endogenous sampling of the spot price process by integrating out the values of the spot prices in periods where they are unobserved. This likelihood is analogous to the Chapman-Kolmogorov equations for computing multi-step transition probabilities from single step transition probabilities. We will discuss some of the drawbacks of this approach in order to motivate computationally simpler but less efficient simulated minimum distance estimator in section 4.

Some form of measurement error or unobserved state variable must be included as one of the state variables x in the model presented in section 2. The reason is that without such "error terms" the model yields a deterministic optimal decision rule $q^o(p,q,x)$ which can be contradicted by any observation (q^o_t, q_t, p_t, x_t) that does not lie on its graph. We will consider the case where x_t can be decomposed as $x_t = (z_t, \varepsilon_t)$, where z_t are variables which are observed by the econometrician and ε_t is a scalar random variable that is unobserved by the econometrician. This is the main reason for the inclusion of the *IID* process $\{\varepsilon_t\}$ in the previous section. We assume that the distribution of ε_t has support on the entire real line and continuous, strictly positive density $\phi(\varepsilon)$. Theorem 2 below derives the implied conditional distribution of q^o given (p,q,z) formed by integrating out ε from the deterministic decision rule $q^o(p,q,z,\varepsilon)$.

Theorem 2: *Let $\{\varepsilon_t\}$ be an IID process whose density ϕ is continuous and strictly positive over the entire real line. Then we have:*

$$
\begin{aligned}
F(q^o|p,q,z) &= \Pr\{q^o(p,q,z,\varepsilon) \le q^o|p,q,z\} \\
&= \int_{-\infty}^{+\infty} I\{q^o(p,q,z,\varepsilon) \le q^o\}\phi(\varepsilon)d\varepsilon \\
&= \int_{s^{-1}(p,z,q)}^{\infty} \phi(\varepsilon)d\varepsilon \\
&+ I\{S(p,z,s^{-1}(p,z,q)) \le q^o + q \le \overline{q}\} \int_{S^{-1}(p,z,q+q^o)}^{s^{-1}(p,z,q)} \phi(\varepsilon)d\varepsilon \\
&+ I\{q^o + q > \overline{q}\} \int_{-\infty}^{S^{-1}(p,z,\overline{q})} \phi(\varepsilon)d\varepsilon,
\end{aligned}
\tag{29}
$$

where

$$
\begin{aligned}
S^{-1}(p,z,q) &= \inf\{\varepsilon|S(p,z,\varepsilon) = q\} \\
s^{-1}(p,z,q) &= \inf\{\varepsilon|s(p,z,\varepsilon) = q\}.
\end{aligned}
\tag{30}
$$

Let $f = dF$ denote the mixed discrete/continuous conditional density of q^o given (p,q,z). It is given by

$$
f(q^o|p,q,z) = \begin{cases} \int_{s^{-1}(p,z,q)}^{\infty} \phi(\varepsilon)d\varepsilon & \text{if} \quad q^0 = 0 \\ \int_{-\infty}^{S^{-1}(p,z,\overline{q})} \phi(\varepsilon)d\varepsilon & \text{if} \quad q^0 = \overline{q} - q \\ \frac{-\phi(S^{-1}(p,z,q+q^o))}{\partial^2 W/\partial^2 q(p,z,q+q^o)} & \text{otherwise} \end{cases}
\tag{31}
$$

Note that Theorem 2 implies that the transition density for q^o is mixed discrete and continuous, with mass points at $q^o = 0$ and $q^o = \overline{q} - q$, and strictly positive density over the interval $[S(p,z,s^{-1}(p,z,q)) - q, \overline{q} - q]$. However there is a "gap" where there is zero density for q^o in the interval $[0, S(p,z,s^{-1}(p,z,q)) - q]$ since the quantity $S(p,z,s^{-1}(p,z,q)) - q$ represents the minimum order size implied by the (S,s) model in teh stae $(p.q,z)$. The gap is problematic for maximum likelihood estimation since a single observation with an order smaller than the predicted minimum order size would result in a zero value for the likelihood function. The formula for the density of q^o in equation (31) can be derived by differentiating the conditional distribution in equation (29) with respect to q^o for q^o in the interval $[S(p,z,s^{-1}(p,z,q)) - q, \overline{q} - q]$ to obtain:

$$
dF(q^o|p,q,z) = -\phi(S^{-1}(p,z,q+q^o)) \frac{\partial S^{-1}}{\partial q^o}(p,z,q+q^o).
\tag{32}
$$

Using the definition of $S(p,z,\varepsilon)$

$$
\frac{\partial W}{\partial q}(p,S(p,z,\varepsilon),z) = p + \varepsilon,
\tag{33}
$$

and the inverse and implicit function theorems we obtain:

$$\frac{\partial S^{-1}}{\partial q^o}(p,z,q+q^o) = \frac{1}{\partial S(p,z,S^{-1}(p,z,q+q^o))/\partial \varepsilon} = \frac{1}{\partial^2 W(p,q+q^o,z)/\partial^2 q}. \tag{34}$$

Corollary 3 of Theorem 1 implies that when $K=0$, $S(p,z,s^{-1}(p,z,q)) = q$, so that q^o has positive conditional density on the entire range $[0,\overline{q}-q]$ of feasible values for q^o. An alternative way to get a non-degenerate likelihood would be to either a) assume that q^o was measured with error, or b) add a second unobserved state variable that provides a distribution over fixed ordering costs K. Since we obtain direct computerized records from the firm on all prices and quantities, we do not think the measurement error model is a plausible solution to this problem. However adding a second stochastic shock complicates the notation. So in what follows we will simply assume that there are no observed orders that are below the minimum order threshold.

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

Theorem 3: *The (mixed discrete/continuous) density of next period inventory q' given (p, p^r, q, q^o, z) is given by:*

$$\mu(q'|p,p^r,q,q^o,z) = \begin{cases} (1-\eta(p^r,p,z))\delta(p^r,p,q+q^o,z) & \text{if} \quad q'=0 \\ \eta(p^r,p,z) & \text{if} \quad q' = q+q^o \\ (1-\eta(p^r,p,z))h(q+q^o-q'|p^r,p,z) & \text{otherwise} \end{cases} \tag{35}$$

Under our setup, it is easy to see that the observables $\{p_t, p^r_t, q_t, q^o_t, z_t\}$ evolve as a joint Markov process which also has a discrete/continuous transition probability density ρ. We state this as Theorem 4:

Theorem 4: *The joint process $\{p_t, p^r_t, q_t, q^o_t, z_t\}$ is Markov with (discrete/continuous) transition density ρ given by:*

$$\begin{aligned} \rho(p_{t+1}, p^r_{t+1}, q_{t+1}, q^o_{t+}, z_{t+1}|p_t, p^r_t, q_t, q^o_t, z_t) &= g(p_{t+1}, z_{t+1}|p_t, z_t) \\ &\times \mu(q_{t+1}|p_t, p^r_t, q_t, q^o_t, z_t) \\ &\times f(q^o_{t+1}|p_{t+1}, q_{t+1}, z_{t+1}) \\ &\times \gamma(p^r_{t+1}|p_{t+1}, q_{t+1}+q^o_{t+1}, z_{t+1}). \end{aligned} \tag{36}$$

Now consider the full information case where all of the variables $\{p_t, p_t^r, q_t, q_t^o, z_t\}$ are observed over the sample period $t = 0, \ldots, T$.

Definition 1: *The full information maximum likelihood (FIML) estimator* $\hat{\theta}_T^f$ *is defined as:*

$$\hat{\theta}_T^f = \underset{\theta \in \Theta}{argmax}\, l_f(\{p_t, p_t^r, q_t, q_t^o, z_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, z_0, \theta), \tag{37}$$

where l_f is given by:

$$l_f(\{p_t, p_t^r, q_t, q_t^o, z_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, z_0, \theta) = \prod_{t=1}^T \rho(p_t, p_t^r, q_t, q_t^o, z_t | p_{t-1}, p_{t-1}^r, q_{t-1}, q_{t-1}^o, z_{t-1}, \theta). \tag{38}$$

where θ denotes a vector compromising the unknown parameters of the densities $\{f, g, h, \eta, \mu, \gamma, \phi\}$ and unknown parameters entering the firm's cost functions $\{c^o, c^h\}$ and the firm's discount factor β, and Θ denotes a compact parameter space.

Now consider the partial information case where we only observe wholesale prices on the subset of n trading days, $T_n \equiv \{t_1, \ldots, t_n\}$ at which purchases occur. To simplify notation we assume that the data are truncated to begin on the day of the first observed purchase, so $t_1 = 0$, and end on the day of the last observed purchase, $t_n = T$. The relevant likelihood in this case is a marginal likelihood l_p formed by integrating the full likelihood function l_f in equation (38) over spot prices p_t for all time indices t in the complement of T_n. For simplicity, we will consider the case where retail sales are observed in every period. Otherwise, an additional set of integrations would need to be performed over the values of p_t^r for business days t where no retail sales occurred. As noted in the Introduction, it is notationally convenient to convert the endogenous sampling problem into a censored sampling problem by defining an observed censored price sequence $\{p_t\}$ in terms of the underlying uncensored price process $\{p_t^*\}$. Thus, the observed prices p_t are given by:

$$p_t = \begin{cases} p_t^* & \text{if} \quad q_t^o > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{39}$$

Definition 2: *The Partial Information Maximum Likelihood (PIML) estimator* $\hat{\theta}_T^p$ *is defined as:*

$$\hat{\theta}_T^p = \underset{\theta \in \Theta}{argmax}\, l_p(\{p_t, p_t^r, q_t, q_t^o, z_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, z_0, \theta), \tag{40}$$

where l_p is given by:

$$l_p(\{p_t, p_t^r, q_t, q_t^o, z_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, z_0, \theta) = \int \cdots \int \prod_{t=1}^T \rho(p_t, p_t^r, q_t, q_t^o, z_t | p_{t-1}, p_{t-1}^r, q_{t-1}, q_{t-}^o, z_{t-1}, \theta) \prod_{t \notin T_n} dp_t. \tag{41}$$

For subsequent reference it is useful to sketch the derivation of the asymptotic distribution of both the FIML and PIML estimators. The asymptotic distribution of the FIML estimator provides a benchmark for assessing the loss of information entailed by only being able to observe p_t when $q_t^o > 0$. In order to derive the asymptotic properties of both the the FIML and PIML estimators, we need to introduce the concepts of an *embedded Markov chain* and a *segmented Markov chain.* Let $\{\xi_t\}$ denote the joint Markov process in Lemma 3, i.e., the process whose value at t is given by:

$$\xi_t \equiv (p_t, p_t^r, q_t, q_t^o, z_t). \tag{42}$$

Definition 3: *The* **purchase set** Γ *is given by:*

$$\Gamma = \{\xi \,|\, q^o = 0\} \tag{43}$$

and the set of purchase dates $T_n = \{t_1, \ldots, t_n\}$ *is defined recursively as:*

$$t_{i+1} = \inf\{t > t_i \,|\, \xi_t \in \Gamma\}. \tag{44}$$

Definition 4: *Let* $\{\zeta_i\}$ *denote the* **embedded process** *associated with* $\{\xi_t\}$ *and* Γ*. This is the discrete time Markov process which is observed at successive purchase dates* $t \in T_n$*, i.e.,*

$$\{\zeta_i\} = \{\xi_{t_i}\} = \{\xi_t \,|\, \xi_t \in \Gamma\}. \tag{45}$$

A simple application of the Chapman-Kolmogorov equation allows us to derive the transition density ν for the embedded process $\{\zeta_i\}$ as a $t_i - t_{i-1}$-step transition density for successive visits to the purchase set Γ.

Lemma 4: *The embedded process* $\{\zeta_i\}$ *is a Markov chain with transition density* ν *given by:*

$$\nu(\zeta_i|\zeta_{i-1}) = \rho(\xi_{t_i}|\xi_{t_{i-1}}) = \int \cdots \int \prod_{t=t_i+1}^{t_{i+1}} \rho(\xi_{t+1}|\xi_t) d\xi_{t_i+1} \cdots d\xi_{t_{i+1}-1}. \tag{46}$$

Definition 5: *Let* $\{\omega_i\}$ *be the* **segmented process** *associated with* $\{\xi_t\}$*, i.e. the process for which* ω_i *is defined as the realized values of* $\{\xi\}$ *for the sequence of* $t_i - t_{i-1}$ *time periods following the purchase at* t_{i-1} *until the purchase at* t_i*:*

$$\omega_i = (\xi_{t_{i-1}+1}, \ldots, \xi_{t_i}). \tag{47}$$

Notice that the number of components in the segment ω_i is a random variable, equal to the difference $t_i - t_{i-1}$, in the successive times that $\{\xi_t\}$ visits the purchase set Γ.

Lemma 5: *The segmented process* $\{\omega_i\}$ *is a Markov chain with transition density* ν *given by:*

$$\nu(\omega_i|\omega_{i-1}, \theta) = \rho(\xi_{t_i}, \xi_{t_i-1}, \ldots, \xi_{t_{i-1}+1}|\xi_{t_{i-1}}, \theta) = \prod_{t=t_i+1}^{t_{i+1}} \rho(\xi_{t+1}|\xi_t, \theta). \tag{48}$$

Notice that due to the Markov property, only the last element of the segment ω_{i-1}, $\xi_{t_{i-1}}$, is needed to fully determine the conditional probability of $\omega_i = (\xi_{1+t_{i-1}}, \ldots, \xi_{t_i})$. Let $\tau = t_{i+1} - t_i$, be the duration between successive purchases, or in the language of Markov processes, the *recurrence time* for successive visits to the purchase set Γ. If the mean recurrence time to Γ is finite, $E\{\tau\} < \infty$, the process $\{\xi_t\}$ will visit Γ infinitely often and the number of visits n observed over any horizon T tends to infinity with probability 1 as $T \to \infty$.

Assumption 4: *The Markov chain $\{\xi_t\}$ is ergodic (i.e. it possesses a unique stationary distribution), and the purchase set Γ is recurrent (i.e. $E\{\tau\} < \infty$), and the embedded and segmented processes $\{\zeta_i\}$ and $\{\omega_i\}$ are also ergodic Markov chains.*

To study the asymptotic properties of the FIML and PIML estimators, it is useful to rewrite the likelihood functions l_f and l_p as follows:

$$l_f(\{p_t, p^r_t, q_t, q^o_t, z_t\}_{t=1}^T | p_0, p^r_0, q_0, q^o_0, z_0, \theta) = \prod_{i=1}^{n-1} \left[\prod_{t=t_i+1}^{t_{i+1}} \rho(\xi_t | \xi_{t-1}, \theta) \right], \tag{49}$$

$$l_p(\{p_t, p^r_t, q_t, q^o_t, z_t\}_{t=1}^T | p_0, p^r_0, q_0, q^o_0, z_0, \theta) = \prod_{i=1}^{n-1} \left[\int \cdots \int \prod_{t=t_i+1}^{t_{i+1}} \rho(\xi_t | \xi_{t-1}, \theta) dp_{t_i+1} \cdots dp_{t_{i+1}-1} \right]. \tag{50}$$

By Assumption 4 and the Renewal Theorem for Markov chains (see, e.g. Resnick 1992), we have with probability 1

$$\lim_{T \to \infty} \frac{n}{T} = \frac{1}{E\{\tau\}}. \tag{51}$$

Thus, as long as $E\{\tau\} < \infty$, the process $\{\xi_t\}$ visits Γ infinitely often and $n \to \infty$ with probability 1 as $T \to \infty$. Therefore we will carry out the asymptotic analysis indexing the sample size by the number of purchases n rather than the total number of time periods that the process is observed, T. To establish consistency of the FIML and PIML estimators, it is convenient to work with the normalized log-likelihood functions. Multiplying and dividing the likelihoods l_f and l_p in equations (49) and (50) by the product term

$$\prod_{i=1}^{n-1} Pr\{t_{i+1} - t_i | \xi_{t_i}, \theta\} = \prod_{i=1}^{n-1} \left[\int \cdots \int \prod_{t=t_i+1}^{t_{i+1}} I\{\xi_t \in \Gamma(\{t_{i+1}\})\} \rho(\xi_t | \xi_{t-1}, \theta) d\xi_t \right], \tag{52}$$

where

$$\Gamma(\{t_{i+1}\}) = \{\xi_t | q^o_t = 0, t \neq t_{i+1}\} \cup \{\xi_t | q^o_t > 0, t = t_{i+1}\}, \tag{53}$$

and then taking logs and dividing by $n-1$ we obtain for $j = f, p$

$$\frac{1}{n-1} \log l_j(\{p_t, p^r_t, q_t, q^o_t, z_t\}_{t=1}^T | p_0, p^r_0, q_0, q^o_0, z_0, \theta)$$

$$= \frac{1}{n-1}\sum_{i=1}^{n-1}\log\rho_j(\omega_{i+1}|\omega_i, t_{i+1}-t_i, \theta) + \frac{1}{n-1}\sum_{i=1}^{n-1}\log Pr\{t_{i+1}-t_i|\omega_i, \theta\}$$

$$= \frac{1}{n-1}\sum_{i=1}^{n-1} v_1(\omega_{i+1}, \omega_i, \theta) + \frac{1}{n-1}\sum_{i=1}^{n-1} v_2(\omega_{i+1}, \omega_i, \theta), \tag{54}$$

where

$$\rho_f(\omega_{i+1}|\omega_i, t_{i+1}-t_i, \theta) = \frac{\prod_{t=t_i+1}^{t_{i+1}} \rho(\xi_t|\xi_{t-1}, \theta)}{\int\cdots\int \prod_{t=t_i+1}^{t_{i+1}} I\{\xi_t \in \Gamma(\{t_{i+1}\})\}\rho(\xi_t|\xi_{t-1}, \theta)d\xi_t} \tag{55}$$

and

$$\rho_p(\omega_{i+1}|\omega_i, t_{i+1}-t_i, \theta) = \frac{\int\cdots\int \prod_{t=t_i+1}^{t_{i+1}} \rho(\xi_t|\xi_{t-1}, \theta)dp_{t_i+1}\cdots dp_{t_{i+1}-1}}{\int\cdots\int \prod_{t=t_i+1}^{t_{i+1}} I\{\xi_t \in \Gamma(\{t_{i+1}\})\}\rho(\xi_t|\xi_{t-1}, \theta)d\xi_t.} \tag{56}$$

Note that both ρ_p and ρ_f are conditional densities for the segment ω_{i+1} given ω_i and the length of the segment $t_{i+1}-t_i$. The main difference is that due to the censoring, ρ_p can be regarded as the marginal density associated with ρ_f, where we have integrated out the unobserved prices over the interval $(t_i + 1, t_i+2, \ldots, t_{i+1}-1)$ when no purchases occurred. Under suitable regularity conditions on the moments of the functions v_j, $j = 1, 2$, Assumption 4 and the Ergodic Theorem for Markov processes imply that as $n \to \infty$ we have with probability 1:

$$\frac{1}{n-1}\sum_{i=1}^{n-1} v_j(\omega_{i+1}, \omega_i, \theta) \Longrightarrow E\left\{v_j(\omega', \omega, \theta)\right\}, \tag{57}$$

where the expectation is taken with respect to the invariant distribution for (ω', ω) and is given by

$$E\left\{v_j(\omega', \omega, \theta)\right\} = \int\int v_j(\omega', \omega, \theta)d\nu(\omega'|\omega, \theta^*)d\psi(\omega, \theta^*), \tag{58}$$

where $\nu(\omega'|\omega, \theta^*)$ is the transition density for the segmented process given in equation (48) and $\psi(\omega, \theta^*)$ is the invariant distribution for the segmented chain $\{\omega_i\}$. We now provide alternative expressions for these expectations that enables us to verify the consistency of the FIML and PIML estimators. Note that as $n \to \infty$ we have

$$\frac{1}{n-1}\sum_{i=1}^{n-1}\log\rho_j(\omega_{i+1}|\omega_i, t_{i+1}-t_i, \theta) \Longrightarrow E\left\{\log\rho_j(\omega'|\omega, \tau, \theta)\right\}, \tag{59}$$

and

$$\frac{1}{n-1}\sum_{i=1}^{n-1}\log Pr\left\{t_{i+1}-t_i|\xi_{t_i},\theta\right\} \Longrightarrow E\left\{\log Pr\{\tau|\omega,\theta\}\right\}, \tag{60}$$

where τ is the recurrence time to the purchase set Γ. We have

$$\begin{aligned}
&E\left\{\log\rho_j(\omega'|\omega,\tau,\theta)\right\} \qquad (61)\\
&= E\left\{\log\rho(\xi_1,\ldots,\xi_\tau|\xi_0,\tau,\theta)\right\}\\
&= \int_{\xi_0}\sum_{\tau=1}^{\infty}\left[\int_{\xi_1}\cdots\int_{\xi_\tau}\log\rho(\{\xi_t\}_{t=1}^{\tau}|\xi_0,\tau,\theta)p(\{\xi_t\}_{t=1}^{\tau}|\xi_0,\tau,\theta^*)\prod_{t=1}^{\tau}d\xi_t\right]\Pr\{\tau|\xi_0,\theta^*\}d\psi(\xi_0).
\end{aligned}$$

Note that given any ξ_0 and τ, the Information Inequality guarantees that the expression in brackets in (61) is maximized at $\theta=\theta^*$. Similarly we have

$$E\left\{\log\left[Pr\left\{\tau|\xi_0,\theta\right\}\right]\right\} = \int_{\xi_0}\sum_{\tau=1}^{\infty}\log\left[\Pr\{\tau|\xi_0,\theta\}\right]\Pr\{\tau|\xi_0,\theta^*\}d\psi(\xi_0) \tag{62}$$

will also be maximized at $\theta=\theta^*$ given any (ξ_0). This implies that the limiting log likelihood is maximized at θ^*, and standard uniform consistency arguments can be used to show that with probability 1 we have $\hat{\theta}_n^j\to\theta^*$ as $n\to\infty$.

We conclude this section with a brief sketch the derivation of the asymptotic distribution of the PIML and FIML estimators. If model is correctly specified and appropriate regularity conditions hold, the first order conditions for $\hat{\theta}_n^p$ and $\hat{\theta}_n^f$ can be expanded in Taylor series about the true parameter θ^*. Applying a Central Limit Theorem for mixing processes to the key score term in this Taylor series expansion one can show that:

$$\sqrt{n}\left[\hat{\theta}_n^j-\theta^*\right]\longrightarrow N(0,I_j{}^{-1}(\theta^*)),\quad j=p,f \tag{63}$$

where

$$I_j(\theta^*) = I_j^1(\theta^*)+I^2(\theta^*)$$

where

$$I_j^1(\theta^*) = E\left\{\frac{\partial^2}{\partial\theta\partial\theta'}\log\left[\rho_j(\omega'|\omega,\tau,\theta)\right]\right\}\quad \theta=\theta^*$$

and

$$I^2(\theta^*) = E\left\{\frac{\partial^2}{\partial\theta\partial\theta'}\log\left[Pr\left\{\tau|\xi_0,\theta\right\}\right]\right\}\quad \theta=\theta^*.$$

Further it is not difficult to show that since ρ_p is a marginal density of ρ_f, we must have that the difference between the informations $I_f^1(\theta^*)-I_P^1(\theta^*)$ is a positive semi-definitive matrix. This implies that there is indeed a loss of information, and therefore an increase in variance, caused by the endogenous sampling problem. However as long as our model is correctly specified, the PIML estimator will be consistent. If the model is misspecified, then a modification of arguments in White (1982) can be used to show that the PIML and FIML still converge and have an asymptotically normal distribution, but they converge to a value of θ^* that minimizes the Kullbeck-Liebler distance between the parametric model and the true data generating process. The formulas for the asymptotic variance of the estimators must be changed to the outer product of the information and the inverse hessians of the log likelihood when the model is misspecified, since in that case the covariance of $\hat{\theta}_n^j$ is no longer given by the inverse of the information matrix, see White (1982).

The drawback of the PIML estimator is that it is computationally intensive due to the high dimensional integrations that are required to evaluate l_p. Since no purchases of steel are observed on the majority of business days in our sample, the mean time between sales is about 10 business days, so that on average 10 dimensional integrals must be calculated for each term entering the likelihood. Although there have been important advances in simulation estimation and low discrepancy methods for computing high dimensional integrals (see, e.g. Rust, Traub and Woźniakowski, 2002), the PIML will still be a fairly computationally burdensome estimator. A second drawback is that if our interest is primarily on making inferences about the law of motion for $\{p_t, z_t\}$, the other structural parameters that must be estimated to adjust for the endogeneity of the sampling process amount to nuisance parameters. Errors in the specification of the firm's optimal investment and speculation problem will result in inconsistent estimates of the parameters of interest in the transition density $g(p_{t+1}, z_{t+1}|p_t, z_t)$.

It is possible to consider the use of flexible reduced-form specifications for the densities entering the overall decomposition of the transition density ρ given in Theorem 3. However without some strong prior parametric restrictions on some of these densities, it is doubtful that an unrestricted model where the densities (g, μ, f, γ) are treated as unknown objects to be estimated non-parametrically is even identified. In particular the (S,s) model combined with the observations of retail transaction prices provides strong identifying restrictions, limiting how far the wholesale price process $\{p_t\}$ can drift away from observed retail price for a given sequence of observed purchases. In particular, as the implied markup gets larger or smaller, the (S,s) model predicts that the number of orders should be increasing and decreasing in a corresponding fashion. Given the observed sequence of purchases, this property enables us to separately identify the parameters of $g(p', z'|p, z)$ and the structural parameters of the (S,s) model. However if a non-parametric model does not impose any sort of profit maximizing or loss minimizing behavioral motivation on the part of the intermediary, then it appears that the wholesale market price $\{p_t\}$ could drift arbitrarily far away from the retail prices $\{p_t^r\}$ without there being any strong effect on the likelihood of the observed sequences of purchases. This is why we think that it is impossible to non-parametrically identify the form of $g(p', z'|p, z)$ and the trading rule used by the firm when we only have access to endogenously sampled data.

4 Simulated Minimum Distance Estimation

This section introduces a *simulated minimum distance estimator* (SMD) that may be less efficient than the PIML estimator, but which does not require the high dimensional integration and is much easier to compute. 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 calibration process described in section 2 that we used to get our initial "ballpark estimates" of the parameter values. 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 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 estimates of the corresponding moments when there is no censoring, this does not prevent us from deriving a consistent SMD estimator for θ^*. In fact we show that 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 will still require a nested fixed point algorithm to solves for the optimal (S,s) policy and a re-simulation of the model using a fixed set of random shocks (see below) each time the parameter θ is updated, so the SMD estimator is still fairly computationally demanding. Its other 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 $\{\xi_t\}$ denote the censored process introduced in section 3 (i.e. with $p_t = 0$ when $q_t^o = 0$), and let θ denote the $K \times 1$ 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:

$$h_T \equiv \frac{1}{T}\sum_{t=1}^{T} h(\xi_t, \xi_{t-1}), \tag{64}$$

where $J \geq K$. By Assumption 3, 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

$$\sqrt{T}\,[h_T - E\{h\}] \Longrightarrow N(0, \Omega(h)), \tag{65}$$

where

$$\Omega(h) = E\left\{(h(\xi', \xi) - E\{h\})(h(\xi', \xi) - E\{h\})'\right\}, \tag{66}$$

where the expectation in (66) 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

$$\{\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0)\}, \quad t = 2, \ldots, T \tag{67}$$

where for each $t > 1$, $\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0)$ is a continuously differentiable 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 $\{\xi_t(\{u_s\}_{s\le t},\theta)\}$ 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 such smooth simulators, consider the unidimensional case where $\xi_t \in R^1$ for all t. Let $\rho(\xi_{t+1}|\xi_t,\theta)$ 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:

$$\xi_1(u_1,\theta,\xi_0) = P^{-1}(u_1|\xi_0,\theta). \tag{68}$$

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 $t=2,4,\dots$

$$\xi_t(\{u_s\}_{s\le t},\theta,\xi_0) = P^{-1}(u_t|\xi_{t-1}(\{u_s\}_{s\le t-1},\theta,\xi_0),\theta). \tag{69}$$

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 $\xi_t=(p_t,p_t^r,q_t,q_t^o,z_t)$, 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 ρ can be factored as

$$\rho(\xi_{t+1}|\xi_t,\theta) = \rho_2(\xi_{2,t+1}|\xi_{1,t+1},\xi_t,\theta)\rho_1(\xi_{1,t+1}|\xi_t,\theta), \tag{70}$$

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

$$\begin{aligned} \xi_{1,1} &= P_1^{-1}(u_{1,1}|\xi_0,\theta) \\ \xi_{2,1} &= P_2^{-1}(u_{2,1}|\xi_{1,2},\xi_0,\theta). \end{aligned} \tag{71}$$

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:

$$\begin{aligned} \xi_{1,t+1} &= P_1^{-1}(u_{1,t+1}|\xi_t,\theta) \\ \xi_{2,t+1} &= P_2^{-1}(u_{2,t+1}|\xi_{1,t+1},\xi_t,\theta). \end{aligned} \tag{72}$$

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

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

$$h_T(\{u_s\}_{s\leq T},\xi_0,\theta)=\frac{1}{T}\sum_{t=1}^{T}h\left(\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0),\xi_{t-1}(\{u_s\}_{s\leq t-1},\theta,\xi_0)\right). \tag{73}$$

Let $(\{u_s^1\}_{s\leq T},\ldots,\{u_s^S\}_{s\leq T})$ 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 $\{\xi_t(\{u_s^i\}_{s\leq t},\theta,\xi_0)\}$, $i=1,\ldots,S$. Define $h_{S,T}(\theta)$ as the average of S independent time averages $h_T(\{u_s^i\}_{s\leq T},\xi_0,\theta)$

$$h_{S,T}(\theta)=\frac{1}{S}\sum_{i=1}^{S}h_T(\{u_s^i\}_{s\leq T},\xi_0,\theta). \tag{74}$$

Definition 6: *The* **simulated minimum distance estimator** $\hat{\theta}_T$ *is defined by:*

$$\hat{\theta}_T=\underset{\theta\in\Theta}{argmin}(h_{S,T}(\theta)-h_T)'W_T(h_{S,T}(\theta)-h_T), \tag{75}$$

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

$$\{\xi_t(\{u_s\}_{s\leq t},\theta^*,\xi_0)\}\sim\{\xi_t\} \tag{76}$$

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

Later we will relax assumption 5 to allow the parametric model to be misspecified. In that case the asymptotic properties of the SMD estimator 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 $T\to\infty$.

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

Assumption 6: *For any $\theta\in\Theta$ the process $\{\xi_t(\{u_s\}_{s\leq t},\theta,\xi_0)\}$ is ergodic with unique invariant density $\psi(\xi|\theta)$ given by:*

$$\psi(\xi'|\theta)=\int\rho(\xi'|\xi,\theta)d\psi(\xi|\theta). \tag{77}$$

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

$$\begin{aligned} E\{h|\theta\} &= \int h(\xi',\xi)d\rho(\xi'|\xi,\theta)d\psi(\xi|\theta) \\ \nabla E\{h|\theta\} &= \frac{\partial}{\partial\theta}E\{h|\theta\} \\ \nabla h_{S,T}(\theta) &= \frac{\partial}{\partial\theta}h_{S,T}(\theta). \end{aligned} \tag{78}$$

Assumption 7: *θ^* is* **identified***; that is, if $\theta \neq \theta^*$, then $E\{h|\theta\} \neq E\{h|\theta^*\} = E\{h\}$. Furthermore, rank$(\nabla E\{h|\theta\}) = K$ 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

$$\lim_{T\to\infty}\sup_{\theta\in\Theta}\left|(h_{S,T}(\theta)-h_T)'W_T(h_{S,T}(\theta)-h_T)-(E\{h|\theta\}-E\{h|\theta^*\})'W(E\{h|\theta\}-E\{h|\theta^*\})\right| = 0. \tag{79}$$

Assumption 6 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

$$(h_{S,T}(\hat{\theta}_T)-h_T)'W_T\nabla h_{S,T}(\hat{\theta}_T) = 0. \tag{80}$$

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

$$h_{S,T}(\hat{\theta}_T) = h_{S,T}(\theta^*) + \nabla h_{S,T}(\tilde{\theta}_T)(\hat{\theta}_T - \theta^*), \tag{81}$$

where $\tilde{\theta}_T$ denotes a vector that is (elementwise) on the line segment between $\hat{\theta}_T$ and θ^*. Substituting (81) into the first order condition for $\hat{\theta}_T$ in equation (80) and solving for $(\hat{\theta}_T - \theta^*)$ we obtain

$$(\hat{\theta}_T-\theta^*) = -\left[\nabla h_{S,T}(\tilde{\theta}_T)'W_T\nabla h_{S,T}(\hat{\theta}_T)\right]^{-1}\nabla h_{S,T}(\hat{\theta}_T)'W_T\left[h_{S,T}(\theta^*)-h_T\right], \tag{82}$$

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 6 and 7. Now multiply both sides of equation (82) by $\sqrt{T}$ and apply a Central Limit theorem to the difference $\sqrt{T}[h_{S,T}(\theta^*)-h_T]$ to obtain

$$\sqrt{T}\left[h_{S,T}(\theta^*)-h_T\right] \Longrightarrow N\left(0,(1+1/S)\Omega(h,\theta^*)\right). \tag{83}$$

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 5 has the same distribution as $\{\xi_t\}$. As a result each of the terms entering $h_{S,T}(\theta^*)$,

$h_T(\{u^i\},\theta^*)$, has the same probability distribution as h_T and are distributed independently of h_T. The Central Limit Theorem applied to h_T yields

$$\sqrt{T}\left[h_T - E\{h|\theta^*\}\right] \Longrightarrow N\left(0,\Omega(h,\theta^*)\right). \tag{84}$$

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

$$\sqrt{T}\left[h_T(\{u_s^i\}_{s\le T},\theta^*) - E\{h|\theta^*\}\right] \Longrightarrow N\left(0,\Omega(h,\theta^*)\right). \tag{85}$$

Note that

$$[h_{S,T}(\theta^*) - h_T] = \left[\frac{1}{S}\sum_{i=1}^{S}[h_T(\{u_s^i\}_{s\le T},\theta^*) - E\{h|\theta^*\}] + E\{h|\theta^*\} - h_T\right], \tag{86}$$

so that we have

$$\sqrt{T}\left[h_{S,T}(\theta^*) - h_T\right] \Longrightarrow \left[\frac{1}{S}\sum_{i=1}^{S}\tilde{X}_i + \tilde{X}_0\right], \tag{87}$$

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 $\sqrt{T}\left[h_{S,T}(\theta^*) - h_T\right]$ is $N(0,(1+1/S)\Omega(h,\theta^*))$. Using this result and equation (82) we have

$$\sqrt{T}\left[\hat{\theta}_T - \theta^*\right] \Longrightarrow N(0,(1+1/S)\Lambda_1^{-1}\Lambda_2\Lambda_1^{-1}), \tag{88}$$

where

$$\begin{aligned}\Lambda_1 &= [\nabla E\{h|\theta^*\}]' W \nabla[E\{h|\theta^*\}]\\ \Lambda_2 &= [\nabla E\{h|\theta^*\}]' W \Omega(h,\theta^*) W [\nabla E\{h|\theta^*\}].\end{aligned} \tag{89}$$

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

$$\sqrt{T}[\hat{\theta}_T - \theta^*] \Longrightarrow N(0,(1+1/S)\Lambda^{-1}) \tag{90}$$

where:

$$\Lambda = \left[\nabla E\{h|\theta^*\}'[\Omega(h,\theta^*)]^{-1}\nabla E\{h|\theta^*\}\right]. \tag{91}$$

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 $\{\xi_t(\{u_s\}_{s\le t},\theta,\xi_0)\}$ is fairly small, the variance of the resulting estimator is only twice as large as an estimator that computes the expectation of $h_T(\{u\},\theta)$ exactly, such as would be done via Monte Carlo integration when $S\to\infty$.

The SMD estimator can be implemented in practice by solving

$$\hat{\theta}_T = \underset{\theta\in\Theta}{argmin}(h_{S,T}(\theta)-h_T)'\left[\hat{\Omega}(h,\theta)\right]^{-1}(h_{S,T}(\theta)-h_T), \tag{92}$$

where

$$\hat{\Omega}(h,\theta)=\frac{1}{T}\sum_{t=1}^{T}\varepsilon_t(\theta)\varepsilon_t(\theta)' \tag{93}$$

where

$$\varepsilon_t(\theta)=h(\xi_t(\{u_s\}_{s\le t},\theta,\xi_0),\xi_{t-1}(\{u_s\}_{s\le t-1},\theta,\xi_0))-h_T(\{u_s\}_{s\le T},\xi_0,\theta). \tag{94}$$

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

More efficient estimators can be obtained by selecting “efficient moment functions” h such as the score of the partial information maximum likelihood function derived in section 2. Such an estimator can attain the Cramer-Rao efficiency bound derived for the PIML estimator in equation (63). However the score involves 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 the PIML estimator directly. This is a topic for future work. We note that the definition of the SMD estimator can be extended to allow moments formed from the segmented Markov chain $\{\omega_i\}$ defined in section 2. This formulation would be required in the case where h is the score of the partial information likelihood function, since the components of the score involve the segmented chain as shown in section 2. Using moments from the segmented chain involves some minor modifications of the arguments given above. We now do the asymptotics as a function of the number of purchases n rather than the total number of time periods T over which the process is observed. In this case we define the sample moments h_n by

$$\frac{1}{n-1}\sum_{i=1}^{n-1}h(\omega_{i+1},\omega_i), \tag{95}$$

and the simulated moments $h_{S,n}(\theta)$ can be defined accordingly, using the simulated process $\{\xi_t(\{u_s^i\}_{s\le t},\theta,\xi_0)\}$, $i=1,\dots,S$ to construct S *IID* realizations of the segmented process.

Finally, we note that it is appears that it is possible to relax Assumption 5 that the parametric model is correctly specified. As long as assumptions 6 and 7 hold, there will still exist well defined limiting moments for the simulated process, $E\{h|\theta\}$, for each $\theta \in \Theta$. Define θ^* as the value that minimizes the distance between the simulated model and the true data generating process:

$$\theta^* = \underset{\theta \in \Theta}{argmin} \left[E\{h|\theta\} - E\{h\}\right]' W \left[E\{h|\theta\} - E\{h\}\right], \tag{96}$$

where $E\{h\}$ denotes the limit of h_T as $T \to \infty$ for the true data generating process. If the value of θ^* that minimizes this distance is interior to the parameter space Θ, then the following first order condition must hold at θ^*:

$$(E\{h|\theta^*\} - E\{h\})' W \nabla E\{h|\theta^*\} = 0, \tag{97}$$

where $E\{h\}$ denotes the long run or ergodic expectation of h with respect to the true data generating process. This implies that as $t \to \infty$ the random vector

$$X_t \equiv \nabla E\{h|\theta^*\}' W (h(\xi_t(\{u_s\}_{s \leq t}, \theta^*, \xi_0)) - h(\xi_t)), \tag{98}$$

satisfies

$$\begin{aligned} \lim_{t\to\infty} E\{X_t\} &= 0 \\ \lim_{t\to\infty} \text{cov}(X_t) &= \Lambda_2 \end{aligned} \tag{99}$$

for some $J \times J$ covariance matrix Λ_2. However in the misspecified case, Λ_2 may not equal the same formula as the Λ_2 given in equation (89). Using a suitable Central Limit theorem for mixing processes, we should have

$$\sqrt{T} \nabla E\{h|\theta^*\}' W \left[h_T(\{u\}, \theta^*) - h_T\right] \Longrightarrow N(0, \Lambda_2). \tag{100}$$

Following a Taylor expansion argument just as in the correctly specified case above, we should be able to derive the same general form for the asymptotic distribution of $\hat{\theta}_T$ in the misspecified case, i.e.

$$\sqrt{T}\left[\hat{\theta}_T - \theta^*\right] \Longrightarrow N(0, (1+1/S)\Lambda_1^{-1}\Lambda_2\Lambda_1^{-1}), \tag{101}$$

where

$$\Lambda_1 = \left[\nabla E\{h|\theta^*\}' W \nabla E\{h|\theta^*\}\right] \tag{102}$$

and where

$$\sqrt{T} \nabla h_T(\{u_s\}, \theta^*)' W_T \left[h_T(\{u_s\}, \theta^*) - h_T\right] \Longrightarrow N(0, \Lambda_3). \tag{103}$$

The main outstanding issue is to actually establish the limiting asymptotic distribution that is conjectured in (103) and relate the asymptotic covariance matrix Λ_3 to the asymptotic covariance matrix Λ_2 in (100).

The similarity of the two expressions in (103) and (100) suggests that $\Lambda_2 = \Lambda_3$, but further work, including careful attention to regularity conditions, would required to determine whether this is the case. This work, together with Monte Carlo tests and an empirical application of the SMD is currently in progress.

5 Estimation

To illustrate the simulated minimum distance estimator, we 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 spot price, $S(p)$ and $s(p)$. We first estimate the model using data generated from the model itself. In this case, we know the model is correctly specified, and we know the true parameter vector. Second, we estimate the model twice using actual data for two products from the steel service center.

Consider a version of the model in which the firm's general manager solves the following problem:

$$\max_{\{q_t^o\}} E \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \left\{ p_t^r q_t^s - c^o(q_t^o, p_t) - c^h(q_t + q_t^o, p_t) - c^g(q_t^r, q_t + q_t^o) \right\} \tag{104}$$

subject to (3) and (4), and where

$$\begin{aligned} c^o(q_t^o, p_t) &= \begin{cases} p_t q_t^o + K & \text{if} \quad q_t^o > 0 \\ 0 & \text{otherwise,} \end{cases} \\ c^h(q_t + q_t^o, p_t) &= \phi(q_t + q_t^o), \text{ and} \\ c^g(q_t^r, q_t + q_t^o) &= \begin{cases} \gamma & \text{if} \quad q_t^r > q_t + q_t^o \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

As before, the manager takes the spot price p_t and quantity demanded q_t^r as given. The manager knows p_t before deciding q_t^o. The manager then draws q_t^r. The order cost function, $c^o(\cdot,\cdot)$, and the stock-out penalty function, $c^g(\cdot,\cdot)$, are as described in section 2. The holding cost function, $c^h(\cdot)$, is linear, so the marginal cost of storage is a constant.

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

$$log(p_{t+1}) = \mu_p + \lambda_p log(p_t) + w_t^p \tag{105}$$

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 spot price distrubution, we can compute

$$\tilde{\sigma}_p = \sqrt{log(\overline{\sigma}_p^2 + \overline{\mu}_p^2) - 2log(\overline{\mu}_p)}. \tag{106}$$

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

$$\mu_p = (1-\lambda_p)(log(\overline{\mu}_p) - \tilde{\sigma}_p^2/2) \text{ and } \sigma_p = \tilde{\sigma}_p \times \sqrt{1-\lambda_p^2}. \tag{107}$$

The firm sets the retail price by using a fixed linear markup rule over the current spot price:

$$p_t^r = \alpha_0 + \alpha_1 p_t. \tag{108}$$

The firm draws a quantity demanded q_t^r each period. The probability that $q_t^r = 0$ is a fixed, time-invariant constant denoted by η. The mean of the demand distribution is decreasing in p_t.

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

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

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

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

Finally θ denotes the $1 \times p$ (with $p = 13$) parameter vector to be estimated:

$$\theta = \{r, K, \alpha_0, \alpha_1, \lambda_p, \bar{\mu}_p, \bar{\sigma}_p, \bar{\mu}_q, \bar{\sigma}_q, \xi, \gamma, \eta, \phi\}.$$

The SMD estimation procedure requires us to solve for the optimal inventory 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). This PPI algorithm amounts to the following iterative procedure:

1. Approximate the value function $V(p,q)$ given in equation (12) with a finite linear combination of basis functions.

2. Discretize the state space into a finite number of (p,q) pairs.

3. For each iteration i use equation (16) to compute the optimal decision rule $q_i^o(p,q)$ at each of the discretized (p,q) pairs. Although we discretized the state variables, q_i^o is a continuous variables though it is subject to the frequently binding constraint: $0 \le q_i^o \le \bar{q} - q$.

4. Perform a policy iteration step. That is compute

$$V_i(p,q) = E\left\{ \sum_{j=0}^{\infty} (\left(\frac{1}{1+r}\right)^j \pi(p_j, p_j^r, q_j^s, q_i^o(p_j, q_j) + q_j^s) \Big| p, q \right\}. \tag{109}$$

5. Regress the updated value function, $V_i(p,q)$, on the discrete set of p and q's to compute a new parameterized approximation of the value function.

6. Iterate over i on steps 3–5 until the coefficients on the parameterized approximation of the value function converge.

We approximated the value function by a complete set of Chebychev polynomials of degree 3 in p and q. We discretized the state space with 225 grid points (15 in the p dimension and 15 in the q dimension). The grid points are fixed at the Chebychev zeros, so the grid points are more heavily weighted toward the boundaries of the state space.

As can be seen by Bellman's equation (13) each iteration of the algorithm requires the solution of a constrained optimization problem involving the three functions $ES(p,q)$, $EG(p,q)$, and $EV(p,p^r,q)$, each of which is a conditional expectation of functions of two continuous variables (sales, $p^r q^s$, and the value function, $V(p,q)$). Since no analytic solutions to these conditional expectations exist, we resorted to numerical integration. We implemented a "quadrature" approach that approximates EV by a probability weighted sum:

$$\hat{EV}(p,p^r,q) = \frac{1}{N_p}\frac{1}{N_q}\sum_{i=1}^{N_p}\sum_{j=1}^{N_q} I\{q_j \le q\}\hat{V}(p_i,q-q_j)h(q_j|p^r)g(p_i|p) \tag{110}$$

where $h(q_j|p^r)$ is a discretized approximation to the conditional probability density $h(q|p^r,q)$, and $g(p_i|p)$ is a discretized approximation to the transition probability density $g(p'|p)$. Further adjustments to this formula were made in order that $\hat{EV}(p,p^r,q)$ reflects that mass points on stockouts and zero sales as in equation (16).

We have considerable freedom in our choice of moments functions, the h vector, to use in the criterion. The most efficient moment functions we could use would be the score of the partial information maximum likelihood function derived in section 2. Such an estimator could attain the Cramer-Rao efficiency bound derived for the PIML estimator. However as discussed in section 4 the score involves a ratio of integrals. It is not clear that these integrals can be replaced by simulation estimates and still obtain a consistent SMD estimator. Instead we match the means and histograms of the p, p^r, q^o, q^s, and q processes. For example for product 2, we use five moments for the wholesale price, p:

$$\frac{1}{N} \quad \textstyle\sum_{t=1}^{T} \quad p_t \quad \text{where } p_t = 0 \text{ if } q_t^o = 0 \tag{111}$$

$$\frac{1}{T} \quad \textstyle\sum_{t=1}^{T} \quad 100 \times I(p_t > 0) \tag{112}$$

$$\frac{1}{N} \quad \sum_{t=1}^{T} \quad 100 \times I(p_t > 15.85 \text{ and } p_t \leq 16.95) \tag{113}$$

$$\frac{1}{N} \quad \sum_{t=1}^{T} \quad 100 \times I(p_t > 16.95 \text{ and } p_t \leq 19.60) \tag{114}$$

$$\frac{1}{N} \quad \sum_{t=1}^{T} \quad 100 \times I(p_t > 19.60) \tag{115}$$

where I is an indicator function, and N is the number of days an order is placed: $N \equiv \sum_{t=1}^{T} I(p_t > 0)$. So the moment in equation (111) is the mean of the censored order price process. The moment in equation (112) measures the percentage of days a purchase is made; it is just $100 \times N/T$. The moments in equations (113), (114), and (115) are the second, third and fourth quartiles of the censored price process. Similarly, we constructed means and quintile histogram bins for the remaining four series, p^r, q^o, q^s, and q. Finally we computed the percentage of days a sale occurs. So the total number of moments used in the criterion, J, is 26. We set the number of simulations, S, to 10.

The indicator functions in the moment conditions are of course not smooth. However the criterion is visually a smooth function of the parameters with the exception of the fixed cost parameter K. In this case, concentrated "slices" of the criterion function have "steps" and "cliffs." Therefore to verfiy MATLAB's constrained minimization routine `fmincon.m` found a global minimum in each case, we visually inspected various concentrated slices of the criterion function after each estimation.

parameter	truth	point estimate	standard error
K	100	81.7	22.2
α_0	0.90	1.34	0.30
α_1	1.16	1.09	0.02
λ_p	0.970	0.965	0.002
$\overline{\mu}_p$	18.00	18.73	0.32
$\overline{\sigma}_p$	3.50	3.08	0.11
$\overline{\mu}_q$	250.0	255.2	9.3
$\overline{\sigma}_q$	300.0	311.19	24.3
ξ	0.70	0.73	0.18
γ	500	593	229
η	0.350	0.334	0.002
r	0.075	0.039	0.036
ϕ	-0.06	-0.039	0.006
criterion	0.0944		

Table 1: Estimation results on data generated by the model.

In our initial 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 a simulated data set of 1013 periods from the model; second using the simulated data, we estimated the model via our simulated minimum distance estimator. The results from this experiment are reported in table 1. The estimation procedure provides a formal criterion of the validity of model. In bottom row of table 1 we report the value of the objective function evaluated at the SMD point estimates: 0.0944. Since the number of moment conditions exceeds the number of parameters estimated ($J > p$) the model is overidentified. Following Hansen (1982), we use the objective function to test the overidentifing restrictions.

$$(h_{S,T}(\hat{\theta}) - h_T)' \left[\hat{\Omega}(h)\right]^{-1} (h_{S,T}(\hat{\theta}) - h_T) \to \chi^2(J-p) \tag{116}$$

We also report in table 1 the point estimates and standard errors for each of the thirteen parameters along with the parameter values used to create the simulated data. 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 \$100 per order. The value of a lost sale, γ, is \$500. The parameter choices for $\bar{\mu}_p$ and $\bar{\sigma}_p$ implies the uncensored price process with a mean of \$18.00 per hundred-weight or 18 cents per pound and standard deviation of \$3.50 dollars per hundred-weight. The parameter values of $\bar{\mu}_q$, $\bar{\sigma}_q$ and ξ imply the average sale is 250 hundred-weight or 2,500 pounds. The interest rate r is set to 7.5 percent per annum. The linear storage cost net of convenience yield, ϕ is set -6 cents per hundred-weight, so the convenience yield dominates the storage cost.

For most of the parameters, the point estimates are within one or two standard errors of the true values. There are several exception, however. The point estimate for ϕ, the linear storage cost net of convenience yield, is biased toward zero. The are also several large deviations between the true parameter value and point estimates. The widest divergence is the interest rate r. Its point estimate is about half the true value. Both ϕ and r capture the "cost" of holding inventories: r captures the opportunity cost of holding inventories, while ϕ captures the physical holding costs and convenience yield of holding inventories. So a positve r and a negative ϕ could offset each other. The point estimate for the fixed cost parameter, K is also about 20 percent lower than its true value. This parameter largely determines the distance between the S and s bands. To accurately identify this parameter requires numerous observations of days in which the firm is holding inventory levels close to s. Given the relatively few days the firm purchases, there are very few days the firm holds inventories close to s.

We now estimate the model for two products independently. In table 2 we report the point estimates

	Product 2		Product 4	
parameter	point estimate	standard error	point estimate	standard error
K	83.0	10.1	95.6	23.9
α_0	1.10	0.50	1.13	0.77
α_1	1.07	0.02	1.11	0.04
λ_p	0.979	0.00006	0.981	0.0003
$\overline{\mu}_p$	18.34	0.17	18.03	0.30
$\overline{\sigma}_p$	2.80	0.08	3.85	0.16
$\overline{\mu}_q$	258.6	7.7	289.8	9.6
$\overline{\sigma}_q$	215.3	12.5	219.1	15.7
ξ	0.68	0.24	0.74	0.30
γ	471	215	501	362
η	0.344	0.006	0.353	0.003
r	0.0796	0.041	0.071	0.072
ϕ	-0.0705	0.005	-0.0423	0.006
criterion	0.58		0.31	

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

and standard errors for the parameters of the model for products we call product 2 and 4. We also report the minimized SMD estimation crtiterion. Although the model is formally rejected for, it is consistent with the main stylized facts of inventory investment behavior that we observe in our data (for further discussion, see Hall and Rust, 1999).

Although we estimated the parameters for each of these products independently, it is reassuring that the point estimates for almost the entire parameter vector are similar across the two products. It is reasonable to expect that the parameters, K, r should be identical across products. We also expect the parameters α_0, α_1, λ_p, ξ, γ, and ϕ to be quite similar, if not identical, across products.[8] Indeed, with the exception of ϕ, this is the case. The estimates fixed order cost parameter, K are well within two standard errors of each other. When we asked the general manager for his estimate of this parameter he first stated $50. We then asked him whether $80 to $90 seemed reasonable, he said yes. The main fixed cost to ordering is the value of the general manager and his administrative assistant's time in takes to complete the paperwork.

The point estimates across the two products of the interest rate used by the firm to discount future cash flows, r, are with one standard deviation of each other, and the point estimates are quite sensible. The estimated nominal interest rate for the two products is between 7 to 8 percent (7.96% for product 2 and

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

7.11% for product 4). The CEO of the company drew the line at providing us 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 they faced. The average of the 3-month LIBOR rate of the four-year period studied is about 5.75, which implies an average borrowing rate for the firm of about 7.5%.

When we asked the general manager for his estimate of the goodwill cost of turning away a customer due to stock out, he stated it was around $200. When asked whether $400 -$500 was reasonable, he felt that this value was a little too high. For both products, the standard errors for this parameter are quite large. The criterion is clearly very flat along this dimension. The marginal cost of storage parameter, ϕ, is negative for both products so the marginal convenience yield dominates an 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 prices.

One diagnostic of the model and the estimation procedure is to compare the estimated value of η in the two models to the fraction of days no sales are made. The percentage of days a sale occurs is one of the moments used the criteria and should pin down the value of η: $\frac{1}{T}\sum_{t=1}^{T} 100 \times I(p_t^r > 0) \approx 100 \times (1-\eta)$. For product 2, the fraction of days in data without a sale is 0.355 while the point estimate of η is .344 with a standard error of .006. For product 4, the fraction of days in data without a sale is 0.366 while the point estimate of η is .353 with a standard error of .003. While the point estimates for η are close to the moment

The endogenous sampling problem is illustrated in figures 3 and 4. In figure 3 we plot we the $S(p)$ and $s(p)$ bands derived from the optimal decisions rules for the manager's problem using the estimated parameter vector for product 2. 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, as implied by Corollary 2, the order size at s is a decreasing function of the price. In figure 3 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 graphs). In the simulation presented, this occurs less than 15 percent of time.

The infrequency with which purchases are made is also illustrated in figure 4. This figure presents the censored and uncensored purchase price series, p_t. The solid line is the analogue 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 325-350, 475-500 and 575-595) the firm aggressively made purchases to build up large levels of inventories. The

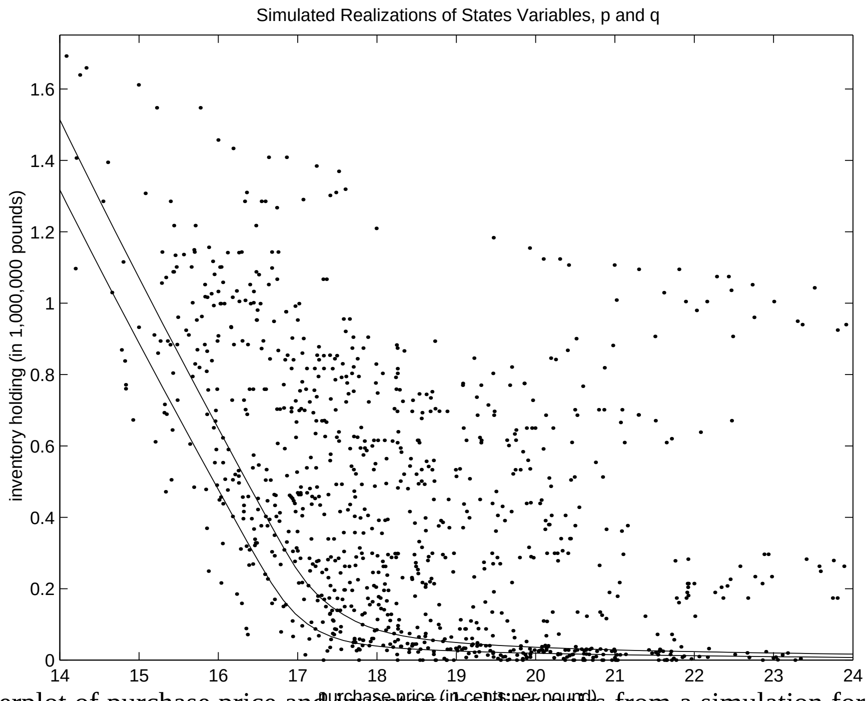

Figure 3: Scatterplot of purchase price and inventory holding pairs from a simulation for product 2. The solid lines are the $S(p)$ and $s(p)$ bands from the model.

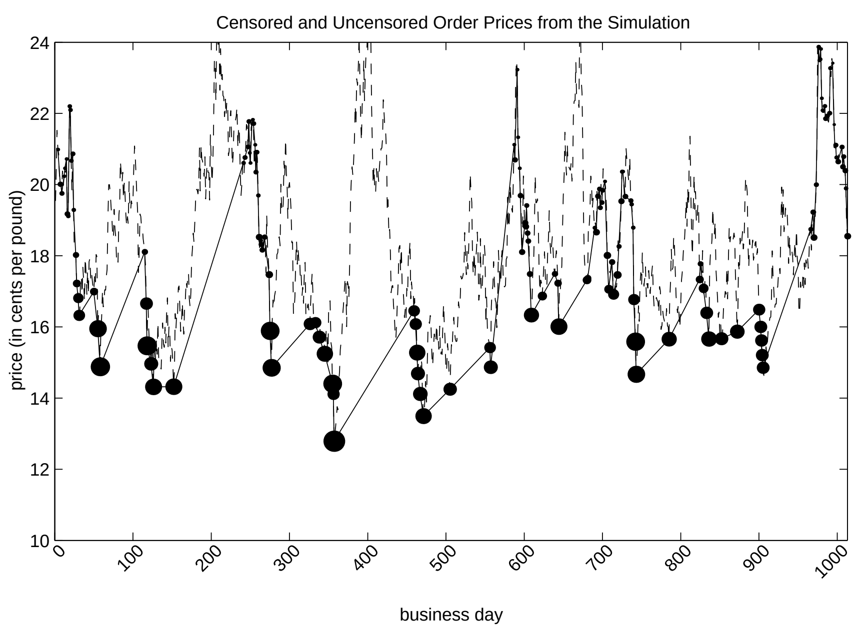

Figure 4: Censored (solid line) and uncensored (dotted line) purchase prices, p_t from simulation for product 2. For the censored series, the size of the marker is proportional to the size of the purchase.

large levels of inventories were slowly drawn down as prices inevitably rose. Thus after exploiting a low price opportunity, the firm may subsequently make no new purchases for many days.

Figures 3 and 4 highlight some of the strengths of the model. The model can capture several of the salient features of the inventory and price data. First, in the data purchases are made infrequently. It is clear from both graphs that the model implies purchases are only made about 15 percent of the time. Second, we observe both small and large purchases in the data. Again this can been in both graphs. In figure 3 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 16 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 4, 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 3 we can see that the model predicts that inventory levels will vary over the sample between 0 and 1.8 million pounds.

The ability of the model to match the data can be seen visually in figure 5-16. In these graphs we plot time series for inventory holdings, prices, and days-supply, histograms for order and sales quantities, and a stem plot of orders. On the left hand side, we plot the actual data for product 2. On the right hand side we plot the output of a simulation for the model using the estimated parameter values for product 2.

The main shortcoming of the estimation is our inability to match the downward trend of the price process. We assume that prices are stationary though highly persistent. Conseqently, as can been seen in the (S, S) band plotted in figure 3, the optimal decision rules imply counterfactually that the firm should make only small purchases and hold low levels of inventories whenever the procurment price is above 17 dollars per hundred-weight. From figures 7 and 9 we see that, for product 2, the firm made large purchases around 19 dollars per hundred-weight in the early part of the sample and around 15 dollars per hundred-weight in the later part of the sample.

If we concentrate out all the other parameters except λ_p the criterion surface is a steep and smooth cup centered around 0.978 so the small standard error associated with the AR(1) is not surprising. But the concentrated criterion surface actualy turns down slighty between .995 and 1.01. (The model still solves numericaly for values of λ_p slightly greater than one.) The global minimum is still centered around 0.978, but there appears to be a local minimum just above 1.00. However if we assume the price process follows a (or a very nearly) random-walk, the optimal decision rules implies frequent small- to medium-size orders

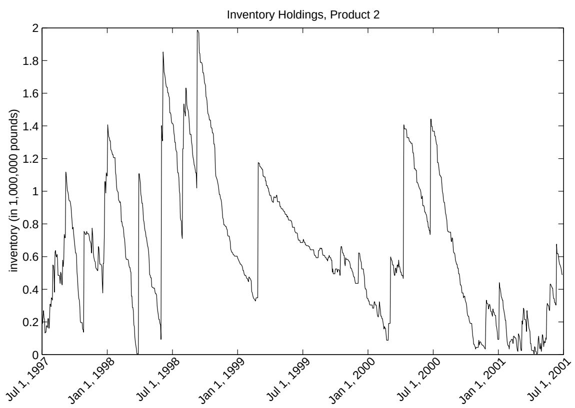

Figure 5: Actual inventory data for product 2.

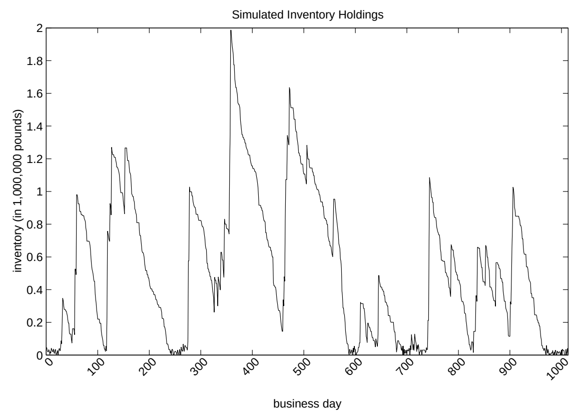

Figure 6: Simulated inventories for product 2.

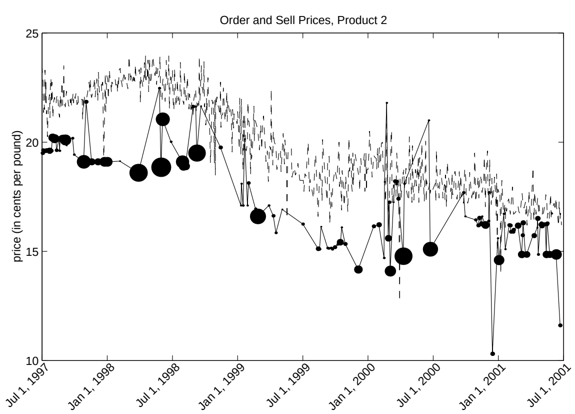

Figure 7: Order and sell prices for product 2.

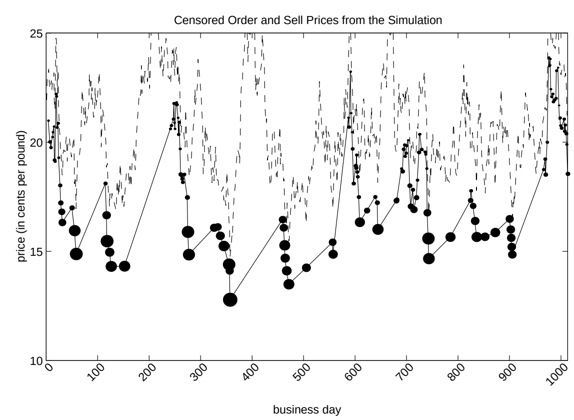

Figure 8: Simulated prices for product 2.

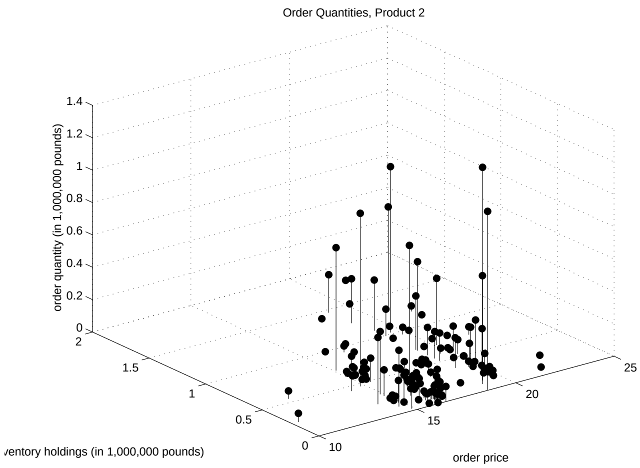

Figure 9: Order size for product 2.

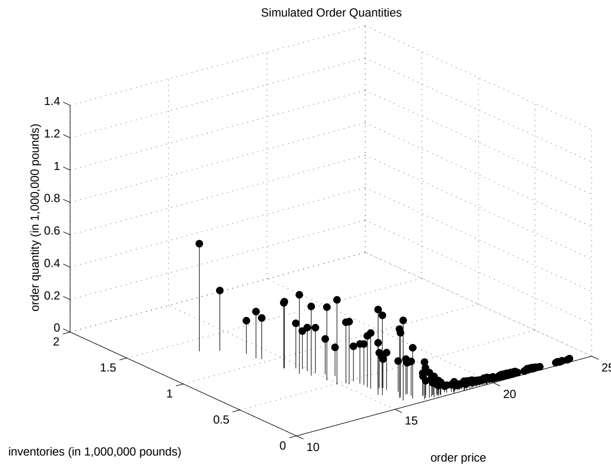

Figure 10: Simulated order size for product 2.

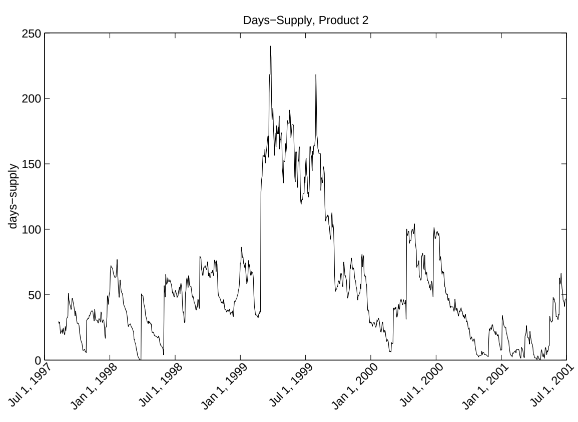

Figure 11: Days-supply for product 2.

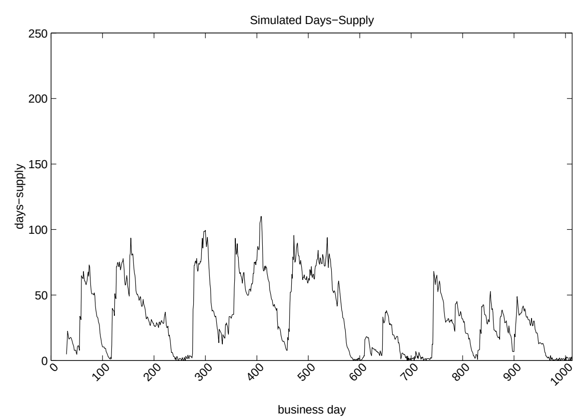

Figure 12: Simulated days-supply for product 2.

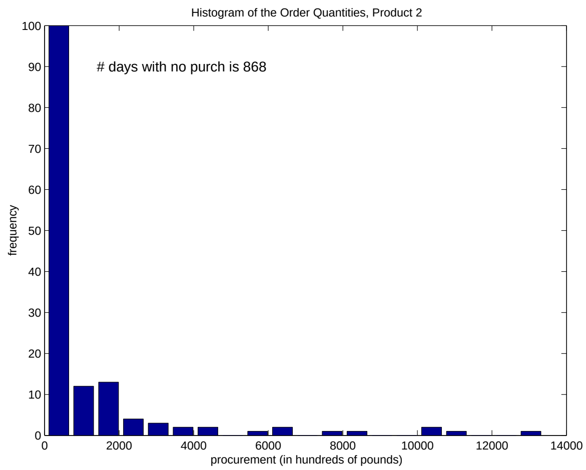

Figure 13: Histogram of procurments for product 2.

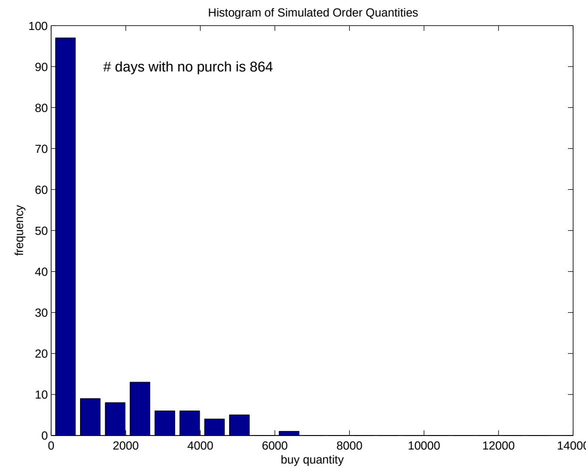

Figure 14: Histogram of simulated orders for prod. 2.

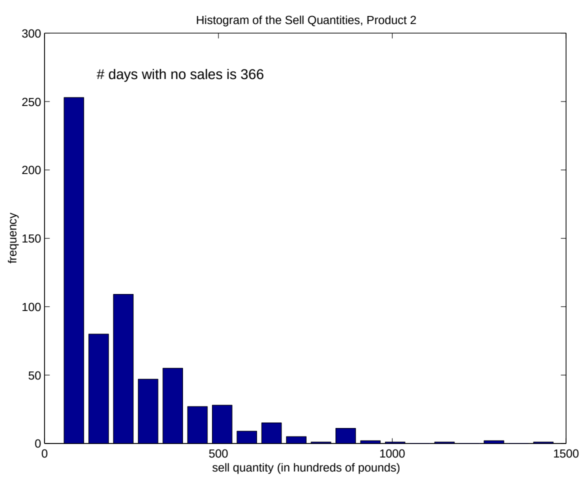

Figure 15: Histogram of daily sales for product 2.

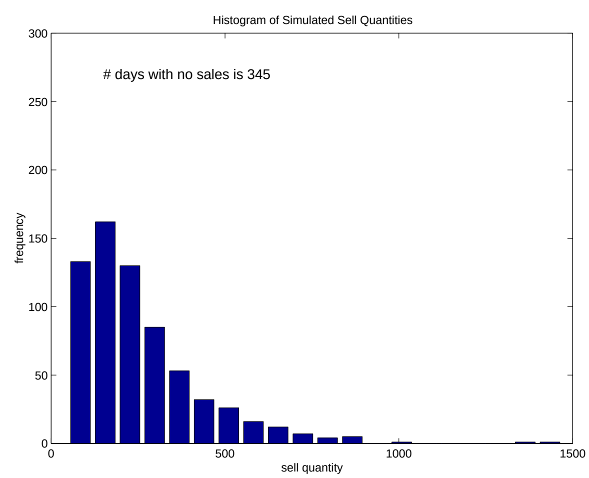

Figure 16: Histogram of simulated sales for prod 2.

	G.M.'s actual performance		Model's Policy Prescription		Long-run simulation	
markup	\$251,493	76%	\$230,108	43%	\$352,129	55%
capital gain	92,503	24%	333,223	58%	278,482	45%
holding cost	447,641		501,680		500,365	
lost sales	0		-28,726		-18,533	
order costs	-10,804		-12,820		-11,607	
total profits	780,833		1,023,464		1,100,836	

Table 3: Profit Decomposition For Product 2 Using Equation 117
Profits in the first two cases cover the 1013 days in our sample and are discounted back to start of the sample.
The third case decomposed the profits implied by the model for a simulation of 1,000,000 periods.
Total profits are the sum of the first five rows.

so that the invnetory level fluctuates closely around a fixed target level. Thus the model which assumes p_t follows a random walk will not imply the large variation in inventory holdings that we see in the data. The other possible solution is to detrend the data. However when we first started working on this project, no one we talked to expected steel price to decline over 25% in four years. To some extent we are just working with too short a sample period.

Finally, from the simulations we can deduce the importance to the firm's profits of the markup relative to the capital gains and losses from the firm's optimal speculation strategy. By substituting the law of motion for inventories, (4), into the firm's objective function, (104) the discounted present value of the firm's profits can be expressed by

$$\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^{t}\pi(p_t,p_t^r,q_t^r,q_t+q_t^o)=\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^{t}(p_t^r-p_t)q_t^s+\sum_{t=2}^{T}\left(\frac{1}{1+r}\right)^{t}(p_t-(1+r)p_{t-1})q_t- \sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^{t}I(q_t^o)K-\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^{t}c^h(q_t+q_t^o,p_t)-\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^{t}c^g(q_t^r,q_t+q_t^o). \quad (117)$$

The first term on the right hand side of equation (117) can be interpreted as the discounted present value of the markup over the current spot price while the second 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 third, fourth, and fifth terms are the discounted present values of the order costs, the holding costs, and the lost sale good will costs incurred by the firm over the sample period.

This decomposition is reported in tables 3 and 4. We first employ this decomposition to evaluate the general manager's actual performance over the four-year sample period for product 2. Since this decomposition depends on the wholesale price path between purchases, we interpolate between the observed purchase dates via importance sampling. That is, we simulate paths conditioning on the intial and terminal

	G.M.'s actual performance		Model's Policy Prescription		Long-run Simulation	
markup	\$287,488	72%	\$238,096	41%	\$502,721	56%
capital gain	116,829	28%	346,861	59%	394,716	44%
holding cost	363,858		166,107		382,479	
lost sales	0		-47,772		-27,168	
order costs	-14,494		-21,134		-15,509	
total profits	753,681		682,158		1,237,239	

Table 4: Profit Decomposition For Product 4 Using Equation 117
Profits are the average for 100 simulated price paths of 1013 business days each.
For the first two cases the simulated price path is constrained to equal the observed prices for the days purchases were made.
Total profits are the sum of the first five rows.

conditions. In particular we note

$$\left[p_{t_{n-1}}, p_{t_n}, p_{t_{n-1}+1}, p_{t_{n-1}+2}, \cdots p_{t_n-1}\right]' = A\left[\eta_{t_{n-1}}, \eta_{t_n}, \eta_{t_{n-1}+1}, \eta_{t_{n-1}+2}, \cdots \eta_{t_n-1}\right]' \tag{118}$$

where $p_{t_{n-1}}$ is the price observed at date and p_{t_n} is the next observed price. The matrix A is the Cholesky decomposition of the variance-covariance matrix of the p vector implied by the autoregressive process estimated above. The vector η is a vector of i.i.d. Normal (0,1) draws. This sampling procedure delivers simulations of $\{p_{t_{n-1}+1}, p_{t_{n-1}+2}, ..., p_{t_n-1}\}$ conditioned on $p_{t_{n-1}}$ and p_{t_n} that are consistent with the autoregressive process estimated above.

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 $p_t < s^{-1}(q_t)$ for any draw whithin the . Using the interpolated price series, we decomposed the firm's profits using the actual data for q_t, q_t^s, and q_t^o and our point estimates for K, r, γ, and ϕ. We report the average decompostion for 100 simulated wholesale price paths. As one can see from the figure 7 the price of product 2 fell steadily over the sample period. Never the less, by our accounting, the firm made well over \$300,000 from the markup and capital gains on each of these two product over the four-year period.[9] Ignoring any returns from the convenience yield, about 75 percent of these profits came from markup, while only about 25 percent came from the capital gains. We find it remarkable that the firm made positive capital gains over this period despite the fact the price of steel fell almost 50 percent during this period.

As a diagnostic of our model we then compared the general manager's performace to the performance of our model. In this experiment, we take as given the 100 interpolated wholesale price series, the quantity

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

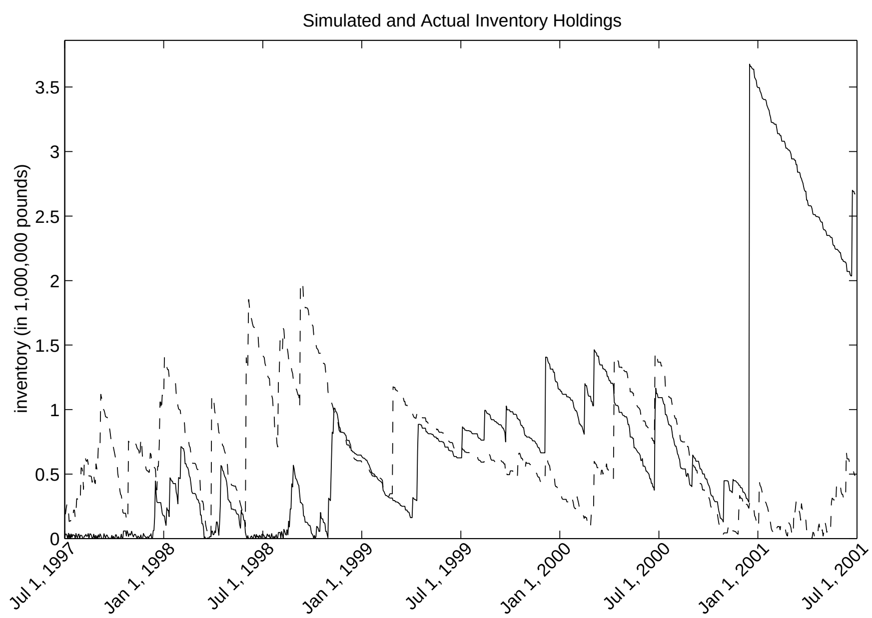

Figure 17: Actual (dashed line) and counter-factual (solid line) inventory holdings for product 2.

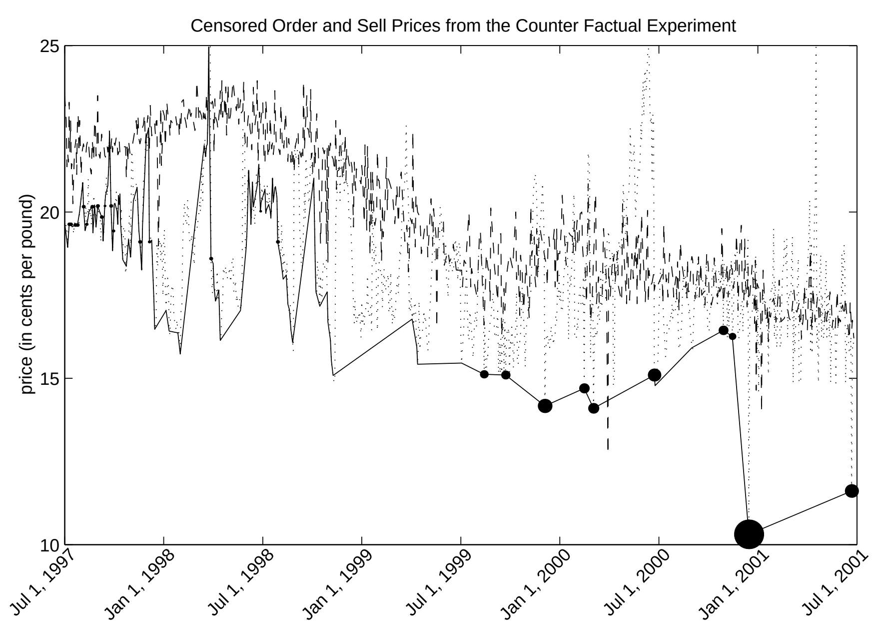

Figure 18: Counter-factual uncensored purchase prices (dotted line), censored purchase prices (solid line), and retail prices (dashed line) for product 2. For the censored purchase price series, the size of the marker is proportional to the size of the purchase.

demanded series, and initial level of inventories for each of the products. But in this case, we let the model's optimal decision rules dictate when and how much to order. Inventories followed the accumulation identity given by equation 4. As reported in tables 3 and 4, had the general manager followed the optimal order strategy implied by our model, his discounted profits from the markup would have been ten to twenty percent lower, but his discounted capital gains would three times larger. The model implies that the firm should aggresively price speculate. In figures 17 and 18 we plot the prices and inventory holdings for one simulation of the model. In figure 17 we plot both the actual inventory holdings along with the implied holdings under the model's decision rules. In the beginning of the sample, 1997-1998, when prices are high, the firm holds almost no inventories. The model recommends that the firm turn away customer and face the lost sale penalty rather than buy when prices are high. Notice the discounted value of the firm's lost sale penalties suggests that the firm's optimally stocks out roughly 50 to 100 days (5-10 percent of the days) during the sample.

During the middle part of the sample period, 1999-2000, the counter-factual inventory holdings track the firm's actual inventory holdings reasonable well. Where the model's inventory policy diverges substantially from the observed policy is at the end of the sample. In December of 2000, the firm had the opportunity to buy steel at 10 cents per pound. The model suggests that the firm should have bought a huge amount (about 3.5 million pounds) of product 2. In fact the firm only bought about 0.5 million pounds. Finally we decomposited the firm's profits for 100 simulations, each of 1013 periods, for the model. We find that conditional on the model being correctly specified, 44 percent of the firm's profits are due the capital gains component while the remaining 56 percent are attributed to the markup. While the model predicts that the firm should aggressively speculate on the price of steel, the returns to price speculation still account for less than half of the firm's profits.

6 Conclusion

In this paper, we motivated our simulated mimimum distance estimator with new database from a single steel wholesaler. While most readers may not be interested in the steel service center industry, the model we presented could be easily modified to apply to any middleman in any asset or commodity market in which the procurement price is stochastic and taken as given. Furthermore, in most markets the only prices recorded are the transaction prices. For example, the optimal trading problem facing a dealer in the market for U.S. Treasury securities may be very similar to one presented here.

References

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

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

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

[4] Chan, H. M. (1999) “Analysis of Price Dispersion in Steel Sales: Evidence of Price Discrimination” manuscript, Yale University.

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

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

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

[8] 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.

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

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

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

[12] Hall, G. and J. Rust (2000) “A (S,s) Model of Commodity Price Speculation” manuscript, Yale University.

[13] 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.

[14] Kaldor, Nicholas (1939) “Speculation and Economic Stability” *Review of Economic Studies* 7, 1-27.

[15] Kalymon, Basil (1971) “Stochastic Prices in a Single-Item Inventory Purchasing Problem” *Operation Research* 19, 1434-1458.

[16] Kingsman, B.G. (1969) “Commodity Purchasing” *Operational Research Quarterly* 20-1, 59-79.

[17] Lee, B. and B.F. Ingram (1991) “Simulation Estimation of Time-Series Models” *Journal of Econometrics* 47197–205.

[18] 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.

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

[20] McFadden, D. (1998) “A Method of Simulated Moments for Estimation of Discrete Models Without Numerical Integration” *Econometrica* 57–5, 995–1026.

[21] 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*

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

[23] Resnick, S. (1992) *Adventures in Stochastic Processes* Birkhäuser, Boston.

[24] 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`

[25] Rust, J. and G. Hall (2003) “Middlemen versus Market Makers: A Theory of Competitive Exchange” *Journal of Political Economy* forthcoming.

[26] Rust, J. Traub, J. and H. Woźniakowski (2002) "Is There a Curse of Dimensionality for Contraction Fixed Points in the Worst Case?" *Econometrica* **70-1** 285–329.

[27] 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.

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

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

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