

TECHNICAL WORKING PAPER SERIES

ECONOMETRIC METHODS FOR ENDOGENOUSLY SAMPLED TIME SERIES:
THE CASE OF COMMODITY PRICE SPECULATION IN THE STEEL MARKET

George Hall
John Rust

Technical Working Paper 278
<http://www.nber.org/papers/T0278>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2002

This paper was originally prepared for the Cowles Foundation Econometrics Conference, October 23-24, 1999. We thank our discussant, Halbert White, and the conference participants including Robert Engle and George Tauchen for helpful comments. We also received helpful feedback from Michael Keane and Kenneth Wolpin and other seminar attendees at subsequent presentations of this paper at the University of Pennsylvania, Boston University, Johns Hopkins, Rice, Penn State, the Universities of Chicago, 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. The views expressed in this paper are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2001 by George Hall and John Rust. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Econometric Methods for Endogenously Sampled Time Series:
The Case of Commodity Price Speculation in the Steel Market
George Hall and John Rust
NBER Technical Working Paper No. 278
August 2002
JEL No. C1, C6, L2

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. However there are no formal exchanges or centralized markets where steel is traded and p_t can be observed. Instead nearly all steel transaction prices are a result of private bilateral negotiations between buyers and sellers, typically intermediated by middlemen known as steel service centers. Even though there is no central record of daily transactions prices in the steel market, we do observe transaction prices for a particular firm -- a steel service center that purchases large quantities of steel in the wholesale market for subsequent resale in the retail market. The endogenous sampling problem arises from the fact that the firm only records p_t on the days that it purchases steel. 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 firm'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 PIML estimator also yields estimates of 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. 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 share of the middleman's discounted profits that are due to markups paid by its retail customers, and the share due to price speculation. The latter measures the firm's success in forecasting steel prices and in timing its purchases in order to "buy low and sell high". The more successful the firm is in speculation (i.e. in strategically timing its purchases), the more serious are the potential biases that would result from failing to account for the endogeneity of the sampling process.

George Hall
Department of Economics
Yale University
P.O. Box 208268
New Haven, CT 06520-8268
and NBER
george.hall@yale.edu

John Rust
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
Tel: 301-405-3489
Fax: 301-405-3542
jrust@gemini.econ.umd.edu
<http://gemini.econ.umd.edu/jrust>

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 $\{t_1, \dots, t_n\}$ 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 are no formal markets or centralized exchanges where steel is traded. Instead nearly all steel transaction prices are a result of private bilateral negotiations between buyers and sellers, typically intermediated by middlemen known as *steel service centers*.¹ Even though there is no central record of daily transactions prices in the steel market, we do observe transaction prices for a particular firm — a steel service center that purchases large quantities of steel in the wholesale market for subsequent resale in the retail market. The endogenous sampling problem arises from the fact that the firm only records p_t on the days that it purchases steel.

We introduce the endogenous sampling problem in the context of price speculation in the steel market in order to provide a concrete example. However we believe that similar endogenous sampling problems arise in many other contexts. Examples include financial applications where transaction prices are observed at randomly spaced intervals (see Aït-Sahalia and Mykland, 2001, Engle and Russell, 1999, and Russell and Engle, 1998), and in marketing applications where the prices of goods that a household purchases are generally only recorded for the items the household purchased and on the dates it purchased them (see Allenby, McCulloch and Rossi 1996, and Erdem and Keane, 1996). However we are not aware of any econometric literature that is directly relevant for handling endogenous sampling problems 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 (Heckman, 1981, Manski and McFadden, 1981, and McFadden, 1997).

We present a parametric analysis of the endogenous sampling problem under the maintained assumption that the timing of steel purchases is part of an optimal trading strategy that maximizes the firm's ex-

¹It 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 occurring via a *market maker* on a centralized exchange, or all trade occurring via decentralized transactions with *middlemen*, or trade segmenting between middlemen and market makers. This theory could explain the variety of different trading institutions that we see in different markets, including the nonexistence of centralized exchanges for steel.

pected 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 can also be viewed as a simulated moments estimator (SME) (Lee and Ingram, 1991 and Duffie and Singleton, 1993), applied to a situation where the data are endogenously sampled. Using the SMD estimator, we estimate the parameters of a truncated lognormal AR(1) model of the wholesale price processes for particular types of steel plate. We use these estimates to infer the share of the firm’s discounted profits that are due to markups paid by its retail customers, and the share due to price speculation. The latter measures the firm’s success in forecasting steel prices and in timing its purchases in order to “buy low and sell high”. The more successful the firm is in speculation (i.e. in strategically timing its purchases), the more serious are the potential biases that would result from failing to account for the endogeneity of the sampling process.

This paper originated from 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 steel 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:

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

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 firm 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 the retail price may not provide information that is directly relevant for estimating the unknown parameters of the wholesale price process.²

The estimation methods we propose requires nested numerical solution of a dynamic programming problem that determine the firm’s optimal trading strategy. This must be done for each trial value for the unknown parameter vector θ , and as a result, the estimators we propose are computationally intensive. However significant computational savings can be achieved by exploiting special features of the solutions to these dynamic programming problems. Extending a seminal result by Scarf (1959) for a simpler class of inventory investment problems, Hall and Rust (2001) showed 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*. In a generalized (S, s) rule, S and s) are functions of the current wholesale price p and a vector of other state variables x such as interest rates, demand shifters, and other variables that affect the firm’s beliefs about future prices and sales levels. The functions $S(p, x)$ and $s(p, x)$ satisfy $S(p, x) \geq s(p, x)$. 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 $S(p, x)$.

²Our 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 firm and evolves as an exogenous stochastic process (i.e. its realizations do not depend on actions of the firm) 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. The main modification is that when there is no spot wholesale market and the “law of one price” does not hold, we would need to estimate a conditional probability distribution representing the firm’s beliefs about the distribution of potential prices available at a given point in time.

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

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

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

$$q_t^o = S(p_t, x_t) - q_t, \quad (2)$$

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

The main idea behind the likelihood based approach to solving the endogenous sampling problem 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 recursive quadrature, or by 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 and thus avoids the need for high dimensional integrations. The SMD estimator only relies on the ability to simulate realizations of the optimal trading model. These simulations are then censored in exactly the same way as the observed data are censored, an approach that is similar in many respects to the strategy of “data augmentation” used in Bayesian inference of latent variable models. The idea behind the SMD estimator is to choose parameter values that result in simulated moments that match the observed moments as closely as possible, where both the real and simulated data are censored according to the same sampling rule; namely the one given in equation (1). Even though the moments entering the SMD criterion are biased and inconsistent due to the endogenous sampling problem, the fact that we can censor the data entering the

simulated and real moments in the same way implies that the SMD estimator itself is consistent. 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 variety of other contexts.

Section 2 describes our data set and introduces the steel speculation and inventory problem that motivates 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 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, $\xi_t = (q_t, q_t^o, p_t, p_t^r, x_t)$. 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 as well as 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 firm’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 firm’s discounted profits are due to the markups it charges its retail customers, and the fraction that is due to pure commodity price speculation, i.e., its success in timing purchases of steel in order to profit from “buying low and selling high.”

2 Description of the Data and the Model of Price Speculation

In this section we introduce the data and describe a generalized version of a model of commodity price speculation introduced by Hall and Rust (1999, 2000, 2001) that allows for additional covariates and unobserved state variables. This model provides the framework for inference and provides the key insights that enabled us to pose and solve the endogenous sampling problem.

2.1 The Data

Via a personal contact with an executive at a large U.S. steel wholesaler, we acquired a new high frequency micro database on transactions in the steel market. This firm has provided us with an ongoing data feed that enables us to observe 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 5 are based on data on every transaction the firm made between July 1, 1997 to March 14, 2002 (1191 business days) for two of its highest volume steel products. For each transaction we observe the quantity (number of units and/or weight in pounds) of steel bought or sold, the sales price, the shipping costs, and the identity of the buyer or seller.

Although this is an exceptionally clean and rich dataset, we only observe prices on the days the firm actually made transactions: the firm does not record any price information on days that it does not transact (either as a buyer or seller of steel). This shortcoming of our dataset is much more important for steel purchases than steel sales, since 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 we present below predicts that purchases are not made at random. Instead, the firm tends to make purchases when prices are low, so that the average price on the days the firm makes purchases will be lower than the average wholesale price on days the firm does not purchase. The exception to this general rule is that the firm may make purchases even when prices are relatively high if its inventories are low. Conversely, the firm may refrain from purchasing even if prices and inventories are low if it expects that the rate of retail sales will be depressed for a long period of time, say due to bad macroeconomics conditions. Thus, while the firm is attempting to “buy low and sell high”, its purchase decisions involve a tradeoff among a number of different considerations.

We illustrate our data by plotting the time series of inventories and prices of one of the firm’s products in figures 1 and 2. This product, which we call product 4, is one of highest volume products sold by this

firm. It is also a benchmark product within the industry since the prices of several other steel products are often computed as a function of this product's price. It is possible to get weekly and monthly survey data on prices for certain classes of steel products through trade publications such as *Purchasing Magazine* and *American Metal Market*. However, since there are no public exchange markets for steel products, transaction in the steel market are carried out in private negotiations. Hence these price surveys rely on participants in the steel market to report truthfully the prices they paid or received for various steel products. The firm often faces considerably different prices than those in the survey data.

As a result, in our plots of wholesale transaction prices in figure 2 (the lower curve with the large black circles), we used straight line interpolations between observed purchase prices at successive purchase dates. The black circle at each purchase date is proportional to the size of the firm's purchase in pounds. This gives us our first visual indication of the endogenous sampling problem. First, we see that even though we have 1191 observations on this firm, we observe purchases in the wholesale market on only 184 days. Second, the patterns of the black dots suggests that the firm is more likely to purchase large quantities of steel when wholesale prices are low, although other economic factors seem to be influencing the firm's purchase decisions as well. One key factor is the level of inventory: the firm tends to make large purchases when its inventory is low. We also see that even though wholesale prices continued to decline during 2000 and 2001, the firm's largest purchases of steel occurred during the "turning point" in prices in early 1998. The firm may have avoided making large purchases in late 2000 and 2001 due to economic uncertainties resulting from the "dot com crash" and the economic uncertainties following the 9/11/2001 terrorist attack on the U.S.

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

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

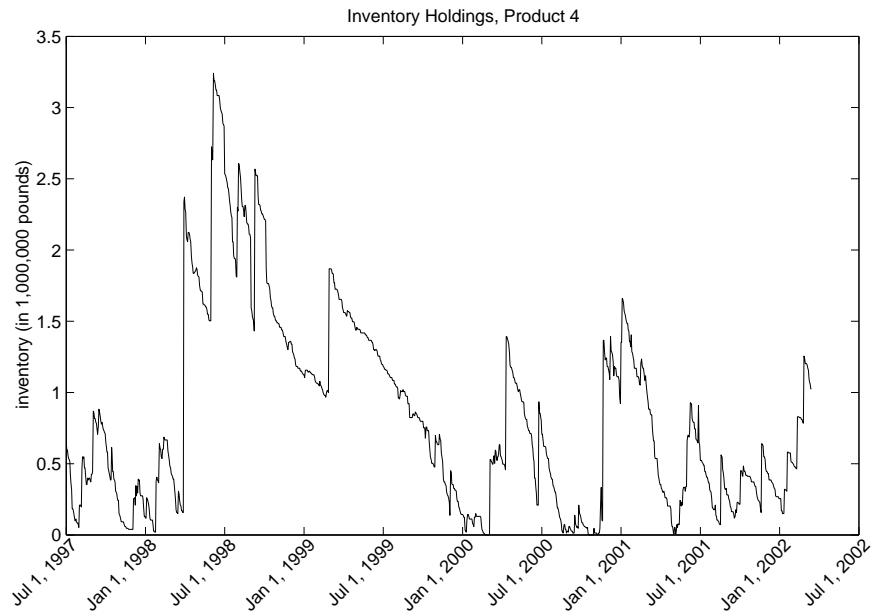


Figure 1: Times series plot of the inventory for product 4.

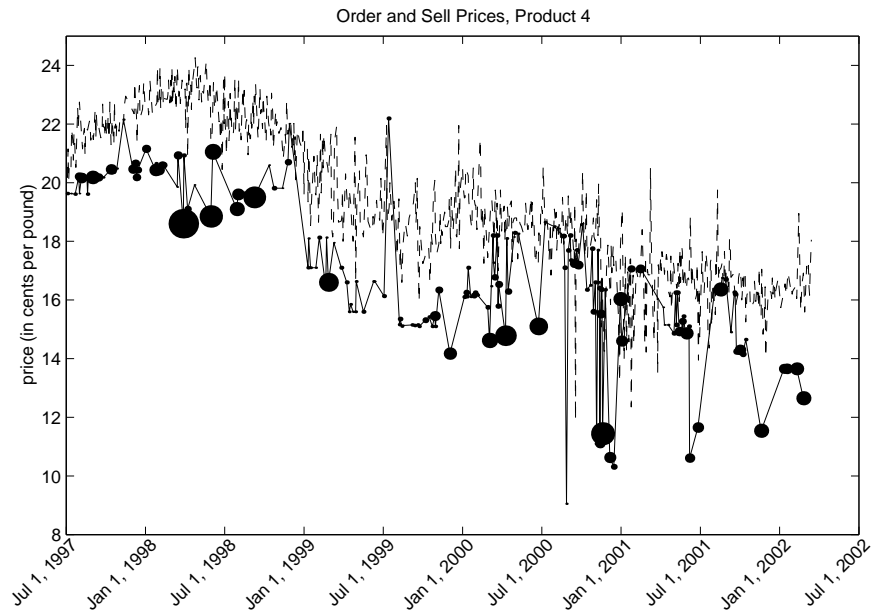


Figure 2: Purchase prices (solid line) and retail prices (dashed line) for product 4. For the purchase price series, the size of the marker is proportional to the size of the purchase.

the high frequency variation is due to observable factors. Athreya (2002) finds that roughly 65% of the high frequency variation in retail prices can be explained by observable customer characteristics such as geographical location and past volume of purchases. The remaining 35% of the variation in retail prices appears to be due either to high frequency fluctuations in wholesale prices or to some sort of “informational price discrimination” in the retail market. Using the limited number of days on which both wholesale and retail prices are available, Chan (2001) finds that at most 50% of the variation in retail prices can be explained by variations in the wholesale price of steel. This conclusion is possible due to the fact that on many days there are multiple retail sales to different customers. These findings suggest that a large share of the high frequency variation in retail prices can be ascribed to price discrimination, i.e. the firm charges higher prices to more impatient or poorly informed retail customers (see Chan, Hall and Rust (2003) for a more detailed analysis of bargaining, price setting, and price discrimination in the retail market for steel). We conclude that even though retail sales occur much more frequently than wholesale purchases, the fact that retail prices involve a number of other different considerations (including price discrimination based on observable and unobservable characteristics of the customer) suggest that the retail price is at best a very noisy and (upward) biased signal of the underlying wholesale price.

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

2.2 The Model

Our model is an extension of previous work by Hall and Rust (2001), who showed that in a broad class of commodity price speculation problems, the optimal trading rule is a generalized version of the classic (S, s) rule from inventory theory. Their work can be viewed as linking contributions by Arrow *et. al.* (1951) and Scarf (1959) who first proved the optimality of (S, s) policies in inventory investment problems to more recent work by Williams and Wright (1991), Deaton and Laroque (1992) and Miranda and Rui (1997) on the rational expectations commodity storage model. The fixed (S, s) thresholds derived by Scarf under the assumption that the price (cost) of procuring (producing) inventories is constant are clearly suboptimal in a speculative trading environment, since the stochastic fluctuations in the price of steel affects the firm's perception of the optimal level of inventory S , and the threshold for purchasing new inventory s . Hall and Rust (2001) showed that the firm's optimal speculative trading strategy is a *generalized the (S, s) rule* where S and s are functions of certain underlying state variables including the wholesale price of steel p .³

Before we describe how the generalized (S, s) rule allows us to formulate and solve the problem of endogenous sampling of steel wholesale prices, we describe the notation and key assumptions underlying Hall and Rust's model of commodity price speculation. Then we formally define the (S, s) trading strategy, and show how in a broad class of models of speculation, the (S, s) rule constitutes the optimal strategy for "buying low and selling high". We assume that a middleman (which we also refer to as the "firm") can purchase unlimited quantities of steel at a time-varying wholesale price p_t that evolves according to a Markov transition density to be specified 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 wholesale price p_t (if we think of the firm as selling to different customers on different business days, this

³This analysis extends previous results in the operations research literature such as Fabian *et. al.* (1959), Kingman (1969), Kalyon (1971), Golabi (1985), Song and Zipken (1993), Moinszadeh (1997), and Ozekici and Parlar (1999) that prove the optimality of generalized versions of the (S, s) rule when the cost (price) of producing (procuring) new inventory fluctuates stochastically. While Hall and Rust (2001) are not the first to prove the optimality of generalized versions of the (S, s) rule, they build on the OR literature by making the connection between models of optimal inventory policies and models of storage and commodity prices. Moreover in the current paper we computationally solve and estimate our model. Thus we can formally compare the model's optimal policies to the inventory policies we see in the data. Besides the work noted above, the most closely related recent work that we are aware of is the ambitious paper by Aguirregabiria (1999) that models price and inventory decisions by a supermarket chain. A supermarket is similar to our steel wholesaler in that both types of firms hold inventories of a substantial number of different products, purchasing them in the wholesale market and selling their inventories at a markup to retail customers. The key difference is that prices in supermarkets are almost always posted so there is no direct price discrimination and there is presumably a larger "menu cost" to changing prices on a day by day basis. Aguirregabiria also did not directly address the endogenous sampling issue, using monthly price averages as proxies for underlying daily prices. For this reason we are unable to directly employ his innovative and ambitious approach to estimation.

randomly varying markup is intended to be a “reduced-form” approach to capturing the pricing and price discrimination decisions by the firm).

On 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 wholesale 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. 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 $\gamma(p_t^r | q_t + q_t^o, p_t, x_t)$.
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 $H(q_t^r | p_t, p_t^r, x_t)$ 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

$$q_t^s = \min[q_t + q_t^o, q_t^r]. \quad (3)$$

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:

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

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

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:

$$\delta(q, p, p^r, x) = 1 - H(q | p^r, p, x). \quad (5)$$

Since retail sales occur much more frequently than purchases of new inventory, the retail sales price p_t^r provides an important source of information about the wholesale price p_t . Presumably for most transactions we should have $p_t^r \geq p_t$, reflecting nonnegative markups over the current wholesale price of steel.

However as noted above markups vary 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 $\gamma(p_t^r|q_t + q_t^o, p_t, x_t)$. This way of modeling prices is sufficiently flexible to be consistent with a variety of theories of bargaining and price discrimination by the firm.⁴

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

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

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

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

In order to state the per period profit function, we need to describe the costs that the firm incurs. The main cost is the cost of ordering new inventory, represented by the *order cost function* $c^o(q^o, p)$. We assume that the firm incurs a fixed cost $K \geq 0$ associated with placing new orders for inventory, which implies that $c^o(q^o, p)$ is given by

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

⁴Hall and Rust (2000) solved 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 (2000) model requires a two-dimensional optimization instead of an one-dimensional optimization at each iteration of the Bellman equation. Second, in models with endogenous price setting, the generalized (S, s) rule is not always guaranteed to be an optimal trading strategy. As a result we cannot restrict our search to the subclass of generalized (S, s) policies as we can when we solve the model presented here. This greatly increases the computational time required to solve models that incorporate either endogenous (uniform) price setting (as in Hall and Rust 2000), or in models of bargaining and price discrimination (as in Chan, Hall and Rust, 2003).

The firm's remaining costs are summarized by the *holding cost function* $c^h(q, p, x)$. These costs include physical storage costs, and “goodwill costs” representing the present value of lost future business from customers whose orders cannot be filled due to a stockout. Goodwill costs can be viewed as the inverse of the “convenience yield” discussed in the commodity storage literature (Kaldor, 1939, Williams and Wright, 1991). In this case a convenience yield emerges from a desire to hold a buffer stock or precautionary level of inventories in order to minimize goodwill costs from stockouts. This allows the model to capture other reasons besides pure price speculation for holding inventories.⁵ The firm's single-period profits π equals its sales revenues, less the cost of new orders for inventory $c^o(q^o, p)$ and inventory holding costs $c^h(q + q^o, p, x)$:

$$\pi(p, p^r, q^r, q + q^o, x) = p^r q^s - c^o(q^o, p) - c^h(q + q^o, p, x). \quad (9)$$

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

$$V(p_t, q_t, x_t) = \max_{q_t^o} E \left\{ \sum_{j=t}^{\infty} \rho^{(j-t)} \pi(p_j, p_j^r, q_j^r, q_j^o + q_j, x_j) \middle| p_t, q_t, x_t \right\}, \quad (10)$$

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

$$V(p, q, x) = \max_{0 \leq q^o \leq \bar{q} - q} \left[W(p, q + q^o, x) - c^o(q^o, p) \right], \quad (11)$$

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

$$W(p, q, x) \equiv \left[ES(p, q, x) - c^h(q, p, x) + \rho EV(p, q, x) \right], \quad (12)$$

and EV denotes the conditional expectation of V given by:

$$\begin{aligned} EV(p, q, x) &= E\{V(\tilde{p}, \max[0, q - \tilde{q}^r], \tilde{x}) | p, q, x\} \\ &= \lambda_1(p, q, x) \int_{p'} \int_{x'} V(p', q, x') g(p', x' | p, x) dp' dx' \\ &+ \lambda_2(p, q, x) \int_{p'} \int_{x'} V(p', 0, x') g(p', x' | p, x) dp' dx' \\ &+ \lambda_3(p, q, x) \int_{p'} \int_{x'} \int_0^q V(p', q - q', x') h(q' | p, q, x) g(p', x' | p, x) dq' dp' dx', \end{aligned} \quad (13)$$

⁵The firm obtains much of its steel from foreign sources. In the model orders occur instantaneously with certainty. In practice, however, delivery lags can be several months and the steel delivered can often be of lower quality than agreed on. The firm does have the option of refusing to take delivery if the steel is not of the quality promised. Having a buffer stock of inventories on hand reduces the cost to firm of exercising this option. Also foreign producers of steel do from time to time renege on previously negotiated deals, failing to deliver the amount of steel originally promised.

where

$$\begin{aligned}
\lambda_1(p, q, x) &= \int_{p^r} \eta(p^r, p, x) \gamma(p^r | p, q, x) dp^r \\
\lambda_2(p, q, x) &= \int_{p^r} [1 - \eta(p^r, p, x)] \delta(p^r, p, q, x) \gamma(p^r | p, q, x) dp^r \\
\lambda_3(p, q, x) &= \int_{p^r} [1 - \eta(p^r, p, x)] \gamma(p^r | p, q, x) dp^r \\
h(q' | p, q, x) &= \int_{p^r} h(q' | p^r, p, q, x) \gamma(p^r | p, q, x) dp^r.
\end{aligned} \tag{14}$$

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

$$q^o(p, q, x) = \inf_{0 \leq q^o \leq \bar{q} - q} \operatorname{argmax} \left[W(p, q + q^o, x) - c^o(q^o, p) \right]. \tag{15}$$

We invoke the inf operator in the definition of the optimal decision rule in equation (15) to handle the case where there are multiple maximizing values of q^o . We effectively break the tie in such cases by defining $q^o(p, q)$ as the *smallest* of the optimizing values of q^o .

In this model the variables q and q^o do not enter as separate arguments in the value function W given in (12): rather they enter as the sum $q + q^o$ as shown in equation (15). 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:

$$\frac{\partial W}{\partial q^o}(p, q + q^o, x) = p. \tag{16}$$

If W were strictly concave in q , there would be a unique value of $q + q^o$ that solves equation (16) for any value of p . Call this solution $S(p, x)$:

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

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

It turns out that if $K > 0$ the function $W(p, q, x)$ will not be strictly concave. However under fairly

general conditions W is K -concave as a function of q for each fixed p .⁶ Using the K -concavity property we can prove that whenever $q \geq s(p, x)$, it is not optimal to order: $q^o(p, q, x) = 0$. When $q < s(p, x)$ the symmetry property implies that $q^o(p, q, x) = S(p, x) - q$ as discussed above. In particular Hall and Rust (2001) proved:

Theorem 1: Consider the function $W(p, q + q^o, x)$ defined in equation (12), where W is defined in terms of the unique solution V to Bellman's equation (11). Under appropriate regularity conditions given in Hall and Rust (2001), the optimal speculative trading strategy $q^o(p, q, x)$ takes the form of an (S, s) rule. That is, there exist a pair of functions (S, s) satisfying $S(p, x) \geq s(p, x)$ where $S(p, x)$ is the desired or target inventory level and $s(p, x)$ is the inventory order threshold, i.e.

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

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

$$S(p, x) = \operatorname{argmax}_{0 \leq q^o \leq \bar{q} - q} [W(p, q^o, x) - c^o(q^o, p)] \quad (19)$$

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:

$$s(p, x) = \inf_{q \geq 0} \{qW(p, q, x) - pq \geq W(p, S(p, x), x) - pS(p, x) - K\}. \quad (20)$$

By a simple substitution of the generalized (S, s) rule in equation (18) into the definition of V in equation (11) we obtain the following corollaries:

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

$$V(p, q, x) = \begin{cases} W(p, S(p, x), x) - p[S(p, x) - q] - K & \text{if } q \in [0, s(p, x)] \\ W(p, q, x) & \text{if } q \in (s(p, x), \bar{q}]. \end{cases} \quad (21)$$

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

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

$$S(p, x) = s(p, x). \quad (22)$$

⁶A function $W(p, q) : [p, \bar{p}] \times [0, \bar{q}] \rightarrow R$ 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 $p \in [p, \bar{p}]$, and for all $z \geq 0$ and $b \geq 0$ such that $q + z \leq \bar{q}$ and $q - b \geq 0$ we have $W(p, q + z) - K \leq W(p, q) + z[W(p, q) - W(p, q - b)]/b$.

3 Maximum Likelihood Estimation

This section derives the likelihood function 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 yield a first order condition that could possibly provide a basis for a generalized methods of moments (GMM) strategy for estimating the unknown parameters of the model:

$$\frac{\partial W}{\partial q}(p, S(p, x), x) - p = 0. \quad (23)$$

If we assume that there is additive measurement error ε in the wholesale price p , or assume that ε represents other unobserved (per unit) components of the cost of ordering new inventory, then it is tempting to treat equation (23) as an ‘‘Euler equation’’ and use GMM to estimate parameters of the model. However there are several big obstacles to this approach. First, we do not have a convenient analytical formula for the partial derivative of the value function, $\partial W / \partial q$. Second, 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 $q < s(p, x)$), is generally nonzero. Finally, there is the issue of endogenous sampling, and the fact that we observe purchases only on a relatively small subset of business days in our overall sample.

These problems motivate a search for an alternative likelihood-based approach that is capable of incorporating other information such as retail sales prices in order to improve our ability to make inferences about the $\{p_t\}$ process. 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 reflects 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 whole price process. It does this by integrating out the unobserved values of the wholesale prices in periods where they are unobserved. This likelihood is the analog of the Chapman-Kolmogorov equation for computing multi-step transition probabilities from a one-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. Without some sort of ‘‘error term’’ the model yields a de-

terministic optimal decision rule $q^o(p, q, x)$ that can be contradicted by any observation (q_t^o, q_t, p_t, x_t) that does not lie on its graph. To avoid the resulting “zero likelihood” problem, consider the case where there is an unobserved component of the per unit cost of steel, denoted by ε_t . 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, x) formed by integrating out ε from the deterministic decision rule $q^o(p, q, x, \varepsilon)$.

Theorem 2: *Let ε_t be an (unobserved to the econometrician) component of the per unit cost of ordering new inventory. Assume that $\{\varepsilon_t\}$ is an IID process whose density ϕ is continuous and strictly positive over the entire real line. Then the optimal trading strategy is still a generalized (S, s) rule and the conditional distribution of the optimal order quantity q^o given (p, q, x) is given by*

$$\begin{aligned}
F(q^o | p, q, x) &= \Pr\{q^o(p, q, x, \varepsilon) \leq q^o | p, q, x\} \\
&= \int_{-\infty}^{+\infty} I\{q^o(p, q, x, \varepsilon) \leq q^o\} \phi(\varepsilon) d\varepsilon \\
&= \int_{s^{-1}(p, x, q)}^{\infty} \phi(\varepsilon) d\varepsilon \\
&+ I\{S(p, x, s^{-1}(p, x, q)) \leq q^o + q \leq \bar{q}\} \int_{S^{-1}(p, x, q+q^o)}^{s^{-1}(p, x, q)} \phi(\varepsilon) d\varepsilon \\
&+ I\{q^o + q > \bar{q}\} \int_{-\infty}^{s^{-1}(p, x, \bar{q})} \phi(\varepsilon) d\varepsilon,
\end{aligned} \tag{24}$$

where

$$\begin{aligned}
S^{-1}(p, x, q) &= \inf\{\varepsilon | S(p, x, \varepsilon) = q\} \\
s^{-1}(p, x, q) &= \inf\{\varepsilon | s(p, x, \varepsilon) = q\}.
\end{aligned} \tag{25}$$

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

$$f(q^o | p, q, x) = \begin{cases} \int_{s^{-1}(p, x, q)}^{\infty} \phi(\varepsilon) d\varepsilon & \text{if } q^o = 0 \\ \int_{-\infty}^{S^{-1}(p, x, \bar{q})} \phi(\varepsilon) d\varepsilon & \text{if } q^o = \bar{q} - q \\ \frac{-\phi(S^{-1}(p, x, q+q^o))}{\partial^2 W / \partial^2 q(p, x, q+q^o)} & \text{otherwise.} \end{cases} \tag{26}$$

The formula for the density of q^o in equation (26) can be derived by differentiating the conditional distribution in equation (24) with respect to q^o for q^o in the interval $[S(p, x, s^{-1}(p, x, q)) - q, \bar{q} - q]$ to obtain:

$$dF(q^o | p, q, x) = -\phi(S^{-1}(p, x, q+q^o)) \frac{\partial S^{-1}}{\partial q^o}(p, x, q+q^o). \tag{27}$$

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

$$\frac{\partial W}{\partial q}(p, S(p, x, \varepsilon), x) = p + \varepsilon, \quad (28)$$

and the inverse and implicit function theorems we obtain:

$$\frac{\partial S^{-1}}{\partial q^o}(p, x, q + q^o) = \frac{1}{\partial S(p, x, S^{-1}(p, x, q + q^o))/\partial \varepsilon} = \frac{1}{\partial^2 W(p, q + q^o, x)/\partial^2 q}. \quad (29)$$

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 $q^o = \bar{q} - q$, and strictly positive density over the interval $[S(p, x, s^{-1}(p, x, q)) - q, \bar{q} - q]$. However there is a “gap” where there is zero density for q^o in the interval $[0, S(p, x, s^{-1}(p, x, q)) - q]$ since the quantity $S(p, x, s^{-1}(p, x, q)) - q$ represents the minimum order size implied by the (S, s) model in the state (p, q, x) . 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. To obtain a fully nondegenerate likelihood function, we would have to introduce a second unobservable, such as an unobservable component υ of the fixed cost K of placing an order. If the distribution of this component is such that there is positive probability that the combined order cost $K + \upsilon$ is arbitrarily close to zero for sufficiently small realizations of υ , then consistent with Corollary 3 of section 2, the gap will be zero, thus eliminating the possibility of a “zero likelihood problem.” In practice for the values of K we encountered in our estimation, the gap is sufficiently small that zero likelihood problems did not arise. Therefore in order to simplify the the model and the exposition we decided to omit the case where there are unobservable components of K as well as p .

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

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

$$\mu(q^l|p, p^r, q, q^o, x) = \begin{cases} (1 - \eta(p^r, p, x))\delta(p^r, p, q + q^o, x) & \text{if } q^l = 0 \\ \eta(p^r, p, x) & \text{if } q^l = q + q^o \\ (1 - \eta(p^r, p, x))h(q + q^o - q^l|p^r, p, x) & \text{otherwise} \end{cases} \quad (30)$$

Under our setup, we can show that the observables $\{p_t, p_t^r, q_t, q_t^o, x_t\}$ evolve as a joint Markov process which also has a discrete/continuous transition probability density λ . We state this as Theorem 4:

Theorem 4: *The joint process $\{p_t, p_t^r, q_t, q_t^o, x_t\}$ is Markov with (discrete/continuous) transition density λ given by:*

$$\begin{aligned} \lambda(p_{t+1}, p_{t+1}^r, q_{t+1}, q_{t+1}^o, x_{t+1} | p_t, p_t^r, q_t, q_t^o, x_t) &= g(p_{t+1}, x_{t+1} | p_t, x_t) \\ &\times \mu(q_{t+1} | p_t, p_t^r, q_t, q_t^o, x_t) \\ &\times f(q_{t+1}^o | p_{t+1}, q_{t+1}, x_{t+1}) \\ &\times \gamma(p_{t+1}^r | p_{t+1}, q_{t+1} + q_{t+1}^o, x_{t+1}). \end{aligned} \quad (31)$$

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

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

$$\hat{\theta}_T^f = \underset{\theta \in \Theta}{\operatorname{argmax}} l_f(\{p_t, p_t^r, q_t, q_t^o, x_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, x_0, \theta), \quad (32)$$

where l_f is given by:

$$l_f(\{p_t, p_t^r, q_t, q_t^o, x_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, x_0, \theta) = \prod_{t=1}^T \lambda(p_t, p_t^r, q_t, q_t^o, x_t | p_{t-1}, p_{t-1}^r, q_{t-1}, q_{t-1}^o, x_{t-1}, \theta). \quad (33)$$

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

Now consider the partial information case where we only observe wholesale prices on the subset of n trading days, $T_n \equiv \{t_1, \dots, t_n\}$ at which purchases occur. To simplify notation we assume (without loss of generality) that the data 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 (33) over wholesale 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:

$$p_t = \begin{cases} p_t^* & \text{if } q_t^o > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (34)$$

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

$$\hat{\theta}_p = \underset{\theta \in \Theta}{\operatorname{argmax}} l_p(\{p_t, p_t^r, q_t, q_t^o, x_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, x_0, \theta), \quad (35)$$

where l_p is given by:

$$l_p(\{p_t, p_t^r, q_t, q_t^o, x_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, x_0, \theta) = \int_{p_{T_n}} \int_{\varepsilon_{T_n}} \cdots \int_{p_{T_1}} \int_{\varepsilon_{T_1}} l_f \prod_{t \notin T_n} I\{q_t > s(p_t, x_t, \varepsilon_t)\} q(\varepsilon_t) dp_t d\varepsilon_t. \quad (36)$$

Thus, the PIML likelihood l_p is derived from the FIML likelihood l_f by integrating out the unobserved wholesale prices over the dates $t \notin T_n$ that purchases do not occur. The region of integration is limited to the region of the state space where making a purchase is not optimal. This is given by the indicator function $I\{1_t > s(p_t, x_t, \varepsilon_t)\}$. Notice that this region involves the unobserved state variable ε_t . Thus the integration must be done over both unobserved variables (p_t, ε_t) over all of the $T - n$ dates $t \notin T_n$ at which purchases do not occur.

We will now sketch the asymptotic properties of the PIML estimator under the assumption that there is only one firm, but $T \rightarrow \infty$. The asymptotic properties of the FIML estimator are well known: the logarithm of l_f can be approximated as a (normalized) sum of random variables. Despite the correlation in these random variables in successive time periods, standard limit theorems for ergodic processes can be used to show that this normalized sum converges to a well defined score function. A standard ‘‘information inequality’’ argument can then be used to show that this score function is maximized at the true parameter value θ^* , assuming that the model is correctly specified. A formal proof would require specification of regularity conditions similar to Billingsley (1961) and White (1982) to ensure that the convergence of these normalized sums to the score function is uniform and that the score function is uniquely maximized at θ^* . These are standard sufficient conditions for the consistency of maximum likelihood.

However the argument for the consistency of the PIML estimator is more complicated. The high-dimensional integrations over the irregularly spaced intervals between successive purchases create linkages between the observations in the PIML estimator. When we take the logarithm of the likelihood it no longer decomposes into a normalized sum of T random variables as in the FIML case. Thus the standard arguments used to prove the consistency and asymptotic normality in the FIML case do not appear to apply

in the PIML case. At best, the logarithm of the PIML likelihood decomposes into a sum of n terms, where each term is the logarithm of a high dimensional integral of the transition probability density λ over the times between successive purchases. However if the joint process $\{p_t, p_t^r, q_t, q_t^o, x_t\}$ is ergodic, we should have $n \rightarrow \infty$ with probability 1 as $T \rightarrow \infty$. Our strategy will be to do the asymptotics for the PIML estimator as a function of the number of purchases n rather than as a function of the number of time periods T over which the firm is observed. In order to derive the asymptotic properties of the PIML estimator, we will use the fact that the state of the process at successive purchase dates is an *embedded Markov chain* and the sequence of realized states between successive purchases forms a *segmented Markov chain*. We will then argue that the segmented Markov chain is ergodic, which will allow us to apply the relevant limit theorems to establish the asymptotic properties of the PIML estimator.

Let $\{\xi_t\}$ denote the joint Markov process in theorem 4, i.e., the process whose value at t is given by:

$$\xi_t \equiv (p_t, p_t^r, q_t, q_t^o, x_t). \quad (37)$$

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

$$\Gamma = \{(\xi, \varepsilon) \mid q^o = 0\} = \{(\xi, \varepsilon) \mid q > s(p, x, \varepsilon)\}, \quad (38)$$

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

$$t_{i+1} = \inf\{t > t_i \mid \xi_t \in \Gamma\}. \quad (39)$$

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.,*

$$\{\zeta_i\} = \{\xi_{t_i}\}. \quad (40)$$

We 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 1: *The embedded process $\{\zeta_i\}$ is a Markov chain with transition density ν_e given by:*

$$\nu_e(\zeta_i \mid \zeta_{i-1}, \theta) = \lambda(\xi_{t_i} \mid \xi_{t_{i-1}}, \theta) = \int_{\xi_{t_{i-1}+1}} \int_{\varepsilon_{t_{i-1}+1}} \dots \int_{\xi_{t_{i-1}}} \int_{\varepsilon_{t_{i-1}}} \prod_{t=t_{i-1}+1}^{t=t_i-1} I\{(\xi_t, \varepsilon_t) \notin \Gamma\} \lambda(\xi_t \mid \xi_{t-1}, \theta) d\xi_t d\varepsilon_t. \quad (41)$$

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 (observed) values of $\{\xi_t\}$ for the sequence of $t_i - t_{i-1}$ time periods following the purchase at t_{i-1} until the purchase at t_i :*

$$\omega_i = (\xi_{t_{i-1}+1}, \dots, \xi_{t_i}). \quad (42)$$

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 2: *The segmented process $\{\omega_i\}$ is a Markov chain with transition density v_s given by:*

$$v_s(\omega_i | \omega_{i-1}, \theta) = \int_{p_{t_i+1}} \int_{\varepsilon_{t_i+1}} \cdots \int_{p_{t_{i+1}-1}} \int_{\varepsilon_{t_{i+1}-1}} \prod_{t=t_i+1}^{t_{i+1}} \lambda(\xi_t | \xi_{t-1}, \theta) \prod_{t=t_i+1}^{t_{i+1}-1} I\{q_t > s(p_t, x_t, \varepsilon_t)\} q(\varepsilon_t) dp_t d\varepsilon_t. \quad (43)$$

Thus, the transition density for the segmented chain $\{\omega_i\}$ is basically the product of the transition densities for the uncensored $\{\xi_t\}$ process between successive purchases at periods t_i and t_{i+1} , $\prod_{t=t_i+1}^{t_{i+1}} \lambda(\xi_t | \xi_{t-1}, \theta)$, but integrated over the region of (p_t, ε_t) space between time periods $t_i + 1, \dots, t_{i+1} - 1$ when purchases are not observed. The appropriate region of integration is defined by the product of the indicator functions $I\{q_t > s(p_t, x_t, \varepsilon_t)\}$ that specify that the relevant price paths are those for which inventories lie above the $s(p, x, \varepsilon)$ band, so that it is not optimal to purchase during this time interval.

Notice that due to the Markov property for $\{x_t\}$, only the last element of the segment ω_{i-1} , $\xi_{t_{i-1}}$, is needed to fully determine the conditional probability of $\omega_i = (\xi_{t_{i-1}+1}, \dots, \xi_{t_i})$. 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 $T \rightarrow \infty$.

Assumption 1: *The Markov chain $\{\xi_t\}$ is ergodic (i.e. it possesses a unique stationary distribution), the purchase set Γ is recurrent (i.e. $E\{\tau\} < \infty$), and the embedded and segmented processes $\{\zeta_i\}$ and $\{\omega_i\}$ are ergodic Markov chains.*

To study the asymptotic properties of the PIML estimator, it is useful to rewrite the likelihood function l_p as a product of $n - 1$ terms, each of which describes the integrated likelihood between the n purchase dates:

$$l_p(\{p_t, p_t^r, q_t, q_t^o, x_t\}_{t=1}^T | p_0, p_0^r, q_0, q_0^o, x_0, \theta) = \prod_{i=1}^{n-1} v_s(\omega_{i+1} | \omega_i, \theta). \quad (44)$$

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

$$\lim_{T \rightarrow \infty} \frac{n}{T} = \frac{1}{E\{\tau\}}. \quad (45)$$

Thus, as long as $E\{\tau\} < \infty$, the process $\{\xi_t\}$ visits Γ infinitely often and $n \rightarrow \infty$ with probability 1 as $T \rightarrow \infty$. 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 PIML estimator, it is convenient to work with the normalized log-likelihood functions. First, we multiply and divide the likelihood l_p by a product of the conditional densities of the recurrence times $\prod_{i=1}^{n-1} Pr\{\tau = t_{i+1} - t_i | \xi_i, \theta\}$ where

$$Pr\{\tau = t_{i+1} - t_i | \xi_i, \theta\} = \int_{\xi_{t_{i+1}}} \int_{\varepsilon_{t_{i+1}}} I\{q_{t_{i+1}} \leq s(p_{t_{i+1}}, x_{t_{i+1}}, \varepsilon_{t_{i+1}})\} \lambda(\xi_{t_{i+1}} | \xi_{t_{i+1}-1}, \theta) \times \quad (46)$$

$$\left[\int_{\xi_{t_{i+1}}} \int_{\varepsilon_{t_{i+1}}} \cdots \int_{\xi_{t_{i+1}-1}} \int_{\varepsilon_{t_{i+1}-1}} \prod_{t=t_i+1}^{t_{i+1}-1} I\{q_t > s(p_t, x_t, \varepsilon_t)\} \lambda(\xi_t | \xi_{t-1}, \theta) d\xi_t d\varepsilon_t \right] d\xi_{t_{i+1}} d\varepsilon_{t_{i+1}}.$$

Taking logs and dividing by $n - 1$ we obtain the following form for the normalized log-likelihood function

$$\begin{aligned} & \frac{1}{n-1} \log l_p(\{p_t, p'_t, q_t, q'_t, x_t\}_{t=1}^T | p_0, p'_0, q_0, q'_0, x_0, \theta) \\ &= \frac{1}{n-1} \sum_{i=1}^{n-1} \log \rho(\omega_{i+1} | \omega_i, t_{i+1} - t_i, \theta) + \frac{1}{n-1} \sum_{i=1}^{n-1} \log Pr\{t_{i+1} - t_i | \omega_i, \theta\} \\ &= \frac{1}{n-1} \sum_{i=1}^{n-1} v_1(\omega_{i+1}, \omega_i, \theta) + \frac{1}{n-1} \sum_{i=1}^{n-1} v_2(\omega_{i+1}, \omega_i, \theta), \end{aligned} \quad (47)$$

where

$$\begin{aligned} \rho(\omega_{i+1} | \omega_i, t_{i+1} - t_i, \theta) &= \quad (48) \\ & \int_{p_{t_{i+1}}} \int_{\varepsilon_{t_{i+1}}} \cdots \int_{p_{t_{i+1}-1}} \int_{\varepsilon_{t_{i+1}-1}} \prod_{t=t_i+1}^{t_{i+1}} \frac{\lambda(\xi_t | \xi_{t-1}, \theta)}{Pr\{t_{i+1} - t_i | \xi_i, \theta\}} \prod_{t=t_i+1}^{t_{i+1}-1} I\{q_t > s(p_t, x_t, \varepsilon_t)\} q(\varepsilon_t) dp_t d\varepsilon_t. \end{aligned}$$

Thus, ρ is the conditional density of the segment ω_{i+1} given the previous segment ω_i , and given that the length of segment ω_{i+1} is $t_{i+1} - t_i$, i.e. the duration between purchases at times t_i and t_{i+1} . Comparing equations (44) and (49) and noting that due to the Markov property we have $Pr\{\tau | \omega\} = Pr\{\tau | \xi\}$ when the last subvector in ω is ξ , we see that

$$v_s(\omega' | \omega, \theta) = \rho(\omega' | \omega, \tau, \theta) Pr\{\tau | \omega, \theta\}. \quad (49)$$

Note that ω_{i+1} implicitly contains the information on $t_{i+1} - t_i$ since this duration is also proportional to the length of ω_{i+1} as we can see in Definition 5. Thus, since the realized value of the duration between successive purchases $t_{i+1} - t_i$ is implicitly determined by ω_{i+1} , we suppress $t_{i+1} - t_i$ in formula (47) in order to emphasize that the normalized log-likelihood function can be written as a normalized sum of random variables that depend on the realizations of an ergodic segmented Markov chain $\{\omega_i\}$. Under suitable regularity conditions on the moments of the functions v_j , $j = 1, 2$, Assumption 1 and the Ergodic Theorem

for Markov processes imply that as $n \rightarrow \infty$ we have with probability 1:

$$\frac{1}{n-1} \sum_{i=1}^{n-1} v_j(\omega_{i+1}, \omega_i, \theta) \rightarrow E \{v_j(\omega', \omega, \theta)\}, \quad j = 1, 2. \quad (50)$$

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

$$E \{v_j(\omega', \omega, \theta)\} = \int \int v_j(\omega', \omega, \theta) v_s(d\omega' | \omega, \theta^*) \psi(d\omega, \theta^*), \quad j = 1, 2. \quad (51)$$

where $v_s(\omega' | \omega, \theta^*)$ is the transition density for the segmented process given in equation (44) and $\psi(\omega, \theta^*)$ is the invariant distribution for the segmented chain $\{\omega_i\}$. Using the alternative representation of l_p in equation (47), we are now able to verify the consistency of the PIML estimator. Note that as $n \rightarrow \infty$, the existence of the ergodic limits in equation (51) imply that the following limits hold

$$\frac{1}{n-1} \sum_{i=1}^{n-1} \log \rho(\omega_{i+1} | \omega_i, t_{i+1} - t_i, \theta) \rightarrow E \{\log \rho(\omega' | \omega, \tau, \theta)\}, \quad (52)$$

and

$$\frac{1}{n-1} \sum_{i=1}^{n-1} \log Pr\{t_{i+1} - t_i | \xi_{t_i}, \theta\} \rightarrow E \{\log Pr\{\tau | \omega, \theta\}\}, \quad (53)$$

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

$$E \{\log \rho(\omega' | \omega, \tau, \theta)\} = \int \sum_{\tau=1}^{\infty} \left[\int \log \rho(\omega' | \omega, \tau, \theta) \rho(d\omega' | \omega, \tau, \theta^*) \right] Pr\{\tau | \omega, \theta^*\} \psi(d\omega, \theta^*). \quad (54)$$

Note that for any ω and τ , the Information Inequality guarantees that the expression in brackets in (54) is maximized at $\theta = \theta^*$. Similarly we have

$$E \{\log [Pr\{\tau | \xi, \theta\}]\} = E \{\log [Pr\{\tau | \omega, \theta\}]\} = \int \sum_{\tau=1}^{\infty} \log [Pr\{\tau | \omega, \theta\}] Pr\{\tau | \omega, \theta^*\} \psi(d\omega) \quad (55)$$

will also be maximized at $\theta = \theta^*$. This implies that the limiting expected log likelihood is maximized at θ^* . Standard uniform consistency arguments can be used to show that with probability 1 we have $\hat{\theta}_p \rightarrow \theta^*$ as $n \rightarrow \infty$.

We conclude this section with a brief sketch the derivation of the asymptotic distribution of the PIML estimator. If model is correctly specified and appropriate regularity conditions hold, the first order conditions for $\hat{\theta}_p$ 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:

$$\sqrt{n} [\hat{\theta}_p - \theta^*] \longrightarrow N(0, I^{-1}(\theta^*)) \quad (56)$$

where

$$I(\theta^*) = I_1(\theta^*) + I_2(\theta^*)$$

where

$$I_1(\theta^*) = E \left\{ \frac{\partial^2}{\partial \theta \partial \theta'} \log [\rho(\omega' | \omega, \tau, \theta^*)] \right\}$$

and

$$I_2(\theta^*) = E \left\{ \frac{\partial^2}{\partial \theta \partial \theta'} \log [Pr\{\tau | \xi_0, \theta^*\}] \right\}.$$

Further it is not difficult to show that the difference between the information matrices for the FIML and PIML estimators 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 Kullback-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 purchases 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 $\{p_t, x_t\}$, 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 $g(p_{t+1}, x_{t+1} | p_t, x_t)$.

It is possible to consider the use of flexible reduced-form specifications for the densities entering the overall decomposition of the transition density λ given in Theorem 4. 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 $g(p', x' | p, x)$ 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 firm, then the wholesale market price $\{p_t\}$ could drift arbitrarily far away from the retail prices $\{p'_t\}$ without there being any strong effect on the likelihood of the observed sequences of purchases. Thus it is impossible to non-parametrically identify the form of $g(p', x' | p, x)$ 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 method of “calibration”. The main difference is that the SMD estimator is based on an explicit statistical criterion function that enables us to compute asymptotic distributions for the parameter estimator, evaluate the fit of alternative specifications, and to conduct goodness of fit tests.

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

from deriving a consistent SMD estimator for θ^* . In fact 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 $S \rightarrow \infty$. This increase in variance seems small in comparison to the substantial reduction in computational burden from using only a single simulation of the model. Estimation still requires a nested fixed point algorithm to solve for the optimal (S, s) policy and a re-simulation of the model using a fixed set of 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 $\mathcal{K} \times 1$ 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:

$$h_T \equiv \frac{1}{T} \sum_{t=1}^T h(\xi_t, \xi_{t-1}), \quad (57)$$

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

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

$$\sqrt{T}[h_T - E\{h\}] \implies N(0, \Omega(h)), \quad (58)$$

where

$$\Omega(h) = E\{(h(\xi', \xi) - E\{h\})(h(\xi', \xi) - E\{h\})'\}, \quad (59)$$

where the expectation in (59) 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

$$\{\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0)\}, \quad t = 2, \dots, T \quad (60)$$

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 $\{\xi_t(\{u_s\}_{s \leq t}, \theta)\}$ depend only on the first t realized values of $\{u_s\}$ and not on subsequent realized values of u_s for $s > t$. Note that we allow the simulated process to depend on the first value ξ_0 of the observed data as an initializing condition.

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

$$\xi_1(u_1, \theta, \xi_0) = P^{-1}(u_1|\xi_0, \theta). \quad (61)$$

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 $t = 2, 4, \dots$

$$\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0) = P^{-1}(u_t|\xi_{t-1}(\{u_s\}_{s \leq t-1}, \theta, \xi_0), \theta). \quad (62)$$

We can see recursively that $\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0)$ 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 $\xi_t = (p_t, p_t', q_t, q_t^o, x_t)$, we can do a similar simulation as in the univariate case described above, using a factorization of the transition density of $\{x_t\}$ into a product of univariate conditional densities such as given in Theorem 4. For example, if ξ_t has two components, $\xi_t = (\xi_{1,t}, \xi_{2,t})$, suppose that its transition density λ can be factored as

$$\lambda(\xi_{t+1}|\xi_t, \theta) = \lambda_2(\xi_{2,t+1}|\xi_{1,t+1}, \xi_t, \theta)\lambda_1(\xi_{1,t+1}|\xi_t, \theta), \quad (63)$$

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

$$\begin{aligned}\xi_{1,1} &= P_1^{-1}(u_{1,1}|\xi_0, \theta) \\ \xi_{2,1} &= P_2^{-1}(u_{2,1}|\xi_{1,1}, \xi_0, \theta).\end{aligned}\tag{64}$$

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:

$$\begin{aligned}\xi_{1,t+1} &= P_1^{-1}(u_{1,t+1}|\xi_t, \theta) \\ \xi_{2,t+1} &= P_2^{-1}(u_{2,t+1}|\xi_{1,t+1}, \xi_t, \theta).\end{aligned}\tag{65}$$

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

$$h_T(\{u_s\}_{s \leq T}, \xi_0, \theta) = \frac{1}{T} \sum_{t=1}^T h(\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0), \xi_{t-1}(\{u_s\}_{s \leq t-1}, \theta, \xi_0)).\tag{66}$$

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

$$h_{S,T}(\theta) = \frac{1}{S} \sum_{i=1}^S h_T(\{u_s^i\}_{s \leq T}, \xi_0, \theta).\tag{67}$$

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

$$\hat{\theta}_T = \underset{\theta \in \Theta}{\operatorname{argmin}} (h_{S,T}(\theta) - h_T)' W_T (h_{S,T}(\theta) - h_T),\tag{68}$$

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

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

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

$$\{\xi_t(\{u_s\}_{s \leq t}, \theta^*, \xi_0)\} \sim \{\xi_t\} \quad (69)$$

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

We believe that it is possible to relax assumption 2 to allow the parametric model to be misspecified, following an analysis similar to that of Hall and Inoue (2002) who characterized the asymptotic properties of the GMM estimator in the misspecified case. We conjecture that their analysis will also apply to the case of SMD estimation and that the asymptotic properties of the SMD estimator that we derive for the correctly specified case will still hold, except that now θ^* is interpreted as the value of θ the minimizes the distance between the moments of the true data generating process and the parametric simulated process, where the expectation is taken in the limit as both $S \rightarrow \infty$ and $T \rightarrow \infty$.⁷

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

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

$$\psi(\xi'|\theta) = \int \lambda(\xi'|\xi, \theta) d\psi(\xi|\theta). \quad (70)$$

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

$$\begin{aligned} E\{h|\theta\} &= \int h(\xi', \xi) d\lambda(\xi'|\xi, \theta) d\psi(\xi|\theta) \\ \nabla E\{h|\theta\} &= \frac{\partial}{\partial \theta} E\{h|\theta\} \\ \nabla h_{S,T}(\theta) &= \frac{\partial}{\partial \theta} h_{S,T}(\theta). \end{aligned} \quad (71)$$

Assumption 4: θ^* is **identified**; that is, if $\theta \neq \theta^*$, then $E\{h|\theta\} \neq E\{h|\theta^*\} = E\{h\}$. Furthermore, $\text{rank}(\nabla E\{h|\theta\}) = \mathcal{K}$ and $\lim_{T \rightarrow \infty} W_T = W$ with probability 1 where W is a $\mathcal{J} \times \mathcal{J}$ 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

⁷When there is misspecification, the standard formula for the asymptotic covariance matrix when the model is correctly specified will generally not be consistent when the model is misspecified. However similar to the case of maximum likelihood estimation of misspecified models (White, 1982), there are alternative estimators of the asymptotic covariance matrix which are consistent when the model is misspecified and when the model is correctly specified.

have

$$\limsup_{T \rightarrow \infty} \sup_{\theta \in \Theta} |(h_{S,T}(\theta) - h_T)' W_T (h_{S,T}(\theta) - h_T) - (E\{h|\theta\} - E\{h|\theta^*\})' W (E\{h|\theta\} - E\{h|\theta^*\})| = 0. \quad (72)$$

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

$$(h_{S,T}(\hat{\theta}_T) - h_T)' W_T \nabla h_{S,T}(\hat{\theta}_T) = 0. \quad (73)$$

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

$$h_{S,T}(\hat{\theta}_T) = h_{S,T}(\theta^*) + \nabla h_{S,T}(\tilde{\theta}_T)(\hat{\theta}_T - \theta^*), \quad (74)$$

where $\tilde{\theta}_T$ denotes a vector that is (elementwise) on the line segment between $\hat{\theta}_T$ and θ^* . Substituting (74) into the first order condition for $\hat{\theta}_T$ in equation (73) and solving for $(\hat{\theta}_T - \theta^*)$ we obtain

$$(\hat{\theta}_T - \theta^*) = - [\nabla h_{S,T}(\tilde{\theta}_T)' W_T \nabla h_{S,T}(\hat{\theta}_T)]^{-1} \nabla h_{S,T}(\hat{\theta}_T)' W_T [h_{S,T}(\theta^*) - h_T], \quad (75)$$

where we assume that $[\nabla h_{S,T}(\tilde{\theta}_T)' W_T \nabla h_{S,T}(\hat{\theta}_T)]$ is invertible, which will be the case with probability 1 for sufficiently large T due to assumptions 3 and 4. Now multiply both sides of equation (75) by \sqrt{T} and apply a Central Limit theorem to the difference $\sqrt{T}[h_{S,T}(\theta^*) - h_T]$ to obtain

$$\sqrt{T}[h_{S,T}(\theta^*) - h_T] \implies N(0, (1 + 1/S)\Omega(h, \theta^*)). \quad (76)$$

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

$$\sqrt{T}[h_T - E\{h|\theta^*\}] \implies N(0, \Omega(h, \theta^*)). \quad (77)$$

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

$$\sqrt{T}[h_T(\{u_s^i\}_{s \leq T}, \theta^*) - E\{h|\theta^*\}] \implies N(0, \Omega(h, \theta^*)). \quad (78)$$

Note that

$$[h_{S,T}(\theta^*) - h_T] = \left[\frac{1}{S} \sum_{i=1}^S [h_T(\{u_s^i\}_{s \leq T}, \theta^*) - E\{h|\theta^*\}] + E\{h|\theta^*\} - h_T \right], \quad (79)$$

so that we have

$$\sqrt{T}[h_{S,T}(\theta^*) - h_T] \implies \left[\frac{1}{S} \sum_{i=1}^S \tilde{X}_i + \tilde{X}_0 \right], \quad (80)$$

where $(\tilde{X}_0, \tilde{X}_1, \dots, \tilde{X}_S)$ are IID $N(0, \Omega(h, \theta^*))$ random vectors. It follows immediately that the asymptotic distribution of $\sqrt{T}[h_{S,T}(\theta^*) - h_T]$ is $N(0, (1 + 1/S)\Omega(h, \theta^*))$. Using this result and equation (75) we have

$$\sqrt{T}[\hat{\theta}_T - \theta^*] \implies N(0, (1 + 1/S)\Lambda_1^{-1}\Lambda_2\Lambda_1^{-1}), \quad (81)$$

where

$$\begin{aligned} \Lambda_1 &= [\nabla E\{h|\theta^*\}]' W \nabla [E\{h|\theta^*\}] \\ \Lambda_2 &= [\nabla E\{h|\theta^*\}]' W \Omega(h, \theta^*) W [\nabla E\{h|\theta^*\}]. \end{aligned} \quad (82)$$

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:

$$\sqrt{T}[\hat{\theta}_T - \theta^*] \implies N(0, (1 + 1/S)\Lambda^{-1}) \quad (83)$$

where:

$$\Lambda = [\nabla E\{h|\theta^*\}]' [\Omega(h, \theta^*)]^{-1} \nabla E\{h|\theta^*\}. \quad (84)$$

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

The SMD estimator can be implemented in practice by solving

$$\hat{\theta}_T = \underset{\theta \in \Theta}{\operatorname{argmin}} (h_{S,T}(\theta) - h_T)' [\hat{\Omega}(h, \theta)]^{-1} (h_{S,T}(\theta) - h_T), \quad (85)$$

where

$$\hat{\Omega}(h, \theta) = \frac{1}{T} \sum_{t=1}^T \varepsilon_t(\theta) \varepsilon_t(\theta)' \quad (86)$$

where

$$\varepsilon_t(\theta) = h(\xi_t(\{u_s\}_{s \leq t}, \theta, \xi_0), \xi_{t-1}(\{u_s\}_{s \leq t-1}, \theta, \xi_0)) - h_T(\{u_s\}_{s \leq T}, \xi_0, \theta). \quad (87)$$

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 3. Such an estimator can attain the Cramer-Rao efficiency bound derived for the PIML estimator in equation (56). 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 3. 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 3. 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

$$\frac{1}{n-1} \sum_{i=1}^{n-1} h(\omega_{i+1}, \omega_i), \quad (88)$$

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

Finally, we note that it appears that it is possible to relax assumption 2 that the parametric model is correctly specified. As long as assumptions 3 and 4 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:

$$\theta^* = \underset{\theta \in \Theta}{\operatorname{argmin}} [E\{h|\theta\} - E\{h\}]' W [E\{h|\theta\} - E\{h\}], \quad (89)$$

where $E\{h\}$ denotes the limit of h_T as $T \rightarrow \infty$ 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^* :

$$(E\{h|\theta^*\} - E\{h\})' W \nabla E\{h|\theta^*\} = 0, \quad (90)$$

where $E\{h\}$ denotes the long run or ergodic expectation of h with respect to the true data generating

process. This implies that as $t \rightarrow \infty$ the random vector

$$X_t \equiv \nabla E\{h|\theta^*\}' W(h(\xi_t(\{u_s\}_{s \leq t}, \theta^*, \xi_0)) - h(\xi_t)), \quad (91)$$

satisfies

$$\begin{aligned} \lim_{t \rightarrow \infty} E\{X_t\} &= 0 \\ \lim_{t \rightarrow \infty} \text{cov}(X_t) &= \Lambda_2 \end{aligned} \quad (92)$$

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 (82). Using a suitable Central Limit theorem for mixing processes, we should have

$$\sqrt{T} \nabla E\{h|\theta^*\}' W[h_T(\{u\}, \theta^*) - h_T] \implies N(0, \Lambda_2). \quad (93)$$

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.

$$\sqrt{T} [\hat{\theta}_T - \theta^*] \implies N(0, (1 + 1/S)\Lambda_1^{-1}\Lambda_2\Lambda_1^{-1}), \quad (94)$$

where

$$\Lambda_1 = [\nabla E\{h|\theta^*\}' W \nabla E\{h|\theta^*\}] \quad (95)$$

and where

$$\sqrt{T} \nabla h_T(\{u_s\}, \theta^*)' W_T [h_T(\{u_s\}, \theta^*) - h_T] \implies N(0, \Lambda_3). \quad (96)$$

The main outstanding issue is to actually establish the limiting asymptotic distribution that is conjectured in (96) and relate the asymptotic covariance matrix Λ_3 to the asymptotic covariance matrix Λ_2 in (93). As we noted above, we believe that results of Hall and Inoue (2002) on GMM estimation of misspecified models can be adapted to establish the asymptotic distribution of the SMD estimator in the misspecified case. However given the space constraints we leave this topic, together with Monte Carlo tests and an empirical application of the SMD estimator for a misspecified model, as a topic for subsequent research.

5 Empirical Application

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

wholesale 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. Finally, we decompose the firm's profits by product into four components. We use this decomposition to infer the share of the firm's profits that are due to markups paid by retail customers and the share due to price speculation. We also use this decomposition to compare the general manager's purchasing decisions to the model's trading rules.

5.1 A special case of the model

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

$$\max_{\{q_t^o\}} E \sum_{t=0}^{\infty} \rho^t \left\{ p_t^r q_t^s - c^o(q_t^o, p_t) - c^h(q_t + q_t^o, p_t) \right\} \quad (97)$$

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

$$\begin{aligned} c^o(q_t^o, p_t) &= \begin{cases} p_t q_t^o + K & \text{if } q_t^o > 0 \\ 0 & \text{otherwise,} \end{cases} \text{ and} \\ c^h(q_t + q_t^o, p_t) &= \phi(q_t + q_t^o)^2. \end{aligned}$$

As before, the manager takes the wholesale 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 holding cost function, $c^h(\cdot)$, are described in section 2. The holding cost function is quadratic so the marginal convenience yield is decreasing in the level of inventories.

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

$$\log(p_{t+1}) = \mu_p + \lambda_p \log(p_t) + w_t^p \quad (98)$$

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

$$\tilde{\sigma}_p = \sqrt{\log(\bar{\sigma}_p^2 + \bar{\mu}_p^2) - 2\log(\bar{\mu}_p)}. \quad (99)$$

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

$$\mu_p = (1 - \lambda_p)(\log(\bar{\mu}_p) - \tilde{\sigma}_p^2/2) \text{ and } \sigma_p = \tilde{\sigma}_p \times \sqrt{1 - \lambda_p^2}. \quad (100)$$

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

$$p_t^r = \alpha_0 + \alpha_1 p_t. \quad (101)$$

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

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

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

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

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

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

5.2 Computation

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). The PPI algorithm involves approximating the value function $V(p, q)$ given in equation (11) as a linear combination of N basis functions, $\{\varphi_1(p, q), \varphi_2(p, q), \dots, \varphi_N(p, q)\}$:

$$V(p, q) \approx \sum_{n=1}^N \vartheta_n \varphi_n(p, q). \quad (102)$$

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

$$\begin{aligned} \sum_{n=1}^N \vartheta_n \varphi_n(p_m, q_m) &= \max_{0 \leq q^o \leq \bar{q} - q_m} \left[ES(p_m, q_m) - c^o(q^o, p_m) - c^h(q_m, p_m) + \right. \\ &\quad \left. \rho E \left\{ \sum_{n=1}^N \vartheta_n \varphi_n(p', \max[0, q_m - q^s + q^o]) \mid p_m, q_m \right\} \right] \quad \text{for } m = 1, \dots, M. \end{aligned} \quad (103)$$

As the name suggests, PPI employs an iterative strategy to find the N coefficients on the basis functions that solve the system of equations in (103). Given an initial guess of the coefficient vector, ϑ , we solve the two-period problem on the right-hand side of (103) for each discretize pair of (p, q) . This yields an $(M \times 1)$ vector of the decision rule $q^o(p, q)$. Note that although we discretized the state variables, q^o is a continuous variable subject to the frequently binding constraint: $0 \leq q_i^o \leq \bar{q} - q$.

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

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

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

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

and let the $(M \times N)$ matrix X be given by $X = (P - \rho EP)$. Then the system of equations (103) can be written in matrix form as $y = X\vartheta$. If $M = N$ and X is invertible, the solution for ϑ is simply $\hat{\vartheta} = y/X$. If $M > N$, we form an approximate solution using ordinary least squares estimation, i.e. $\hat{\vartheta} = (X'X)^{-1}X'y$. Using $\hat{\vartheta}$ as our updated coefficient vector, we iterate on this procedure until the coefficient vector converges to a fixed point.

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

5.3 Estimation

We have considerable freedom in our choice of moment functions, the h vector, to use in the criterion. As discussed above, the most efficient moment functions we could use would be the score of the partial information maximum likelihood function derived in section 3. However given the difficulties in computing the

high dimensional integrals involved in evaluating the score, we instead match the means and histograms (four of the five quintile binds) of the p , p^r , q^o , q^s , and q processes for a total of 25 moment conditions. We set the number of simulations, S , to 10.

Computing histogram bins requires the use of indicator functions. To ensure these indicator functions do not add discontinuities into the criterion function, we use the exponential transformation of the indicator functions. In general, the criterion is (at least visually) a smooth function of the parameters. However we did find regions in the parameter space in which concentrated “slices” of the criterion function had “steps” and “cliffs.” Therefore to verify that MATLAB’s constrained minimization routine `fmincon.m` found a global minimum in each case, we visually inspected concentrated slices of the criterion function after each estimation.

As presented in equations (86) and (87) of the previous section, the inverse of the optimal weighting matrix, $\hat{\Omega}(h, \theta)$ is the variance-covariance of the residuals from the simulation sequence. However if the model is correctly specified, then when $\theta = \theta^*$, the simulated sequence will have the probability distribution as the observed sequence; therefore we use inverse of the variance-covariance matrix of the residuals of the observed sequence as our weighting matrix, W . Since this weighting matrix is just a function of the sample moments, it remains fixed throughout the estimation.

parameter	truth	Simulation 1		Simulation 2		Simulation 3	
		point estimate	standard error	point estimate	standard error	point estimate	standard error
K	100	108.6	11.6	138.7	16.5	87.4	10.5
α_0	1.50	1.46	0.66	1.80	0.47	1.45	0.45
α_1	1.15	1.13	0.04	1.11	0.03	1.20	0.03
λ_p	0.990	0.991	0.0003	0.989	0.0007	0.990	0.0003
$\bar{\mu}_p$	19.50	20.06	0.60	20.10	0.49	19.78	0.55
$\bar{\sigma}_p$	5.60	6.39	0.29	6.18	0.27	5.33	0.30
$\bar{\mu}_q$	150.0	137.1	6.5	157.1	4.7	130.3	3.6
$\bar{\sigma}_q$	300.0	363.5	25.2	270.3	12.1	250.5	11.4
ζ	1.50	1.31	0.17	1.41	0.17	1.62	0.21
ϕ	-2.5	-2.69	1.36	-1.87	1.37	-2.67	1.01
r	0.075/365	0.075/365		0.075/365		0.075/365	
η	0.35	0.35		0.36		0.33	
$\chi^2(15)$		381		187		217	

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

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

In our 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 three simulated data sets of 1191 periods from the model; second using these simulated data sets, we estimated the model using our simulated minimum distance estimator. The point estimates and standard errors for each of the ten parameters are reported in table 1. Prior to estimation, we set the interest rate equal to its' true value and η equal to the fraction of days in which no sale occurred.

The quantity data are in hundred-weight (i.e. in 100's of pounds) so the price parameters are in dollars-per-hundredweight (or cents per pound). The fixed cost, K , is set to \$100 per order. The parameter choices for \bar{m} and $\bar{\sigma}_p$ imply the uncensored price process has a mean of \$17.60 per hundred-weight or 17.6 cents per pound and a standard deviation of \$3.70 dollars per hundred-weight. The parameter values of \bar{m} , $\bar{\sigma}_q$ and ζ imply the average sale is 107 hundred-weight or 1,070 pounds. The interest rate r is set to 7.5 percent per annum. The storage cost net of convenience yield, ϕ is set -2.75 dollars per squared hundred-weight, so the convenience yield dominates the storage cost.

For most of the parameters, the point estimates seem reasonably close to their true values. For example, all three of the point estimates of the AR(1) coefficient of the wholesale price process, λ_p , are within two-tenths of one percent of the true value. All three point estimates of the fixed cost, K , are sensible particularly given the difficulty in estimating K . The fixed cost of ordering largely determines the distance between the S and s bands and thus the minimum order size. To accurately identify this fixed cost requires numerous observations of days in which the firm is holding inventory levels close to s particularly at low prices. Given the relatively few days the firm purchases, particular at low prices, there are very few days the firm holds inventories close to s .

While we feel the SMD estimator delivers sensible point estimates, only two-thirds of the point estimates are within two standard errors of the true values. Unfortunately, several of the numerical standard errors seem implausibly small, particularly given the variation in the three point estimates. For example, the standard deviations of the point estimates for K and λ_p are considerably larger than the standard errors. Moreover, the estimation procedure provides a formal criterion of the validity of model. Since the number of moment conditions exceeds the number of parameters estimated ($\mathcal{J} > \mathcal{K}$) the model is overidentified. Following Hansen (1982), we use the objective function to test the overidentifying restrictions:

$$\frac{T}{(1+1/S)^2} (h_{S,T}(\hat{\theta}) - h_T)' [\hat{\Omega}(h)]^{-1} (h_{S,T}(\hat{\theta}) - h_T) \rightarrow \chi^2(\mathcal{J} - \mathcal{K}) \quad (104)$$

In bottom row of table 1 we report the value of this χ^2 statistic for each of three estimates. In each case, the

model is decisively rejected. We have run several diagnostics to check that these small standard errors and large χ^2 statistics are not due to computational or approximation errors. Ideally, we would like to compare our numerical standard errors with empirical standard errors computed from a full-fledged Monte Carlo exercises. However, since it take several days to estimate the model, at the moment a true Monte Carlo exercise is computationally infeasible. We are working on speeding up the code. In particular, we want to experiment further with various choices of the weight matrix, W .

parameter	Product 2		Product 4	
	point estimate	standard error	point estimate	standard error
K	39.2	6.1	59.6	6.9
α_0	1.33	0.98	0.99	1.10
α_1	0.98	0.04	1.10	0.05
λ_p	0.992	0.0006	0.984	0.003
\bar{h}	24.40	0.66	18.55	0.60
$\bar{\sigma}_p$	7.98	0.25	4.83	0.41
\bar{h}	215.2	7.7	301.8	6.9
$\bar{\sigma}_q$	747.6	41.8	496.5	31.1
ζ	1.48	0.20	0.92	0.15
ϕ	-2.70	3.65	-2.72	2.74
r	0.075/365		0.075/365	
η	0.34		0.34	
$\chi^2(15)$	522		334	

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

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

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

Although we estimated the parameters for each of these products independently, it is reassuring that several of the point estimates are similar across the two products. It is reasonable to expect that the

parameters, K , α_0 , α_1 , λ_p , ζ , and ϕ to be quite similar, if not identical, across products.⁸ In general this is case. After we estimated the models, we asked the general manager for his estimate of K . He stated \$50 – the midpoint of our two estimates. 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 marginal cost of storage parameter, ϕ , is negative for both products so the marginal convenience yield dominates the physical costs of storage. This result is consistent with the observation in the commodity storage literature that negative storage costs are a key determinate of the autocorrelation in commodity prices. We experimented with various function forms for the holding cost function and stock-out penalty functions. If the marginal value of holding inventories is small when inventories are close to zero (i.e. when the wholesale price is high), the optimal strategy is for the firm to effectively shut down by holding no inventories until the wholesale price falls. In other words, the $s(p)$ band equals zero for p greater than some threshold. While we do observe near-zero levels inventories in the data from time to time, these near stockout levels do not persist for more than a few days. If the marginal value of holding inventories is “too large” even when the firm is holding large levels of inventories, the model implies the firm should (counterfactually) always hold inventories near its capacity constraint. Hence we found having some convexity in the holding cost helpful in matching mean and spread of inventories holdings we see in the data.

The endogenous sampling problem is illustrated in figures 3, 4, and 5. 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 4. Due to the fixed costs of ordering, the $S(p)$ band is strictly above the $s(p)$ band although the difference between the two bands decreases as the price increases. In other words, the minimum order size is a decreasing function of the price. In figure 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 graph). In the simulation presented, this occurs less than 16 percent of time.

In table 2 we also report the minimized SMD estimation criterion. Although both models are formally rejected, the models at the estimated parameter values capture several of the salient features of the inventory and price data. facts of inventory investment behavior that we observe in our data (for further discussion, see Hall and Rust, 1999). Figures 3, 4, 5 highlight some of the strengths of the model. First, in the data purchases are made infrequently. Figure 5 presents the censored and uncensored purchase price

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

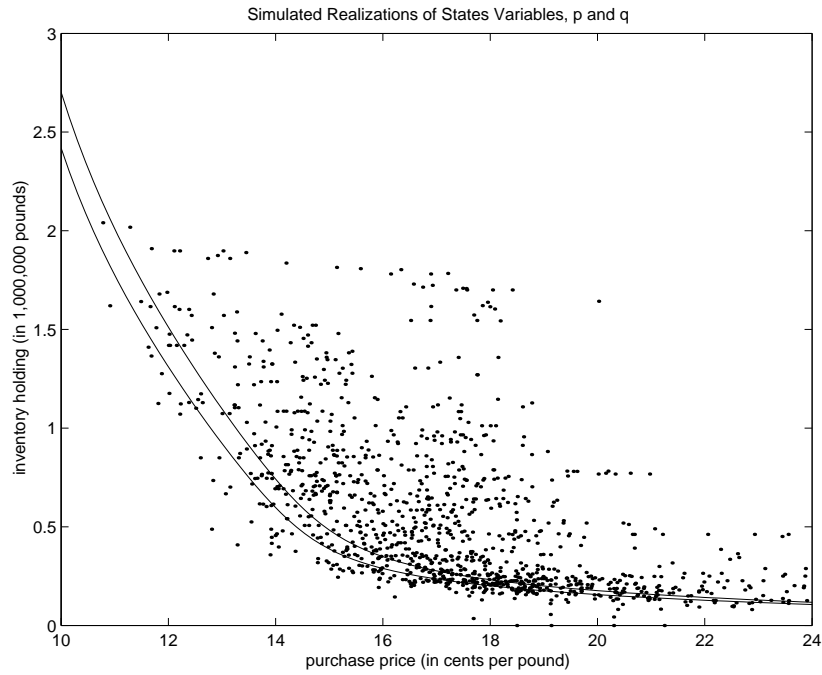


Figure 3: Scatterplot of purchase price and inventory holding pairs from a simulation for product 4. The solid lines are the $S(p)$ and $s(p)$ bands from the model.

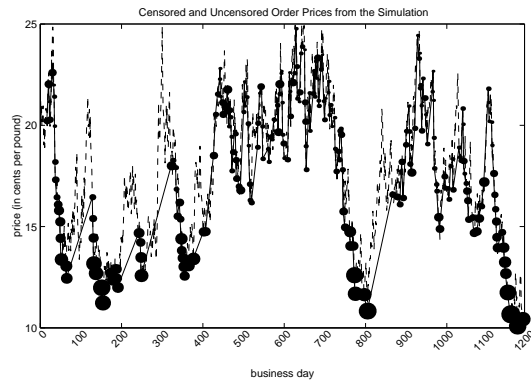
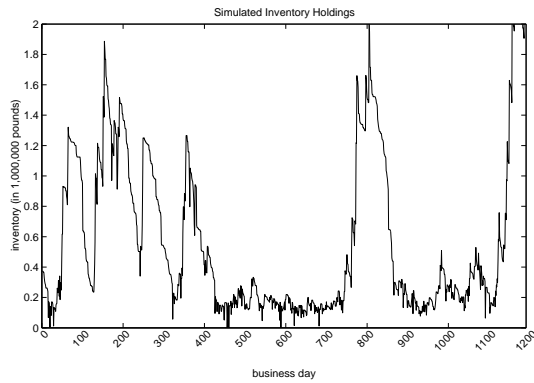


Figure 4: Simulated inventory data from the estimated Figure 5: Censored (solid line) and uncensored (dotted line) purchase prices, p_t from a simulation for product 4.

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 100-200, 350-400 and 750-800) the firm aggressively made purchases to build up large levels of inventories. The large levels of inventories were slowly drawn down as prices inevitably rose. Note there were only four purchases made between business days 200 and 320. Thus after exploiting a low price opportunity, the firm may subsequently make no new purchases for many days. Second, we observe both small and large purchases in the data. Again this can be seen 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 5, the size of the marker is proportional to the size of the purchase. Again one 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 and the time path of inventories plotted in figure 4 we can see that the model predicts that inventory levels will vary over the sample between almost zero and 2.0 million pounds.

The main shortcoming of the estimation is our inability to match the downward trend of the price process that we see in almost all of the firm's products. As illustrated in 2 the wholesale price for product 4 fell from 20 cents per pound in 1997 to about 12 cents per pound in 2002. No such trend is evident in simulations such as the one presented in figures 4 and 5. In our model, prices are stationary though highly persistent. Consequently, as can be seen in the (S, s) bands 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 procurement price is above 17 cents per pound. From figures 1 and 2 we see that, for product 4, the firm made large purchases around 18.5 cents per pound in April 1998, and around 15 dollars per hundred-weight in the later part of the sample.

An often suggested solution to this trend problem is that we assume prices follow a random walk. For product 4, if we concentrate out all the other parameters except λ_p , the criterion surface is a steeply sloped and smooth cup centered around 0.984 so the small standard error associated with the AR(1) coefficient is not surprising. But the concentrated criterion surface actually turns down slightly between .995 and 1.01. (The model still solves numerically for values of λ_p slightly greater than one.) The global minimum is still located at 0.984, 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 such that the inventory level fluctuates closely around a fixed target level. A version of 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. A second potential solution is to detrend the data. However when we first started working on this project, no one we talked to expected steel price to decline 40% in four years. To some extent we are just working with too short a sample period. A third candidate solution is to add an additional macroeconomic state variable. Such a variable could allow for “high price” regimes and “low price” regimes. As we discuss below, we view this third solution as the most promising.

5.4 A profit decomposition exercise

Finally, we use simulations of the estimated model to deduce the relative importance of capital gains versus markups for the overall profitability of the firm. By substituting the law of motion for inventories (4) into the firm’s objective function, (97), the discounted present value of the firm’s profits can be expressed by

$$\begin{aligned} \sum_{t=1}^T \rho^t \pi(p_t, p_t^r, q_t^r, q_t + q_t^o) &= \sum_{t=1}^T \rho^t (p_t^r - p_t) q_t^s + q_1 p_1 + \sum_{t=2}^T \rho^t (p_t - (1+r)p_{t-1}) q_t - \\ &\quad \sum_{t=1}^T \rho^t I(q_t^o) K - \sum_{t=1}^T \rho^t c^h(q_t + q_t^o, p_t) \end{aligned} \quad (105)$$

The first term on the right hand side of equation (105) can be interpreted as the discounted present value of the markup paid by the firm’s retail customers over the current wholesale price while the third term can be interpreted as the discounted present value of the capital gains or loss from holding the steel from period $t - 1$ into period t . The fourth, and fifth terms are the discounted present values of the order costs and the holding costs incurred by the firm over the sample period.

Since this decomposition depends on the wholesale price path between purchases, we simulate between purchase dates via importance sampling. That is, for each interval between successive purchase dates, we simulate wholesale price paths that are consistent with the estimated law of motion (98) and the observed purchase prices at the beginning and end of the interval. Since our theory implies that the firm places an order anytime the quantity falls below the order threshold, $s(p)$, we truncate the simulated price process by rejecting any paths such that $q_t < s(p_t)$ for any draw within the simulated paths. We discuss our simulation method in more detail in the appendix.

We first employ this decomposition to evaluate the general manager’s actual performance over the four-and-a-half year sample period for products 2 and 4. For a given interpolated price series, we decomposed

	Product 2						Product 4					
	G.M.'s actual Performance			Model's Policy Prescription			G.M.'s actual Performance			Model's Policy Prescription		
markup	\$264,671	(26,216)	71%	\$262,288	(24,379)	60%	\$370,864	(35,798)	85%	\$347,023	(31,908)	54%
capital gain	110,148	(28,561)	29%	173,309	(16,524)	40%	66,354	(38,248)	15%	291,500	(20,825)	46%
holding cost	214,397	(0)		226,031	(4,212)		243,179	(0)		218,630	(6,097)	
order costs	-5,951	(0)		-6,981	(321)		-9,742	(0)		-13,456	(524)	
total profits	583,264	(3,479)		654,647	(34,067)		670,655	(3,825)		875,695	(43,840)	

Table 3: Profit Decomposition For Product 2 and 4 Using Equation (105)

Both the actual and the counter-factual profits cover the 1191 days studied and are discounted back to the start of the sample period, July 1, 1997. The profit numbers reported are the average across 100 simulations. The numbers in parentheses are the standard deviations from the 100 simulations. Total profits are the sum of the first four rows.

the firm's profits using the actual data for q_t , q_t^s , and q_t^o , our fixed value for the interest rate, r , and our point estimates for K , and ϕ . In table 3 we report the average decomposition from 100 simulated wholesale price paths. As discussed in the introduction, the price of steel fell steadily over the sample period. Never the less, by our accounting, the firm made \$375,000 (product 2) and \$435,000 (product 4) from the markup and capital gains on each of these two product over the four-and-a-half year period.⁹ Ignoring the fixed order cost and the returns from the convenience yield, about 71 percent (product 2) and 85 percent (product 4) of these profits came from the markup, while the remaining 29 and 15 percent came from capital gains. We find it remarkable and evidence of the general manager's acumen in steel trading that the firm made positive capital gains over this period despite the price of steel falling about 40 percent. While the firm's success in price speculating is good for its profits, it increasing the potential biases from failing to account for the endogeneity of the sampling process.

As a diagnostic of our model, we compare the general manager's performance to the model's predictions. In this exercise we take as given the 100 interpolated wholesale price series, the firm's quantity demanded series, and the firm's initial level of inventories for each product. But in this case, we let the model's optimal decision rule dictate when and how much to order.¹⁰ Inventories follow the accumulation identity given by equation (4). As reported in table 3, had the general manager counter-factually followed the optimal order strategy implied by our model, his discounted profits from the markup would have been modestly smaller: \$2,000 less for product 2; \$23,000 less for product 4. However, his capital gains would have been considerable larger: \$63,000 more for product 2; \$225,146 more for product 4.

The model implies that the firm should aggressively price speculate. In figures 6 and 7 we plot the prices and inventory holdings for one simulation of the model. In figure 6 we plot both the actual inventory holdings along with the implied holdings under the model's decision rules. In figure 7 we plot the corresponding retail and wholesale price paths. The model's counter-factual inventory path differs considerably from the firm's actual inventory path. In the beginning of the sample, years 1997 and 1998, when prices were high, the model implies the firm should have made frequent small purchases and held relatively low levels of inventories. As was discussed in the introduction, in April 1998 when the wholesale price of steel dropped from 20 cents per pound to 18.5 cents per pound, the firm built up its inventory of product 4

⁹Profits are discounted back to July 1, 1997.

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

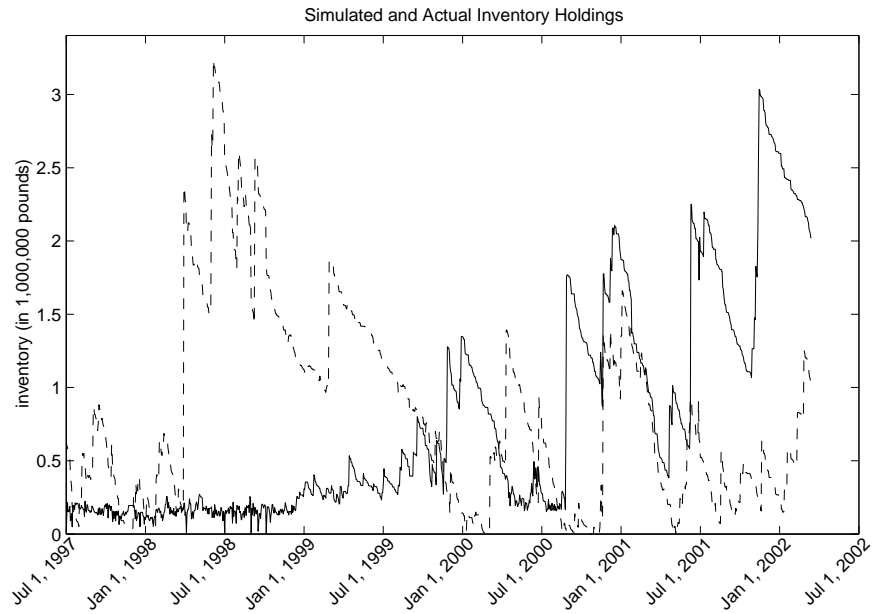


Figure 6: Actual (dashed line) and counter-factual (solid line) inventory holdings for product 4.

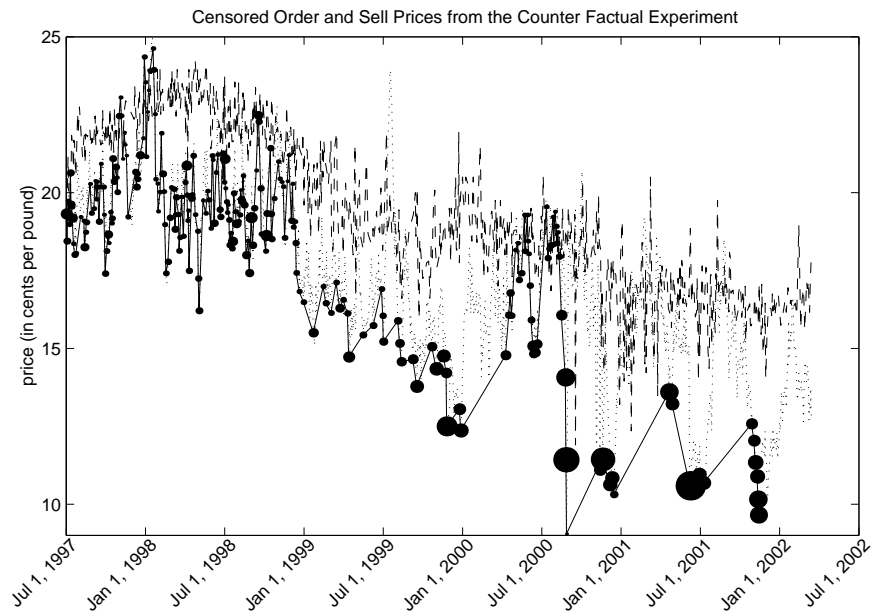


Figure 7: Counter-factual uncensored purchase prices (dotted line), censored purchase prices (solid line), and retail prices (dashed line) for product 4. For the censored purchase price series, the size of the marker is proportional to the size of the purchase.

substantially. In contrast the model does not view 18.5 cents as a particularly good price; as can be seen in the (S, s) bands plotted in figure 3, the target inventory level at 18.5 cents is around 300,000 pounds. In April 1998, the firm's inventory of product 4 exceeded 2,000,000 pounds.

It is not until December 1999 when prices fell below 13 cents a pound that the model recommends holding more than 1,000,000 pounds of inventory. However during December 1999 and January 2000, the general manager let his inventory of product 4 fall to almost zero. The sharp contrast between model's counter-factual inventory policy and the firm's behavior is also evident during the second half of the sample. In this period, the firm held relatively low levels of inventories, whereas the model's inventory was often in excess of 2,000,000 pounds. The only time during the sample that the model's inventory holdings tracked well the firm's inventory holdings was in the first half of 2001. Basically, the model recommends the firm's purchasing strategy should have been the opposite of what it did: the firm should have held low inventory levels in 1997, 1998 and 1999, and high inventory levels in 2000, 2001, and the start of 2002.

This counter-factual exercise is "rigged" in the model's favor in one dimension and "rigged" against the model in another. Since we used the entire sample period to both estimate the model and evaluate the model's performance, the model "knows" the mean and the standard deviations of prices and quantity demanded for the entire period. The model knows, whereas the general manager did not know, that a price of 18.5 cents per pound in the Spring of 1998 was an above-average price for the 1997-2002 period. In this way the model has an advantage over the manager. However the model is constrained to sell at most the quantity of steel the general manager actually sold. The model does not get the opportunity make any sales the general manager might have had the option to make but decided to turn down.

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

percent in the Spring of 2002 to the 12 to 14 cent range. In the Spring of 2002, we reminded the general manager that in the fall our model recommended he build up inventories. He sighed, “I wish I had.”

In this case, our model “got it right” but perhaps not for the right reasons. Our model was predicting an increase in prices since our model always expects prices to return to the sample mean. Our model had no way of predicting where the economy was going. It just expects demand to be stationary; moreover, our model could provide no help in forecasting the President’s actions regarding steel tariffs. In future work, we hope to add a macroeconomic state variable to the analysis. Thus we could use our model jointly with a macroeconomic forecasting model to provide conditional inventory level recommendations to the firm such as “If you expect the economy to remain strong, the model recommends holding inventories in a range from X to Y; if you expect the economy to weaken, ...”

6 Conclusion

In this paper we develop two econometric procedures for estimating an endogenously-sampled Markov process. We first derive a parametric partial information maximum likelihood (PIML) estimator that solves the endogenous sampling problem. While the PIML estimator efficiently estimates the unknown parameters of a Markov transition probability, it requires repeatedly computing numerical approximations to high dimensional integrals. Therefore we introduce an alternative consistent, less efficient, simulated minimum distance (SMD) estimator. This estimation method is computationally simpler than the PIML estimator, but it still requires solving the dynamic programming problem at each trial value of the unknown parameter vector for the endogenous sampling rule. Using this sampling rule, the SMD estimator is able to consistently estimate the unknown parameters of the Markov process even though the econometrician has incomplete information on the process.

While this research was motivated by a new dataset from a single steel wholesaler, most datasets in which agents have the choice of whether and when to participate in a market activity will be endogenously sampled. In most markets, the only prices recorded are the transaction prices – econometricians almost never get to observe prices offered but not transacted on. For example, econometricians rarely get to observe the wages unemployed job seekers are offered but refuse.¹¹ It should be straightforward to apply the SMD estimator to other types of endogenous sampling problems that arise in time series contexts.

¹¹A counter-example is the limit order books for equities posted on ECNs such as www.island.com. But specialists on the NYSE are very reluctant to reveal any information about their limit order books.

References

- [1] Ait-Sahalia, Y. and P. A. Mykland (2001) “The Effects of Random and Discrete Sampling When Estimating Continuous-Time Diffusions” working paper, Princeton University.
- [2] Aguirregabiria, V. (1999) “The Dynamics of Markups and Inventories in Retailing Firms” *Review of Economic Studies* 66, 275–308.
- [3] Allenby, G., McCulloch, R. and P. Rossi (1996) “On the Value of Household Information in Target Marketing” *Marketing Science* 15, 321–340.
- [4] Arrow, K.J., Harris, T. and J. Marschak (1951) “Optimal Inventory Policy” *Econometrica* 19-3, 250–272.
- [5] Athreya, R. (2002) “Price Dispersion in the Wholesale Market for Steel” manuscript, Yale University.
- [6] Benitez-Silva, H., G. Hall, G. Hitsch, G. Pauletto and J. Rust (2001) “A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems” manuscript, State University of New York at Stony Brook.
- [7] Billingsley, P. (1961) *Statistical Inference for Markov Processes* University of Chicago Press.
- [8] Chan, H. M. (2001) “Analysis of Variance in Steel Sales Price” manuscript, Yale University.
- [9] Chan, H., G. Hall and J. Rust (2003) “A Model of Bargaining and Retail Price Discrimination in the Market for Steel” manuscript, Yale University.
- [10] Deaton, A. and G. Laroque (1992) “On the Behavior of Commodity Prices” *Review of Economic Studies* 59, 1–23.
- [11] Duffie, D. and K.J. Singleton (1993) “Simulated Moments Estimation of Markov Models of Asset Prices” *Econometrica* 61-4, 929–952.
- [12] Engle, R. and J. Russell (1999) “Autoregressive Conditional Duration: A New Model for Irregularly Spaced Data” forthcoming, *Econometrica*.
- [13] Erdem, T. and M. Keane (1996) “Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets” *Marketing Science* 15–1, 1–20.

- [14] Fabian, T., J.L. Fisher, M.W. Sasieni, and A. Yardeni (1959) “Purchasing Raw Material on a Fluctuating Market” *Operations Research* 7, 107-122.
- [15] Golabi, K. (1985) “Optimal Inventory Policies when Ordering Prices are Random” *Operations Research* 33-3, 575-588.
- [16] Hall, A. and A. Inoue (2002) “The Large Sample Behavior of the Generalized Method of Moments Estimator in Misspecified Models” manuscript NC State University.
- [17] Hall, G. and J. Rust (1999) “An Empirical Model of Inventory Investment by Durable Commodity Intermediaries” *Carnegie-Rochester Conference Series on Public Policy*, 52.
- [18] Hall, G. and J. Rust (2000) “A (S, s) Model of Commodity Price Speculation” manuscript, Yale University.
- [19] Hall, G. and J. Rust (2001) “The (S, s) Rule is an Optimal Trading Strategy in a Class of Commodity Price Speculation Problems” manuscript, Yale University.
- [20] Hansen, L. P. (1982) “Large Sample Properties of Generalized Method of Moments Estimators” *Econometrica* 50, 1929-1954.
- [21] Keckman, J.(1981) “Statistical Models for Discrete Panel Data” in C. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data* MIT Press.
- [22] Kaldor, N. (1939) “Speculation and Economic Stability” *Review of Economic Studies* 7, 1-27.
- [23] Kalymon, B. (1971) “Stochastic Prices in a Single-Item Inventory Purchasing Problem” *Operation Research* 19, 1434-1458.
- [24] Kingman, B.G. (1969) “Commodity Purchasing” *Operational Research Quarterly* 20-1, 59-79.
- [25] Lee, B. and B.F. Ingram (1991) “Simulation Estimation of Time-Series Models” *Journal of Econometrics* 47, 197-205.
- [26] 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.

- [27] McFadden, D. (1997) “On the Analysis of Endogenously Recruited Panels” manuscript, University of California at Berkeley.
- [28] McFadden, D. (1998) “A Method of Simulated Moments for Estimation of Discrete Models Without Numerical Integration” *Econometrica* 57–5, 995–1026.
- [29] Miranda, M. and Rui, X. (1997) “An Empirical Reassessment of the Nonlinear Rational Expectations Commodity Storage Model” manuscript, Ohio State University, forthcoming, *Review of Economic Studies*
- [30] Moinzadeh, K. (1997) “Replenishment and Stocking Policies for Inventory Systems with Random Deal Offerings” *Management Science* 43–3, 334–342.
- [31] Resnick, S. (1992) *Adventures in Stochastic Processes* Birkhäuser, Boston.
- [32] Russell, J. and R.F. Engle (1998) “Econometric Analysis of Discrete-Valued Irregularly-Spaced Financial Transactions Data Using a New Autoregressive Conditional Multinomial Model” UCSD Discussion Paper 98-10 <ftp://weber.ucsd.edu/pub/econlib/dpapers/ucsd9810.pdf>
- [33] Rust, J. and G. Hall (2003) “Middlemen versus Market Makers: A Theory of Competitive Exchange” *Journal of Political Economy* forthcoming.
- [34] 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.
- [35] 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.
- [36] Song, J. and P. Zipkin (1993) “Inventory Control in a Fluctuating Demand Environment” *Operations Research* 41-2, 351-370.
- [37] White, H. (1982) “Maximum Likelihood Estimation of Misspecified Models” *Econometrica* 50–1, 1–26.
- [38] Williams, J.C. and B. Wright (1991) *Storage and Commodity Markets* Cambridge University Press, New York.

Appendix: Simulating Price Paths with Fixed Starting and Ending Points

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

We assume the spot procurement price follows the AR(1) process given in equation (98) of the paper. To simplify the presentation in this appendix, let p_t denote the $\log(p_t)$. Assume we observe p_{t_1} and p_{t_2} on dates t_1 and t_2 , but we do not observe any prices on dates in between. We want to simulate realizations of $\{p_{t_1+1}, p_{t_1+2}, \dots, p_{t_2-1}\}$ that are consistent with both p_{t_1} and p_{t_2} and the law of motion (98). Let $\tau = t_2 - t_1$, be the recurrence time.

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

$$\begin{bmatrix} 1 \\ p_{t_1} \\ p_{t_2} \\ p_{t_1+1} \\ p_{t_1+2} \\ \vdots \\ p_{t_2-1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & \dots & 0 & 0 \\ \mu_p & 0 & \lambda_p & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ p_{t_1-1} \\ p_{t_2-1} \\ p_{t_1} \\ p_{t_1+1} \\ \vdots \\ p_{t_2-2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \sigma_p \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} w_{t_2}^p. \quad (106)$$

We rewrite this equation using more compact notation as:

$$\mathbf{p}' = \mathbf{A}\mathbf{p} + \mathbf{C}w^{p'} \quad (107)$$

where the \mathbf{p} denotes the vector of logged prices and the prime denotes the next period's values.

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

$$\Omega = \sum_{j=0}^{\tau+1} \mathbf{A}^j \mathbf{C} \mathbf{C}' \mathbf{A}^{j'}$$

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

$$\begin{bmatrix} p_{t_1} - \mu_p \\ p_{t_2} - \mu_p \\ p_{t_1+1} - \mu_p \\ \vdots \\ p_{t_2-1} - \mu_p \end{bmatrix} = \begin{bmatrix} \mathfrak{v}_{11} & 0 & 0 & \dots & 0 \\ \mathfrak{v}_{21} & \mathfrak{v}_{22} & 0 & \dots & 0 \\ \mathfrak{v}_{31} & \mathfrak{v}_{32} & \mathfrak{v}_{33} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathfrak{v}_{\tau+11} & \mathfrak{v}_{\tau+12} & \mathfrak{v}_{\tau+13} & \dots & \mathfrak{v}_{\tau+1\tau+1} \end{bmatrix} \begin{bmatrix} \eta_{t_1} \\ \eta_{t_2} \\ \eta_{t_1+1} \\ \vdots \\ \eta_{t_2-1} \end{bmatrix}. \quad (108)$$

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

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

$\{\eta_{t_1+1}, \eta_{t_1+2}, \dots, \eta_{t_2-1}\}$ are random draws from a standard normal distribution. Once the η vector is constructed, we use equation (108) to compute the simulated price vector $\mathbf{p}' = Y\eta + \mu_p$. Note that each of the simulated prices is a function of η_{t_1} and η_{t_2} .

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