NBER WORKING PAPER SERIES

STICKY PRICES, INVENTORIES, AND MARKET POWER IN WHOLESALE GASOLINE MARKETS

Severin Borenstein
Andrea Shepard

Working Paper 5468

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 1996

For helpful comments, we thank Tim Bresnahan, Dennis Carlton, Robert Pindyck, and seminar participants at U.C. Berkeley, University of Toronto, University of Chicago, Purdue University, and the National Bureau of Economic Research. Wedad Elmaghraby provided excellent research assistance. This paper is part of NBER's research program in Industrial Organization. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper 5468
February 1996

STICKY PRICES, INVENTORIES, AND MARKET POWER IN WHOLESALE GASOLINE MARKETS

ABSTRACT

We present and test an explanation for lags in the adjustment of wholesale gasoline prices to changes in crude oil prices. Our simple model with costly adjustment of production and inventories implies that output prices will respond with a lag to cost shocks even in the absence of menu costs, imperfect information, and long-term buyer/seller relationships. The model predicts that futures prices for gasoline will adjust incompletely to crude oil price shocks occurring close to the expiration date of the futures contract. We test and confirm this implication. The model also predicts that firms with market power will choose a different price adjustment path than would perfectly competitive firms. We examine the responses of prices in 188 local wholesale gasoline markets and find evidence that greater market power leads to slower output price adjustment.

Severin Borenstein
University of California Energy Institute
2539 Channing Way
Berkeley, CA 94720-5180
and NBER

Andrea Shepard
Graduate School of Business
Stanford University
Stanford, CA 94305-5015
and NBER

I. Introduction

The adjustment of market prices to new information has been a central question in both industrial organization and macroeconomics. A primary concern in industrial organization has been the connection between market structure or behavior and the response of output prices to shocks. The macroeconomics literature has focused more on the effect that "sticky" prices have on aggregate economic measures. The combined research has documented lagged price adjustment in a variety of markets, including the markets for wholesale gasoline.[1] Related theoretical work has suggested a number of reasons for prices adjusting slowly or incompletely to cost or demand shocks, including menu costs,[2] long-term relationships between buyers and sellers that lead to non-price allocation mechanisms,[3] search costs,[4] and production stickiness.[5]

In this paper, we explore the combined effects of production stickiness and market power on the adjustment of wholesale gasoline prices to input price shocks. We begin by developing a theory closely related to Pindyck (1994) in which production stickiness, in conjunction with a declining marginal benefit of holding inventory, makes slow price adjustment to cost shocks profit-maximizing for perfectly competitive firms. When the shock occurs, each firm wants to change its output, but adjusting production is costly, so some of the change in sales is accomplished by selling off inventory when input prices fall and storing output when input prices rise. A firm adjusts production and inventory until the marginal value of inventory equals the marginal cost of adjusting production this period (versus adjusting next period), and both are equal to the (expected) inter-period

---

[1] See Borenstein, Cameron and Gilbert (1995), General Accounting Office (1993), Karrenbrock (1991) and Bacon (1991) for studies of lags in the adjustments of petroleum product prices.

[2] See Rotemberg (1982), Mankiw (1985), and Ball and Mankiw (1994).

[3] See Carlton (1986, 1991).

[4] See Benabou and Gertner(1993).

[5] See Thurman (1988) and Pindyck (1993, 1994).

difference in prices. When marginal production adjustment costs and the net marginal value of inventory are not equal to zero, the aggregate behavior of competitive firms leads to prices that do not immediately adjust fully to the cost shock.

We then demonstrate that in the same circumstances, a firm with market power will also adjust price slowly, but will have a different price adjustment path than would a competitive firm. The intuition is simple: competitive firms adjust production and inventory levels until marginal production adjustment and net inventory costs equal the intertemporal difference in *prices*. A firm with market power, however, will adjust production and inventory levels to equate marginal production adjustment and net inventory costs to the intertemporal difference in ***marginal revenue***.

We empirically test the prediction that sticky production leads to sticky prices in competitive markets by examining prices in a market in which menu costs, long-term relationships and search costs do not affect price: the futures market for gasoline. Our model of production stickiness and inventories implies incomplete adjustment of the futures price for gasoline to changes in the futures price of crude oil when the contract is near its delivery date. When the gasoline futures contract is not close to its delivery date, its price will respond immediately and completely to changes in the price of the crude oil contract with the same expiration date. Because traders expect the change in crude futures price to be fully incorporated into the spot gasoline price by the contract delivery date, trading immediately and fully incorporates the change into the futures price. However, as the contracts approach the delivery date, costly adjustment of production and inventory becomes relevant. The theory we present implies that shocks to crude oil price that occur near the delivery date of a give gasoline contract will not be fully passed through into the price of that contract. As a result, the futures price of gasoline will adjust incompletely. This prediction is confirmed by the empirical analysis.

We empirically test for a link between adjustment speed and market power using data from 188 local wholesale gasoline markets. We find that in markets with higher price-cost margins, wholesale gasoline prices adjust more slowly to crude oil price shocks. This pattern holds for both branded and unbranded gasoline. We also show that the prices of branded gasoline, for which product differentiation provide relatively more market power, respond more slowly than the prices of unbranded gasoline within the same local markets.

The remainder of the paper is organized as follows: Section II develops a model of price adjustment when adjusting production is costly and changing inventory levels affects the cost of distribution. Section III presents a test of costly adjustment in the gasoline futures market. In Section IV, we test for a relationship between adjustment speed and market power using data from local wholesale markets. Some concluding comments are offered in section V.

II. A Model of Production, Inventory, and Price Adjustment

This research is motivated by the observation that wholesale gasoline prices respond slowly to crude oil price shocks. These prices – called "terminal" prices because they are for transactions occurring at the terminals where refiners sell gasoline to local wholesalers – are formed in well-organized, relatively frictionless markets. Buyers and sellers are well informed about price and quality. Prices are available on-line from private data firms and the products' specifications are well-established. The flow of sales is continuous and smooth, and rationing is extremely rare. The primary stochastic variable in production costs – the price of crude oil – is public knowledge and established in well-organized markets. Even so, it is well known in the industry and has been shown in prior empirical work that gasoline prices at terminals respond slowly to crude oil price shocks.[6]

Given the nature of the market, these lags are initially surprising. A change in crude prices changes either the direct cost or the opportunity cost of the primary input and under most standard models of firm behavior would lead to an immediate change in the equilibrium price. Consider, for instance, a competitive refiner who realizes that the price of crude oil has

---

[6] See, for example, Borenstein, Cameron, and Gilbert, 1995.

increased by an amount sufficient to cause a 5¢ increase in gasoline prices in the long run. If the firm had been producing where short-run marginal cost was equal to price, then it's marginal cost is now above price and it has an incentive to reduce production. If production can be costlessly adjusted, this refiner, and all others, will cut back production immediately, reducing industry output sufficiently to increase the price of gasoline immediately by 5¢. Even if production cannot be adjusted immediately, a refiner has the option of putting its output into inventory. If increasing inventory had no effects on cost, the refiner would do so, holding output until the price adjusts. Since all refiners have a similar incentive, each would withhold supply and the price would increase immediately by 5¢.

Perfect competition may not be the best characterization of these terminal markets. The number of sellers at any given terminal is less than the "many" sellers in perfectly competitive models and a substantial share of gasoline sold in the U.S. is branded, and therefore not a homogeneous product. Introducing market power, however, does not immediately suggest that prices will adjust slowly. A monopoly firm faced with a change in its costs would want to change production immediately to achieve the new profit-maximizing price. If changing production immediately is prohibitively costly but the firm can costlessly change its inventory level, then it will immediately adjust the quantity it sells by increasing or decreasing its inventory.

This line of argument suggests that price adjustment lags might be attributed to costs associated with adjusting production and inventory levels. To understand how these effects can cause lags even when markets are efficient, information is perfect, and trade is anonymous, we develop a simple multiperiod model. In the model, we focus exclusively on adjustment to unanticipated cost shocks. In fact, refiners are responding to demand and supply changes. Some of these movements are anticipated: there is, for example, a seasonal demand cycle for gasoline. Others are unanticipated shocks: movements in the world price of crude oil cause unanticipated cost shocks. To simplify the exposition, we assume that that demand is constant and that there are no anticipated cost changes.

The timing of the model is as follows. At the beginning of each period, the firm has some level of inventory carried over from the previous period. Before it makes its production decision for the current period, it realizes a shock to its marginal cost of production. Having realized the cost shock, the firm decides how much it wants to sell this period. Implementing the sales decision requires deciding how much to produce this period and how much inventory it wants to carry over into the next period. When it makes its production decision, it takes into account the fact that adjusting production levels increases the cost of production. When it decides on how much inventory to hold, it takes into account the effect of inventory on the cost of distribution.

We include production adjustment costs to capture the idea that production is sticky: altering the production level from period-to-period imposes costs on the firm that lead it to adjust production less than completely in a single period. In particular, we assume quadratic adjustment costs. The assumption of costly production adjustment is clearly appropriate for the petroleum refining industry, although the actual problem facing refiners is much more complicated than suggested by this model. Refinery production schedules are set by solving a complex algorithm that takes into account the demand for the various refined products and the cost and quality of the crude oil input. Refineries operate most efficiently when the product and input mixes are constant. Adjusting output proportions and level is costly. As a result, although refinery output might be tweaked slightly when input prices change, refiners spread substantive adjustments over time, coordinating changes in response to unforeseen shocks with scheduled changes. The time it takes for an individual refiner to reach full adjustment depends on the flexibility of the individual refinery and the stage of the refinery's current production "run," as well as the magnitude of the change.

For simplicity, we assume that the marginal cost of production is constant and that the cost shock affects only the marginal cost, not the cost of adjustment. Marginal production costs in period t are represented by $c_{t-1} + \mu_t$, where μ_t is the shock realized in period t. We assume μ is i.i.d., with mean zero.[7] The total cost of production in period t, therefore, is

$$C_t = (c_{t-1} + \mu_t)y_t + a(\Delta y_t)^2,$$

where y_t is the level of production in period t and $\Delta y_t = y_t - y_{t-1}$.

In addition to the cost of production, the firm bears the cost of distribution. Distribution costs are those associated with getting the product from the refinery to the consumer and include transportation, scheduling and inventory costs. Inventory has two effects on cost. On the one hand, holding inventory increases cost through the cost of storage. The cost of storage includes the cost of storage services and the interest cost of selling product tomorrow rather than today. For simplicity, we abstract from interest cost.[8] Let n_t be the level of inventory carried over from period t to period $t+1$. The cost of storage services in period t is, then, $s(n_t)$. Because marginal storage costs must increase as inventory approaches the limit of storage capacity, we assume $s_n > 0$ and $s_{nn} > 0$.

On the other hand, holding inventory reduces the cost of distribution by reducing the transportation and scheduling costs that would be required in the absence of inventory. In gasoline refining, the distribution system imposes a requirement that firms hold inventory to make sales. A given refiner typically sells gasoline in substantially more locations than it has refineries. As an example, we can consider a refiner with a single refinery on the Gulf Coast and sales throughout the Eastern U.S. To have product available for sale in, say, Newark, New Jersey next period, it must have product in transit to Newark this period. Gasoline is shipped long distances by common carrier, typically a pipeline. Depending on the weather and the demand for carrier services by other users, it can take anywhere from 14 to 22 days to ship gasoline from the Gulf Coast to Newark.[9] In each period,

---

[7] This is consistent with our empirical application in which costs shocks are shocks to the world price of crude oil and follow (approximately) a random walk.

[8] We assume that interest rates are zero in the model and in our empirical work. This has no substantive effect on the central predictions of the model and, given the very short time periods considered in the empirical work, no substantive effect on our estimates.

[9] See National Petroleum Council, 1989, p.45.

then, the refiner must have many days of supply in transit to Newark. Thus, the distribution system imposes a requirement that inventories (including product in pipelines) be non-zero.

In addition, because problems with scheduling and variations in transportation speed mean pipeline delivery dates are uncertain, refiners have an incentive to hold inventory at terminals, where wholesale sales are made, as a buffer against delivery delays. There also are "within period" demand shocks that might leave aggregate demand in that period unaffected but create a transitory mismatch between the flow of output and the flow of consumption. Similarly, local demand shocks affect the spatial distribution of demand and, therefore, the optimal distribution of output. Holding inventory at key distribution points minimizes the cost of geographic variation in demand.

Given the requirements of the distribution system, a refiner always holds some inventory and increasing its inventory reduces the probability that it will be unable to meet demand because delivery has been delayed or local demand is unexpectedly high.[10] We represent this benefit as a decline in distribution costs as inventory increases. This benefit, called "convenience yield" in the inventory literature, is given by $\phi(n_t, E_t q_{t+1})$, where q is sales and the subscript on the expectations operator indicates the period in which expectations are formed. For a given level of (expected) sales next period, distribution costs decline as inventory increases: $\phi_n < 0$. The negative of this derivative $(-\phi_n)$ is called "marginal convenience yield." As inventories increase, the marginal contribution to reducing the probability the firm will be unable to meet demand probably declines; we therefore assume $\phi_{nn} > 0$. Because it is the level of inventory relative to expected sales that determines the probability of inadequate supply, we assume $\phi_q > 0$ and $\phi_{qq} > 0$. These assumptions are consistent with with prior theoretical and empirical work (see, for example, Pindyck (1994), Thurman (1988) and Fama and French (1987)). We adopt the accounting convention that the firm bears the distribution costs for sales made in $t+1$

---

[10] Of course, a firm can always "meet demand" by raising its price. This expression is simply a shorthand for the value of being able to set the price and sell the quantity that would be profit-maximizing absent the distribution constraints.

during period t. To capture the fact that inventories are never drawn down to zero, we assume that $\lim_{n\to 0} \; -\phi_n(n,q) = \infty$.

Total distribution costs, then, are given by $\phi(n_t, E_t q_{t+1}) - s(n_t)$. The change in net distribution costs as a result of a change in inventory level $(-\phi_n - s_n)$ is called "net marginal convenience yield" or "NMCY." Combining the costs of production and distribution, the firm's total cost function is:

$$C_t = (c_{t-1} + \mu_t)y_t + a(\Delta y_t)^2 + \phi(n_t, E_t q_{t+1}) + s(n_t). \qquad [1]$$

In a world without cost shocks ($\mu_t = 0$ for all t), sales would be the same each period. Because there would be no need to adjust production, production adjustment costs would be zero. In addition, since the firm would hold inventory only to minimize distribution costs, the firm would adjust inventory level to make $NMCY = 0$ at all times. That is, given the optimal level of sales, the firm would choose the inventory level that equates the marginal convenience yield and the marginal cost of storage.

When a firm realizes a cost shock, it wants to adjust sales. If, for example, the period begins with a positive shock to oil prices, the firm will want to reduce its sales relative to the previous period. It will do so by reducing its production level (producing less this period than it produced the period before) and increasing its inventory (carrying forward more inventory into the next period than it carried into this period). For a price-taking firm the optimal adjustment to a cost shock realized in period t is found by maximizing the value function:

$$\begin{aligned} V_t = & E_t \sum_{j=0}^{\infty} [p_{t+j} q_{t+j} - (c_{t+j-1} + \mu_{t+j}) y_{t+j} - a(\Delta y_{t+j})^2 \\ & - \phi(n_{t+j}, q_{t+j+1}) - s(n_{t+j})]. \end{aligned} \qquad [2]$$

In each period, there is an adding up constraint: $q_k = y_k - n_k + n_{k-1}$. Rewriting the objective function to incorporate this constraint gives us:

$$\begin{aligned} V_t = & E_t \sum_{j=0}^{\infty} [p_{t+j}(y_{t+j} - n_{t+j} + n_{t+j-1}) - (c_{t+j-1} + \mu_{t+j}) y_{t+j} \\ & - a(\Delta y_{t+j})^2 - \phi(n_{t+j}, q_{t+j+1}) - s(n_{t+j})]. \end{aligned} \qquad [2']$$

The firm maximizes [2'] with respect to inventory and production.

The first order condition for y_t is obtained by holding n_t and all future $y's$ and $n's$ constant. It is

$$p_t - (c_{t-1} + \mu_t) - 2a(\Delta y_t - E_t \Delta y_{t+1}) = 0. \qquad [3a]$$

To interpret this condition, note that changing y_t while holding n_t and all future $n's$ and $y's$ constant implies that the change in production equals the change in sales. Condition [3a] defines the optimal tradeoff between the revenue from selling an additional unit this period and the cost of producing an additional unit this period. The marginal cost of production includes the realized marginal cost plus the marginal adjustment cost. This latter term is the change in adjustment costs from, for example, increasing production by one more unit this period ($2a\Delta y_t$) rather than making the adjustment next period ($2aE_t\Delta y_{t+1}$). Notice that in the absence of adjustment costs ($a = 0$), [3a] is simply the standard condition that a competitive firm adjusts its sales until its marginal cost of production equals the market price.

The first order condition for inventories is obtained by differentiating with respect to n_t holding y_t and all future $y's$ and $n's$ constant. It is

$$-p_t + E_t p_{t+1} - \phi_n(n_t, E_t q_{t+1}) - s_n(n_t) = 0. \qquad [3b]$$

This condition defines the optimal trade-off between the cost and benefit of holding an additional unit of inventory, taking into account possible changes in output prices. As noted above, the net marginal benefit of holding inventory is the NMCY. When output prices are constant in expectation, ($E_k p_{k+1} = p_k$), the firm chooses the inventory level that makes the NMCY equal to zero. When output prices are expected to change, the price change introduces an additional cost: the cost of holding an incremental unit of inventory now includes the (expected) price next period minus the current price. If firms expect prices to be falling (rising) they will want to reduce (increase) inventory. If NMCY had been equal to zero in the previous period, expecting prices to fall (increase) will lead to inventory changes that make the current NMCY positive (negative)

Conditions [3a] and [3b] must hold for all periods j. In particular, there is a first order condition for y_{t+1} analogous to condition [3a]:

$$E_t p_{t+1} - (c_t + E_t \mu_{t+1}) - 2aE_t \Delta y_{t+1} + 2aE_t \Delta y_{t+2} = 0. \qquad [3a']$$

Jointly, conditions [3a], [3a'] and [3b] imply the following equalities:

$$\begin{aligned} p_t - E_t p_{t+1} &= -\phi_n(n_t, E_t q_{t+1}) - s_n(n_t) \\ &= 2a(\Delta y_t - 2E_t \Delta y_{t+1} + E_t \Delta y_{t+2}) + E_t \mu_{t+1}, \end{aligned} \qquad [4]$$

where $E_t \mu_{t+1} = 0$. Optimal response to an input cost shock requires that the expected difference in prices will equal both the net marginal convenience yield and the marginal adjustment cost.

Note that condition [4] implies that instantaneous adjustment of output prices , *i.e.*, $p_t = E_t p_{t+1}$, would occur when marginal production adjustment costs and NMCY are each equal to zero. This could occur if $a = 0$. In that case, production adjustment by each firm would be instantaneous. As a result, prices would adjust instantaneously. Furthermore, firms could costlessly adjust production so that the inventory carried forward into the next period minimized (expected) distribution costs, implying that NMCY would equal zero. [4] could also hold with instantaneous output price adjustment if the net marginal convenience yield were zero at all levels of inventory. In that case, the firm would absorb all cost shocks through inventory changes. It would never alter production.[11]

These counterfactuals make clear that lagged adjustment requires both sticky production and an NMCY that is a declining function of inventory and an increasing function of expected sales.[12] In this case, a positive cost shock leads a firm to reduce production and to increase inventories in period t. As each of the competitive firms undertakes these changes, industry

---

[11] Clearly, this is not reasonable since there are non-negativity and capacity constraints on inventory. If NMCY were fairly flat, however, the firm would adjust to a shock rapidly through inventory changes and then would very gradually change production levels, so as to minimize adjustment costs.

[12] As noted above, these functional form assumptions for the NMCY are consistent with prior research. In the appendix to this paper, we present evidence that NMCY in wholesale gasoline declines in the ratio of inventory to sales.

supply contracts, pushing up the current price. Each firm reduces output and increases inventory until the expected difference in prices $(p_t - E_t p_{t+1})$ equals its NMCY and its interperiod difference in marginal adjustment cost. Adjustment becomes more rapid as the NMCY is less responsive to changes in inventories and expected sales or adjustment costs are smaller. The speed of adjustment will also depend on the change in quantity necessary for price adjustment to occur. If demand were very inelastic, then only small changes in production or inventory would be necessary for price to adjust fully. More elastic industry demand will lead to slower adjustment, *i.e.*, the more elastic the industry demand, the lower the proportion of the eventual passthrough that is accomplished in the first period.

While the model thus far has focused on a perfectly competitive industry, perfect competition may not be the right model for wholesale gasoline transactions at a given terminal. In section IV, we provide institutional details that suggest refiners have some market power. For now, we want only to determine how market power might affect our analysis. To explore the effects of market power on the adjustment path of prices, we take the simple approach of modeling the monopoly solution. We comment below on how this simplification might affect the results of the analysis.

The monopoly problem is a simple variation on the competitive problem. The optimization problem is essentially the same as in [2], except p_t is now a function of q_t, *i.e.*, the firm is not a price taker. The analogs to [3a] and [3b] are given by:

$$p_t + \frac{\partial p_t}{\partial q_t} q_t - (c_{t-1} + \mu_t) - 2a\Delta y_t + 2aE_t\Delta y_{t+1} = 0, \qquad [5a]$$

and

$$-p_t - \frac{\partial p_t}{\partial q_t} q_t + E_t[p_{t+1} + \frac{\partial p_{t+1}}{\partial q_{t+1}} q_{t+1}] - \phi_n(n_t, E_t q_{t+1}) - s_n(n_t) = 0. \quad [5b]$$

The first order condition with respect to this period's production [5a] differs from [3a] only because marginal revenue is not equal to price for the monopolist. Similarly, the first order condition for inventories ([5b]) equates NMCY with the (expected) difference in marginal revenues rather than prices.

Two things are immediately clear. First, like the competitive firm, a monopoly will adjust prices gradually. Second, the rate of adjustment will not, in general, be the same for a monopoly firm as it would be for a competitive firm facing the same shock. In competitive markets, the intertemporal difference in prices in periods t and $t+1$ is equated with the NMCY; in monopoly markets the intertemporal difference in marginal revenues is equated with the NMCY.

But while the adjustment speeds will differ, the theory does not allow us to predict whether a monopolist would adjust more quickly or more slowly than a competitive industry without more information on the industry demand curve. Demand affects adjustment rates in much the same way as it affects the proportion of a cost increase that a monopolist will pass through: passthrough can be greater or smaller with monopoly than with competition depending on the second derivative of the demand curve. Compared to a competitive firm, a monopolist facing a linear demand curve will adjust price less, but a monopolist facing a constant elasticity demand curve will adjust price more. In the first case, marginal revenue is steeper than the demand curve and in the second it is flatter. Similarly, a monopolist adjusting through inventory will equate the $NMCY$ to the difference in marginal revenue. Whether that implies a period j adjustment greater or less than if the firm had equated $NMCY$ to the difference in price depends on the relative slopes of marginal revenue and demand. Thus, while the theory tells us that the price paths will be different, the sign of the difference is an open empirical question.

This analysis has assumed that it is impossible for competitive arbitrageurs to buy and sell the monopolist's product. The assumption matters only if the monopoly adjustment is slower than the speed a competitive arbitrageur would impose (meaning, for a cost increase, that the monopoly $E_t p_{t+1} - p_t$ is greater than the net marginal convenience yield of competitive arbitrageurs). If arbitrage were possible in this case, then arbitrage will move prices in the monopoly market more quickly than the monopolist would like. In fact, branded refiners report that they attempt to limit arbitrage at city terminals. When crude oil prices rise very rapidly, refiners may ration purchases. In particular, they might limit the purchases made by independent resellers.[13] That refiners attempt to control arbitrage is in itself evidence that they prefer a different and slower rate than competitive arbitrage would imply. If their attempts are at least partially successful, then they will be able to maintain a slower adjustment rate.

The monopoly model assumes there are no close substitutes for a refiner's product. In our application, however, the gasoline offered by any single refiner faces competition from the product of other refiners. Thus, all products have close, albeit imperfect, substitutes. Because consumers will substitute among refiners' products as relative prices change, a refiner adjusting the terminal price of its product will take into account the price paths of competing refiners. This will tend to reduce the difference in output price adjustment paths. In the extreme case of perfect substitutes, no firm has market power and the price adjustment path will be the path consistent with perfect competition. As the substitutes become more imperfect, the cross-price elasticities decline and the resulting price adjustment path will more closely resemble that predicted by the monopoly model than that predicted by the competitive model.

This suggests two empirical tests of the model's predictions about the effect of market power. First, if competitive conditions differ across markets, adjustment speeds will differ across markets. Markets where competition is less intense will have price paths closer to the monopoly path than markets where competition more nearly approaches perfect competition. Second, to the extent that sellers of branded product have more market power than sellers of unbranded product, the price paths of branded and unbranded gasoline within the same market should be different. We test these predictions in the empirical work.

III. Evidence of Costly Adjustment from Futures Markets

The theory that we have presented could explain the lagged adjustment of gasoline prices to crude oil price changes in the absence of menu costs,

---

[13] As described in the section IV, refiners sell gasoline at the terminal to independent wholesalers who truck the gasoline to stations, but who also maintain some storage capacity. Because they have storage capacity, they might be able to arbitrage the price adjustment process. See Borenstein and Gilbert (1993) for details.

long-term relationships, or imperfect information. To examine this theory empirically, we begin by studying the response of price of gasoline futures contracts to changes in the price of the futures contract for crude oil.

Because futures markets are very efficient, the prices in these markets will not be affected by menu costs, long-term relationships between buyers and sellers, or imperfect information. Futures prices are an unbiased predictor of the market clearing price for the contract delivery date. In the absence of market stickiness, the (expected) market clearing price will equal the (expected) spot price. If market stickiness prevents the spot price from equating demand and supply, however, the (expected) market clearing price will be formed in the futures market and will not equal the spot price. Consider, for example, a situation in which the price of the crude oil contract rises and traders know that the increase in crude oil prices cannot be passed through to spot prices at the contract delivery date because menu costs, for instance, prevent spot price adjustment. If all transactions must occur in the spot market, there would be excess demand for gasoline. Instead, however, traders will transact in the futures market so that the futures price of gasoline equals the market clearing price, which will be above the spot price in this case. The important point here is that the futures price of gasoline will fully and immediately incorporate the cost shock even if the (expected) spot price does not. Spot market stickiness (due to menu costs or long-term relationships) cannot slow or reduce the responsiveness of gasoline futures prices.

The situation is quite different, however, when there is production stickiness rather than market stickiness. In this case, the sticky spot price is the market clearing price, *i.e.*, the price that equates expected demand and expected, sticky supply. If the observed lag in spot price adjustment is the result of production stickiness as modeled in the previous section, the futures price of gasoline will also be sticky.

To be concrete, consider the contract for delivery of unleaded regular gasoline to the New York harbor during July. That contract is generally traded beginning about one year before the delivery date, as is the contract for July crude oil. Toward the end of June, traders close out their positions in both contracts. Now consider a shock to July crude oil that occurs in May or earlier. One would expect that the price of the July gasoline contract would adjust fully and rapidly (a day or two) to that information. Whatever gasoline production or inventory adjustments are dictated by the crude price changes could be implemented by July, so the price of the July gasoline contract should adjust fully and quickly. But now consider a shock occurring in June. If production stickiness causes a lagged adjustment in production, then the price of the July gasoline contract will not adjust *completely* to the change in the price of the July crude oil contract during June, though whatever adjustment occurs will still take place quickly. Furthermore, we would anticipate that the adjustment would be less complete for trading days later in June than for days earlier in June.

To test these implications of the theory, we use a dataset of New York Mercantile Exchange prices for futures contracts for New York harbor delivery of light sweet crude oil and unleaded regular gasoline. The dataset includes daily prices for contracts with a delivery month between December 1985 and January 1995, a total of 110 contracts and more than 21,000 trading day observations. The estimation is based on a simple lagged adjustment model. The dependent variable is the change in the price of the gasoline contract (from the close on day $t-1$ to the close on day t) for a given month. The independent variables are the change in the price of the same month's crude oil contract on the same trading day and on the two prior trading days. Changes over the preceding trading days are included to capture any lagged response that might occur if, for example, the change in the crude oil contract price occurs very late in the trading day. In that case, the corresponding change in the gasoline futures contract may not be completed on the same trading day.[14]

---

[14] Only two lags are included because with lags up to six days, none of the coefficients on lags 3 through 6 was significant and lags 3 through 6 were not jointly significant.

The estimating equation is:

$$\begin{aligned}\Delta GAS_{s,t} = {} & \beta_0 \\ & + \beta_1 \Delta CRUDE_{s,t} + \beta_2 \Delta CRUDE_{s,t-1} + \beta_3 \Delta CRUDE_{s,t-2} \\ & + D * (\gamma_1 \Delta CRUDE_{s,t} + \gamma_2 \Delta CRUDE_{s,t-1} + \gamma_3 \Delta CRUDE_{s,t-2}) \\ & + \epsilon_{s,t},\end{aligned} \qquad [6]$$

where s indexes the month of delivery specified in the contract, t indexes the trading day, $\Delta GAS_{s,t} = GAS_{s,t} - GAS_{s,t-1}$, $\Delta CRUDE_{s,t} = CRUDE_{s,t} - CRUDE_{s,t-1}$, and D is a dummy variable that equals one when trades occur in the last month before contract expiration. The theory implies very rapid adjustment for shocks that occur prior to the last month before contract expiration. For these shocks, the sum of the coefficients on crude contract price changes (*i.e.*, $\sum_{j=1}^{3} \beta_j$) is the estimated long-run passthrough rate. The sum of the γ coefficients are an estimate of the reduction in responsiveness that occurs for shocks that occur in the last month of contract trading. More precisely, because crude oil contracts stop trading on about the 22nd day of the month prior to delivery, this sum is the offset that applies to shocks occurring during the last 22 calendar days before the crude contract stops trading. The theory predicts that the sum of these coefficients should be negative. Because the theory also predicts that the stickiness in the futures price will increase for days late in the last month, we also estimate [6] setting $D = 1$ only after the 10th day of the month before delivery, about the last 10 calendar days. The sum of these coefficients should also be negative and should be larger in absolute value than the sum of the coefficients for the last 22 calendar days.

Two econometric issues arise in estimating [6]. First, crude oil prices might be endogenous. It is possible that demand shocks in the United States, for example, might both affect NYMEX futures prices for gasoline and create short run excess demand or supply in U.S. crude oil markets that would affect the NYMEX futures price of crude oil.[15] To test and correct for endogeneity, we estimate by two-stage least squares. The instruments we

---

[15] Endogeneity is suggested by occasional reports in the press that movements in the nearest gasoline futures contract affect the associated crude oil contract.

use are spot crude oil prices in England[16] and the 6-month-ahead futures price of crude oil on the NYMEX. Both should be relatively immune to transitory, North American demand shocks. A Hausman test using these instruments confirms the endogeneity.

The second econometric issue arises from the distribution of the error term. The reported standard errors have been adjusted for heteroskedasticity using White's method. There is also possible correlation of residuals of different contracts on the same trading day, *e.g.,* the observation of the July 1993 contract on April 18, 1993 and the observation of the June 1993 contract on April 18, 1993 are likely to have correlated errors. We have used Huber's adjustment for group correlation to correct the standard errors for this possible bias.

The results from estimating equation [6], reported in Table 1, support the production stickiness hypothesis. The first two columns report estimates when we allow different coefficients for crude price changes in the last month before delivery, and the last two columns report the estimates when we allow different coefficients for about the last ten calendar days.

Focussing on the 2SLS estimation, the results indicate that the price of the futures contract for a given month goes up by about 1.16 cents (the sum of the coefficients in the first three rows) for every one cent that the crude price goes up when the crude change occurs prior to the last month the contract is traded. The same one cent change in crude price that occurs in the last trading month has an effect that is about 0.14 cents smaller (the sum of the coefficients on the $D \cdot \Delta CRUDE$ variables). The difference is statistically significant at the 2% level. When we allow for different adjustment to changes that occur after the 10th day of the last trading month, the adjustment is even less complete, as the theory predicts. The 2SLS estimates imply that a one cent change in crude prices in the last ten trading days has a 0.20 cents smaller effect on gasoline prices than changes occurring earlier in the contract life, which is statistically different from zero at the 1% level.

---

[16] These are the same-day spot prices for Brent and North Sea Forties crude oil.

Some of the earlier work motivating this study demonstrated not only that wholesale gasoline prices respond with a lag to crude oil price changes, but also that the speed of response is asymmetric. Decreases in crude prices are passed along more slowly than increases. If this phenomenon also existed in the commodity markets, it would manifest as less complete adjustment to decreases than to increases when the gasoline contract is near expiration. We tested for asymmetry by allowing the adjustments to crude price changes occurring near the contract expiration to differ for crude price increases and decreases. The magnitudes of the parameter estimates (not reported) suggest less complete adjustment for crude contract price decreases than for increases, but the difference was never statistically significant.

IV: Lagged Price Adjustment and Market Structure

To explore the relationship between lagged responses to cost shocks and market structure, we turn to data on adjustment of wholesale prices in local markets to crude oil price changes. There are approximately 370 terminals in the U.S. at which refiners maintain facilities for selling gasoline to resellers. Each terminal serves a local geographic area determined by the economics of trucking gasoline from the terminal to gasoline stations. Sales at the terminal are made to independent wholesalers who deliver and resell the product to gasoline stations. A wholesaler may store some portion of its purchases in its holding tanks prior to resale. Refiners also truck gasoline from the terminal to gasoline stations.

The number of refiners supplying product to a given terminal ranges from as few as 5 to more than 20. Some are major brand refiners – such as Shell, Chevron, or Exxon – while others are less well-known refiners with little or no retail presence – such as Tosco, Crown, or Hill. The product of the less well known refiners is sold to consumers through "unbranded" service stations, *i.e.*, stations not carrying a major brand name. Stations supplied directly by a branded refiner sell branded gasoline. Some "branded" stations are supplied by independent wholesalers who purchase gasoline sold at the terminal under the brand name. A Chevron station, for example, supplied by an independent wholesaler resells gasoline sold at the terminal as "Chevron" gasoline. Wholesalers usually supply some branded and some unbranded stations, and often own many stations themselves. A major refiner may sell unbranded as well as branded gasoline at the terminal, *e.g.,* Chevron might sell both "Chevron" gas to be resold at Chevron stations and unbranded gasoline that can be resold through any unbranded station.[17]

Refiners who sell at the terminal post a "terminal" price for wholesale transactions. If a refiner sells both branded and unbranded product, it will post a price for each. These are the prices at which independent wholesalers purchase gasoline from refiners. Transactions between refiners and the retail outlets they directly supply occur at a "dealer tankwagon" price that includes a delivery charge. Similarly, wholesalers charge the stations they serve a delivered price that is the terminal price plus some markup. While systematic data on terminal prices are available, there are no reliable data on delivered prices. We therefore use terminal prices in our empirical work.

The model developed in Section II indicates that a firms with market power will respond differently to a cost shock than would a perfectly competitive firm. One way this difference might be uncovered in data is through a cross-sectional comparison of terminal price responses to a crude oil price shock. Terminals where the firms exercise greater market power should have a different price response path than terminals where margins are lower. The model predicts that price response paths will vary systematically with variation in market power, *i.e.,* with variation in price-cost margins.

A second way the relationship between market power and adjustment rates might become apparent is in a comparison of the price paths of branded versus unbranded gasoline at a given terminal. Because the product sold at unbranded stations is not identified with an individual refiner, wholesalers are free to purchase unbranded gasoline from any seller of unbranded product. In addition, consumers of unbranded product are prob-

---

[17] Branded gasoline contains the additive package associated with the brand. Unbranded gasoline need not, and often does not, contain the same additive package.

ably more prone to switching stations. These two facts imply that, at a given terminal, sellers of branded product are likely to exercise more market power than sellers of unbranded product at the same terminal. The greater the difference in market power, the greater the difference we expect to see in the price adjustment paths. To test this prediction we regress the difference in price adjustment rates on the difference in price-cost margins.

For both tests, we begin by estimating the adjustment speed of branded and unbranded terminal prices on a city-by-city basis. Descriptive statistics on the variables used appear in Table 2. The terminal price data are the average prices for 87 octane, unleaded gasoline on each Friday at the given terminal.[18] The data include separate averages for branded and unbranded product. We match these price observations to the Friday gulf coast spot price for West Texas Intermediate (WTI) crude oil.[19] The dataset contains weekly observations for the period January 1, 1986 through November 20, 1992 at 188 terminals. We restrict the estimation to terminals at which we have a sufficient number of observations and are located in petroleum administration defense districts (PADDs) I, II and III, which include all areas of the continental U.S. east of the Rocky Mountains. These PADDs are most clearly in the same crude and gasoline distribution system.[20]

We detrend and deseasonalize these data to remove the effects of inflation and any seasonal pattern. There is a marked seasonal pattern in

---

[18] The terminal price data used here are collected and reported by Lundberg Survey, Inc.

[19] The spot price of West Texas Intermediate (and other crudes) is constructed from a survey of traders conducted daily by Dow Jones International Petroleum Report.

[20] PADD V is the west coast, a region that refines primarily Alaska North Slope crude oil and does not trade significant quantities of gasoline with the other four PADDs. PADD IV, which includes Colorado, Idaho, Montana, Utah and Wyoming, trades small quantities of refined products with PADD II, but it has no significant pipeline links and is largely self-sufficient. In contrast, PADDs I, II and III are closely linked by major pipeline and water transportation for both crude and gasoline products. The economic isolation of gasoline and crude markets for PADDs IV and V and the integration of PADDs I, II, and III indicated by physical distribution systems has been confirmed empirically by Slade (1986). We do not have complete time series for each terminal in these three PADDs. A terminal is included only if we have at least 200 usable observations, *i.e.*, with sufficient lagged price information to include the observation in the data.

the demand for refined petroleum products that is reflected in price of gasoline.[21] Detrending and deseasonalizing the data is consistent with assuming that all prices fully and immediately adjust to these effects. This seems reasonable since seasonal patterns and calendar time can be fully anticipated.[22]

There are many econometric approaches to estimating response lags, and our estimates suggest that the basic findings are robust to the technique employed. We report below the results from two models: a partial adjustment model (PAM) and a vector autoregression (VAR) model.[23] A partial adjustment model is attractive because it generates a single parameter estimate of the adjustment rate. The structure imposed by the model tends to yield reasonably precise estimates. However, the structure also constrains price adjustments to be of equal proportions in all adjustment periods, a restriction at odds with prior research.[24]

The PAM we estimate for each terminal is:

$$\Delta TERM_t = \delta(TERM_t^* - TERM_{t-1}) + \mu_t \qquad [7]$$

where $TERM_t^*$ is the target level for the terminal price in period t, *i.e.*, the terminal price that would be observed in the long run if crude prices remained at their period t level. $TERM_t^*$ is the predicted value from the regression

$$TERM_t = \alpha_0 + \alpha_1 CRUDE_t + \epsilon_t. \qquad [7']$$

---

[21] Gasoline demand is higher in periods of higher leisure travel, summer for most cities. Heating oil is in higher demand in the winter months. Since gasoline and heating oil are joint products of the refining process, the interaction of the demand cycles causes seasonal patterns in inventories and prices for both products.

[22] We remove these effects from the crude series by regressing weekly crude prices on time and 52 weekly seasonal effects. These effects are removed from terminal prices by regressing weekly terminal prices on time and 52 seasonal effects on a city-by-city basis to allow for different seasonal patterns across terminals. We have also carried out the analysis by including seasonal effects and a time trend in the regressions, with virtually identical results.

[23] We also estimated a lagged adjustment model, regressing current change in terminal price on current and lagged changes in crude oil price and got results substantively similar to those reported below.

[24] See Borenstein, Cameron, and Gilbert.

The long-run adjustment proportion – or passthrough rate – is α_1, a parameter that prior research suggests should be near one.[25]

As noted in the previous section, prior research has suggested that the price adjustment path may be asymmetric with respect to the sign of the crude price change. In particular, the response of terminal prices to crude oil price increases may be faster than the response to crude oil price decreases. To insure that our estimates of market structure effects are not driven by pooling different price response paths, we also estimate a version of [7] that allows for an asymmetric response. The first step is to estimate the long-run relationship *i.e.*, $TERM_t^*$ as before. The asymmetric PAM regression is then:

$$\Delta TERM_t = \delta_1^+(TERM_t^* - TERM_{t-1})^+ + \delta_2^-(TERM_t^* - TERM_{t-1})^- + \mu_t. \qquad [8]$$

The + (−) superscript indicates that the target minus the lagged terminal price is positive (negative), *i.e.*, that terminal prices are increasing (decreasing).[26]

Compared to the PAM, the VAR model imposes fewer restrictions on the data. In particular adjustment rates can change over the adjustment period. The cost of this freedom is a set of estimated cumulative response rates.[27] The VAR model we estimate is:

$$\begin{aligned}\Delta TERM_t = &\sum_{j=0}^{5} \theta_j \Delta CRUDE_{t-j} + \sum_{k=1}^{5} \phi_k \Delta TERM_{t-k} \\ &- \gamma(TERM_{t-1}^* - TERM_{t-1}) + \mu_t\end{aligned} \qquad [9]$$

---

[25] We use this two step estimation procedure in which we first estimate the long-run relationship between crude and terminal prices using [7'] and use the predicted value in the partial adjustment model because it is useful when we estimate the asymmetric model discussed below. For the symmetric model in [7], it is possible to do the estimation in a single step. Doing so has no substantive effect on the results.

[26] To be more precise, $(TERM_t^* - TERM_{t-1})^+ = \max\{0, TERM_t^* - TERM_{t-1}\}$ and $(TERM_t^* - TERM_{t-1})^- = \min\{0, TERM_t^* - TERM_{t-1}\}$

[27] For a description of calculating cumulative response functions from a VAR model, see Borenstein, Cameron and Gilbert.

This VAR model allows us to calculate the cumulative response of terminal prices for the weeks following a cost shock.[28] Analogous to the two step estimation process for the PAM, the error correction term in the VAR, $TERM^*_{t-1} - TERM_{t-1}$, contains the target terminal price level that is the (lag of) the predicted values from the regression:

$$TERM_t = \alpha_0 + \alpha_1 CRUDE_t + \epsilon_t. \qquad [9']$$

The asymmetric version of the VAR uses the same long-run estimates for the error correction term, but differentiates between positive and negative price changes. The asymmetric model is:

$$\begin{aligned} \Delta TERM_t = & \sum_{j=0}^{5} [\theta_j^+ \Delta CRUDE_{t-j}^+ + \theta_j^- \Delta CRUDE_{t-j}^-] \\ & + \sum_{k=1}^{5} [\phi_k^+ \Delta TERM_{t-k}^+ + \phi_k^- \Delta TERM_{t-k}^-] \qquad [10] \\ & - \gamma (TERM^*_{t-1} - TERM_{t-1}) + \mu_t \end{aligned}$$

where $\Delta X^+ = \max\{0, \Delta X\}$ and $\Delta X^- = \min\{0, \Delta X\}$.

Summary statistics for the estimated adjustment rates are reported in Tables 3A and 3B. All estimates are from a 2SLS procedure in which we instrument for West Texas Intermediate crude prices ($CRUDE$) as described in the previous section. Table 3A reports the summary statistics for the city-by-city adjustment rate estimates from the PAM. Table 3B reports the summary statistics for the city-by-city estimates of the weekly, normalized cumulative response rates from the VAR model. The cumulative response functions for a VAR report the response as a proportion of the underlying shock. Thus, these functions record the cumulative movement of terminal prices as of, for example, the end of the second week as a proportion of the movement in crude prices. To make these comparable to the adjustment rates for the PAM, we divide by the estimated estimated long-run adjustment to the crude shock. This gives us a rate that is the ratio of

---

[28] We truncate the estimation at six weeks (0 through 5) because coefficients on longer lags were insignificant in all but a few cities.

cumulative terminal price adjustment through the *n*th to the full terminal price adjustment.

As reported in Table 3A, the average adjustment rate from the symmetric PAM model is 18% for branded and 21% for unbranded terminal prices. Consistent with lagged adjustment, these rates are well below unity. The difference in the branded and unbranded adjustment rates is significantly different from zero at the 1% level. A 21% (18%) adjustment rate implies a passthrough of about 76% (70%) at the end of six weeks. Allowing for asymmetry in price responses has no substantive effect on the parameter estimates. These results do not suggest that the response of branded gasoline price is asymmetric, but give some indications that the price of unbranded product adjusts more quickly to increases in crude oil prices than to decreases. Although the difference in the mean adjustment rates of unbranded gasoline to increases versus decreases is not large, it is significantly different from zero at the 1% level.

The average cumulative adjustment rates for each week from the VAR model are reported in Table 3B for the first three weeks. Like the PAM estimates, the VAR results show faster adjustment for unbranded than branded prices. The estimates, however, suggests a faster adjustment than is implicit in the PAM estimates. In the symmetric model, adjustment of both branded and unbranded prices appears to be complete by the end of three weeks. In the VAR, there also is much stronger evidence of an asymmetry between price increases and decreases. Both branded and unbranded prices increase more quickly than they decline.

Using the adjustment rates summarized in these tables, we can test for market structure effects across terminals. As noted above, we want to test for a relationship between price-cost margins and adjustment rates. To do this, we regress average adjustment rates on average prices (for branded and unbranded products) and measures of marginal cost. A major component of marginal cost is, of course, the cost of crude oil. But since we are interested in identifying cross-sectional variation, we need to include only measures of the components of marginal cost that vary across terminals.

Perhaps the largest of these is transportation. The cost of selling gasoline at a given terminal will increase with the cost incurred to transport the product from the refinery to the terminal. We therefore include three variables that proxy differences in transportation. A large share of the gasoline in the three PADDs included in our sample is transported by pipeline. Terminals with a pipeline link should be served at lower cost than those for which gasoline must be trucked from the pipeline to the terminal. We therefore construct a dummy variable "Off pipeline" which takes on a value of one when the city is more than 50 miles from the nearest pipeline. We also construct a variable "Distance from Refinery" that estimates the distance from the terminal to the closest refinery. We also include PADD dummies, which allow separate intercept estimates for the east coast (PADD I), the upper midwest (PADD II), and the middle south (PADD III). PADD I, for example, is served primarily by major pipelines coming up from the gulf coast. The PADD dummies could, therefore, capture gross differences in transportation costs.

In addition to these cost variables, we include data on the number of refiners offering product at the terminal. Many standard oligopoly models suggest that competition increases with the number of competitors. Because we observe cost imperfectly, including "Number of Refiners" in the regressions may help to identify market power effects. Controlling for price, the effect of number of competitors on adjustment rates should capture unobserved differences in cost. For example, consider two terminals with the same price for branded gasoline and the same values for the direct cost indicators, but different numbers of competitors. It is reasonable to infer that the terminal with more competitors has (unobserved) higher marginal costs than the terminal with fewer competitors. If this is correct, we would expect the number of competitors, to have a coefficient of the same sign as the direct cost indicators. Holding price constant, higher costs imply lower margins and lower market power. Our data on the number of competitors comes from a single cross-section in 1990.[29]

---

[29] We thank Bob Town for supplying us with these data. We use the total number of companies selling at the terminal for the "Number of Refiners" variable, whether a company is selling branded, unbranded, or both. Breaking out the variable in whether the fuel was branded did not increase explanatory power.

The results from the symmetric models are reported in Table 4A (columns 1 and 2) for the PAM and in Table 4B (columns 1 and 2) for the VAR model. To reduce the number of reported regressions, we present the VAR results using only the adjustment rate after two weeks. We want to focus on a period in which adjustment lags are evident – *i.e.*, before adjustment is complete – which argues for using no more than the first three weeks. We have chosen to report results for week two. Using only the first week leads to similar results.

The parameter estimates from both symmetric models indicate that higher price-cost margins are associated with slower adjustment. Using the estimates from the partial adjustment model (Table 4A), the branded adjustment rate increases in cost and declines in price. A one cent increase in the price of branded gasoline (holding cost constant) results in a 0.5 percentage point decline in the adjustment rate. Holding price constant, higher transportation costs or more competitors reduces the price-cost margin and increases the adjustment rate. Adjustment rates for unbranded products also are affected by price levels, and the magnitude of the response is approximately the same as for branded products. The cost variables have the same signs as for branded product as well, but number of competitors and off the pipeline are not significantly different from zero. In general, the regression for unbranded adjustment rates yields noisier parameter estimates; it also explains only about half as much of the variation in adjustment rates as does the branded gasoline regression. The symmetric VAR estimates (Table 4B) tell approximately the same story: adjustment rates decline in price and increase in cost and the number of competitors. The effects of price and all cost indicators are statistically significant. The coefficients on the price variables imply that a one cent increase in the wholesale price lowers the second week cumulative adjustment (which has a mean of 0.89) by about 0.02.

The evidence from the asymmetric models is slightly more mixed, but generally confirms the implications from the symmetric models. The regressions using the VAR estimates (Table 5B) produce consistently negative coefficients on price and positive coefficients on cost and number of competitors for both increases (columns 1 and 2) and decreases (columns

4 and 5). The coefficients on the price variables are significantly different from zero at the 5% level or better and most of the other coefficients are also precisely estimated. The asymmetric PAM estimates (Table 5A) support the market power hypothesis less strongly. While the results for price increases (columns 1 and 2) are consistent with higher margins being associated with slower adjustment, the results from price decreases (columns 4 and 5) are not. In particular, the response of unbranded prices to a cost decrease increases with costs, but it also appears to increase with price.

The bulk of the evidence suggests that higher margins are associated with slower adjustment to cost shocks. This is consistent with our model in that more market power leads to different price paths. It is also consistent with the anecdotal evidence suggesting that brnaded prices adjust more slowly. As mentioned in Section II, branded refiners report attempting to limit arbitrage by wholesalers who purchase at the terminal and have the capacity to store gasoline. Refiners would want to prevent arbitrage of this sort only if they wanted to slow down the rate of the price response.

Our second test of the market power hypothesis looks at the difference in adjustment rates between branded and unbranded products at the same terminal. If branded sellers have more market power than unbranded, the preceding results suggest that they will want to adjust prices more slowly. Looking at prices at the same terminal has the advantage of eliminating most cost differences. Because our cost variables in the "levels" regressions above controlled for cost imperfectly, we potentially measured price-cost margins with error. In this test, however, we regress the *difference* between branded and unbranded adjustment rates at a terminal on the *difference* in branded and unbranded average prices. This differencing removes the effects of omitted cost variables that are the same for branded and unbranded gasoline terminal, which is probably most of the cost variation. It also eliminates the impact of the cost variables we had previously included.[30]

---

[30] More precisely, differencing removes these effects if they have the same coefficients in the branded as in the unbranded "levels" regressions. Since some coefficients are different, we also ran this regression including the cost variables. The price results were unaffected and none of the included cost variables had a statistically significant effect.

The results of this approach strongly support the conclusion that the branded adjustment rate is low relative to the unbranded rate where the branded price is high relative to the unbranded price. The coefficient on the price difference variable is significantly different from zero for both of the symmetric adjustment regressions (at 5% for PAM and 1% for VAR). The estimate in the PAM difference regression implies that a one cent increase in average branded relative to average unbranded price in the city leads to a 0.0138 decrease in the adjustment speed parameter for branded relative to unbranded. Likewise, the VAR result implies that a one cent increase in average branded relative to average unbranded price leads to the cumulative adjustment of branded being 0.027 lower relative to unbranded in the second week after the crude price change. In the asymmetric models, the coefficients from the PAM are significantly different from zero at the 5% level for cost increases and at the 10% level for price decreases. The VAR coefficients are significantly different from zero for price increases at the 10% level and for price decreases at the 1% level.[31]

The results of both the levels regressions and the difference regressions are consistent with lags caused by production stickiness and mediated by market power. Menu costs are not a likely explanation for the terminal price stickiness we observe: Terminal prices are generally reviewed by refiners on a daily basis and are changed many times each week.[32] Furthermore, menu costs could not reasonably be higher for branded than for unbranded product. Some of the results might, however, reflect long-term buyer/supplier

---

[31] The theory of costly production and inventory adjustment also has implications for the response of prices to demand shocks. For instance, heavy rain on a holiday weekend would lead to unexpectedly low sales and high inventories. Given that production cannot be adjusted quickly, this should temporarily depress wholesale gasoline prices. Unfortunately, precise empirical identification of the resulting speed of price adjustment is probably not possible because demand changes — especially weather-driven changes — tend to be serially correlated and difficult to measure, unlike the crude price movements we study. While it might be possible to identify the general direction of such responses, cross-sectional comparison of adjustment lags would not be reliable.

[32] Henly, Potter and Town (1995) report that terminal sellers – of branded or unbranded gasoline – typically change their about every other day. Given that the adjustments we see take place over weeks, it is very unlikely that they could be due to menu costs.

relationships. Wholesalers who distribute branded product buy exclusively and repeatedly from the same branded refiner.[33] But wholesalers are free to purchase unbranded product from any refiner and reportedly switch unbranded supplies frequently. Long-term relationships may, therefore, be more important for branded than unbranded sales and this difference may explain some of the slower adjustment of branded prices. It is less clear, however, how long-term relationships would explain the relationship between adjustment speed and the level of price-cost margins.

V. Conclusion

We have argued that wholesale gasoline prices respond with a lag to cost shocks because it is costly for firms to adjust production and inventory. Rather than incurring the cost of immediate adjustment, firms minimize the cost of production and distribution by spreading adjustment over several periods. If production and inventory adjustment causes lagged price changes, we also show that the adjustment rate will be a function of market power. In the empirical work, we first examine the gasoline futures market where alternative explanations for price stickiness do not hold. We confirm the theoretical prediction that the futures price of gasoline responds incompletely to changes in the futures price of crude oil when futures contracts are near expiration. Using data from local wholesale gasoline markets, we then show that prices adjust more slowly in markets where there are higher price-cost margins.

We would not, however, argue that costly production and inventory adjustment is the entire cause of lags in gasoline price adjustment. While the behavior of local wholesale prices is consistent with our model's predictions, it also might be consistent with other explanations. In particular, long-term relationships between buyers and sellers are likely to play a role in these markets. Refiners have long-term relationships with independent wholesalers and with the stations selling branded gasoline to end users. An

---

[33] It is important to note that the branded stations do have a special relationship with branded refiners that may affect the rate at which the dealer tankwagon price is adjusted in response to cost shocks. However, the terminal prices are the prices at which the wholesalers buy.

interesting avenue for further research is to explore how these relationships might affect price dynamics.

Another interesting issue not directly addressed here is raised by the apparent asymmetry in the response of local wholesale prices to cost increases and decreases. This asymmetry is consistent with a convex NMCY curve. If, as shown in the appendix, the NMCY is convex, reducing inventory (increasing sales) will have a larger effect on NMCY than increasing inventory (reducing sales). Because reducing sales is easier (has less effect on marginal distribution cost), cost increases will be accomodated more quickly. While this is consistent with the model, we have no direct evidence that it is the convexity of the NMCY function that causes the asymmetry.

References

Bacon, Robert W. (1991). "Rockets and Feathers: The Asymmetric Speed of Adjustment of U.K. Retail Gasoline Prices to Cost Changes," *Energy Economics*, July.

Ball, Laurence, and N. Gregory Mankiw (1994). "A Sticky-Price Manifesto," *Carnegie-Rochester Conference Series on Public Policy*, Vol. 41, 127-151.

Bénabou, Roland and Robert Gertner (1993). "Search with Learning from Prices - Does Increased Inflationary Uncertainty Lead to Higher Markups?", *Review of Economic Studies*, January, 60(1), 69-93.

Borenstein, Severin, A. Colin Cameron and Richard Gilbert (1995). "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?", Program on Workable Energy Regulation working paper #PWP-001R, October.

Borenstein, Severin and Richard Gilbert (1993). "Uncle Sam At the Gas Pump: Causes and Consequences of Regulating Gasoline Distribution," *Regulation*, Spring.

Bresnahan, Timothy F. and Pablo T. Spiller (1986). "Futures Market Backwardation Under Risk Neutrality," *Economic Inquiry*, July, 24(3), 429-441.

Bresnahan, Timothy F. and Valerie Y. Suslow (1985). "Inventories as an Asset: The Volatility of Copper Prices," *International Economic Review*, June, 409-424.

Carlton, Dennis W. (1986). "The Rigidity of Prices," *American Economic Review*, September, pp. 637-658.

Carlton, Dennis W. (1991). "The Theory of Allocation and Its Implication for Marketing and Industrial Structure: Why Rationing is Efficient," *Journal of Law and Economics*, October, pp. 231-262.

Dahl, Carol and Thomas Sterner (1991). "Analyzing Gasoline Demand Elasticities: A Survey," *Energy Economics*, July, 203-210.

Fama, Eugene R. and Kenneth R. French (1987). "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage," *Journal of Business*, January, 55-73.

Henly, John, Simon Potter, and Robert Town (1995). "Price Rigidity, the Firm, and the Market: Evidence from the Wholesale Gasoline Industry During the Iraqi Invasion of Kuwait," mimeo, February.

Karrenbrock, Jeffrey D. (1991). "The Behavior of Retail Gasoline Prices: Symmetric or Not?," *Federal Reserve Bank of St. Louis Review*, July/August.

Mankiw, N. Gregory (1985). "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly," *Quarterly Journal of Economics*.

National Petroleum Council, *Petroleum Liquids Transportation*, report to the Secretary of Energy, 1989.

Pindyck, Robert S. (1993). "The Present Value Model of Rational Commodity Pricing," *Economic Journal*, May, pp. 511-530.

Pindyck, Robert S. (1994). "Inventory and the Short-Run Dynamics of Commodity Prices," *RAND Journal of Economics*, Spring.

Rotemberg, Julio J. (1982). "Sticky Prices in the United States," *Journal of Political Economy*, December, pp. 1187-1211.

Rotemberg, Julio J. and Garth Saloner (1987). "The Relative Rigidity of Monopoly Prices," *American Economic Review*, December, pp. 917-926.

Slade, Margaret (1986). "Exogeneity Tests of Market Boundaries Applied to Petroleum Products," *Journal of Industrial Economics*, March, pp. 291-303.

Thurman, Walter N. (1988). "Speculative Carryover: An Empirical Examination of the U.S. Refined Copper Market," *RAND Journal of Economics*, Autumn, pp. 420-437.

Table 1: Adjustment of Gasoline to Crude Oil Futures Prices

Dependent Variable: $\Delta GAS_{s,t} = GAS_{s,t} - GAS_{s,t-1}$

	Different Adjustment in Last Trading Month		Different Adjustment in Last 10 Days	
	OLS	2SLS	OLS	2SLS
$\Delta CRUDE_{s,t}$	+0.990	+1.029	+0.987	+1.027
	(0.022)	(0.025)	(0.021)	(0.025)
$\Delta CRUDE_{s,t-1}$	+0.087	+0.091	+0.084	+0.089
	(0.023)	(0.022)	(0.021)	(0.021)
$\Delta CRUDE_{s,t-2}$	+0.043	+0.043	+0.044	+0.044
	(0.017)	(0.018)	(0.017)	(0.018)
$D \cdot \Delta CRUDE_{s,t}$	−0.142	−0.100	−0.202	−0.144
	(0.025)	(0.038)	(0.031)	(0.033)
$D \cdot \Delta CRUDE_{s,t-1}$	−0.070	−0.054	−0.091	−0.068
	(0.031)	(0.029)	(0.033)	(0.039)
$D \cdot \Delta CRUDE_{s,t-2}$	+0.007	+0.012	+0.010	+0.009
	(0.021)	(0.023)	(0.026)	(0.032)
$CONSTANT$	+0.005	+0.005	+0.005	+0.005
	(0.009)	(0.009)	(0.009)	(0.009)
Observations:	21108	21108	21108	21108
R^2:	0.78		0.78	
Durbin-Watson:	2.08	2.07	2.08	2.07

Standard errors are corrected for heteroskedasticity and correlation among errors of different contracts on the same day

Table 2
Descriptive Statistics for Estimating Terminal Price Adjustment Rates

(68,320 observations)

(All Variables in ¢/gallon)

Variable	Mean	Std Dev	Min	Max
Branded Terminal Price	62.35	11.04	34.40	106.28
Unbranded Terminal Price	61.39	11.33	33.55	116.46
Crude Oil Price	46.94	10.23	25.83	96.19
Wkly Change in Branded	-0.04	2.32	-21.03	15.50
Wkly Change in Unbranded	-0.04	2.54	-18.00	26.00
Wkly Change in Crude	-0.02	3.07	-15.47	11.55

Table 3A: Partial Adjustment Estimates

(188 observations)

Variable	Mean	Std Dev	Min	Max
Symmetric Adjustment Rates				
Branded Adjustment Rate	0.184	0.028	0.056	0.227
Unbranded Adjustment Rate	0.208	0.032	0.113	0.473
Branded–Unbranded Adj Rate	-0.024	0.027	-0.282	0.035
Asymmetric Adjustment Rates – UP				
Branded Adjustment Rate	0.185	0.040	0.052	0.227
Unbranded Adjustment Rate	0.218	0.054	0.072	0.318
Branded–Unbranded Adj Rate	-0.033	0.035	-0.154	0.115
Asymmetric Adjustment Rates – DOWN				
Branded Adjustment Rate	0.184	0.037	0.034	0.275
Unbranded Adjustment Rate	0.202	0.046	0.113	0.556
Branded–Unbranded Adj Rate	-0.018	0.041	-0.377	0.071
Passthrough Rates				
Branded Passthrough Rate	0.877	0.042	0.675	0.979
Unbranded Passthrough Rate	0.907	0.048	0.740	1.044

able 3B: Cumulative Adjustment Estimates from Vector Autoregressions

(188 observations)

Variable	Mean	Std Dev	Min	Max
Normalized Symmetric Adjustment Rates				
Branded 1-Week	0.48	0.09	0.22	0.63
Branded 2-Week	0.89	0.14	0.45	1.10
Branded 3-Week	1.01	0.13	0.57	1.20
Unbranded 1-Week	0.65	0.11	0.17	0.83
Unbranded 2-Week	1.01	0.13	0.38	1.16
Unbranded 3-Week	1.09	0.11	0.77	1.34
Branded–Unbranded 1-Week	-0.17	0.08	-0.37	0.12
Branded–Unbranded 2-Week	-0.11	0.07	-0.32	0.19
Branded–Unbranded 3-Week	-0.09	0.06	-0.25	0.16
Normalized Asymmetric Adjustment Rates – UP				
Branded 1-Week	0.72	0.13	0.30	0.98
Branded 2-Week	1.16	0.16	0.49	1.47
Branded 3-Week	1.18	0.12	0.48	1.40
Unbranded 1-Week	0.92	0.14	0.31	1.17
Unbranded 2-Week	1.34	0.16	0.62	1.54
Unbranded 3-Week	1.28	0.13	0.66	1.47
Branded–Unbranded 1-Week	-0.20	0.10	-0.41	0.16
Branded–Unbranded 2-Week	-0.18	0.09	-0.41	0.10
Branded–Unbranded 3-Week	-0.11	0.08	-0.35	0.20
Normalized Asymmetric Adjustment Rates – DOWN				
Branded 1-Week	0.26	0.08	-0.07	0.40
Branded 2-Week	0.64	0.14	0.21	0.86
Branded 3-Week	0.80	0.14	0.33	1.03
Unbranded 1-Week	0.42	0.10	-0.01	0.60
Unbranded 2-Week	0.72	0.12	0.09	0.95
Unbranded 3-Week	0.88	0.13	0.47	1.11
Branded–Unbranded 1-Week	-0.16	0.09	-0.52	0.06
Branded–Unbranded 2-Week	-0.07	0.09	-0.33	0.16
Branded–Unbranded 3-Week	-0.07	0.07	-0.43	0.16
Passthrough Rates				
Branded Passthrough Rate	0.877	0.042	0.675	0.979
Unbranded Passthrough Rate	0.907	0.048	0.740	1.044

Table 4A
Effects of Market Structure on Gasoline Price Adjustment
Using Adjustment Rates from Partial Adjustment Model

Dependent Variable:	Branded	Unbranded	Branded–Unbranded
Branded Price	−0.0051 (0.0009)		
Unbranded Price		−0.0049 (0.0013)	
Branded–Unbranded Price			−0.0138 (0.0057)
Number of Refiners	+0.0013 (0.0003)	+0.0009 (0.0007)	
Distance From Refinery (00 miles)	+0.0028 (0.0007)	+0.0025 (0.0009)	
Off Pipeline	+0.0103 (0.0030)	+0.0031 (0.0040)	
PADD I	−0.0117 (0.0039)	−0.0071 (0.0048)	
PADD II	+0.0303 (0.0031)	+0.0283 (0.0048)	
Constant	+0.4618 (0.0568)	+0.4768 (0.0841)	−0.0124 (0.0060)
Observations:	188	188	188
R^2:	0.63	0.35	0.10

White's Heteroskedastic-Consistent Standard Errors in Parentheses

Table 4B
Effects of Market Structure on Gasoline Price Adjustment
Using Adjustment Rates from VAR Model

Dependent Variable:	Branded	Unbranded	Branded–Unbranded
Branded Price	−0.0216 (0.0035)		
Unbranded Price		−0.0228 (0.0035)	
Branded–Unbranded Price			−0.0270 (0.0103)
Number of Refiners	+0.0047 (0.0016)	+0.0078 (0.0017)	
Distance From Refinery (00 miles)	+0.0132 (0.0040)	+0.0099 (0.0034)	
Off Pipeline	+0.0565 (0.0152)	+0.0491 (0.0166)	
PADD I	−0.0633 (0.0220)	−0.0285 (0.0227)	
PADD II	+0.1805 (0.0185)	+0.1361 (0.0185)	
Constant	+2.0561 (0.2195)	+2.2011 (0.2069)	−0.0879 (0.0101)
Observations:	188	188	188
R^2:	0.66	0.55	0.05

White's Heteroskedastic-Consistent Standard Errors in Parentheses

Table 5A: Effects of Market Structure on Gasoline Price Adjustment
Using Asymmetric Adjustment Rates from Partial Adjustment Model

Dependent Variable:	UPWARD ADJUSTMENT			DOWNWARD ADJUSTMENT		
	Branded	Unbranded	Brand–Unbrand	Branded	Unbranded	Brand–Unbrand
Branded Price	−0.0105			−0.0012		
	(0.0016)			(0.0015)		
Unbranded Price		−0.0178			+0.0039	
		(0.0015)			(0.0019)	
Branded–Unbranded Price			−0.0128			−0.0134
			(0.0053)			(0.0075)
Number of Refiners	+0.0017	+0.0012		+0.0009	+0.0008	
	(0.0006)	(0.0007)		(0.0005)	(0.0010)	
Distance From Refinery(00 miles)	+0.0023	+0.0013		+0.0032	+0.0031	
	(0.0010)	(0.0014)		(0.0011)	(0.0015)	
Off Pipeline	+0.0059	−0.0025		+0.0137	+0.0067	
	(0.0060)	(0.0073)		(0.0054)	(0.0072)	
PADD I	−0.0130	+0.0040		−0.0109	−0.0141	
	(0.0073)	(0.0089)		(0.0059)	(0.0087)	
PADD II	+0.0122	+0.0317		+0.0425	+0.0253	
	(0.0067)	(0.0076)		(0.0053)	(0.0084)	
Constant	+0.8053	+1.2709	−0.0206	+0.2161	−0.0670	−0.0043
	(0.0976)	(0.0972)	(0.0051)	(0.0903)	(0.1264)	(0.0082)
Observations:	188	188	188	188	188	188
R^2:	0.49	0.54	0.05	0.44	0.18	0.04

White's Heteroskedastic-Consistent Standard Errors in Parentheses

Table 5B: Effects of Market Structure on Gasoline Price Adjustment
Using Asymmetric Adjustment Rates from VAR model

Dependent Variable:	UPWARD ADJUSTMENT			DOWNWARD ADJUSTMENT		
	Branded	Unbranded	**Brand–Unbrand**	**Branded**	Unbranded	**Brand–Unbrand**
Branded Price	−0.0288			−0.0126		
	(0.0047)			(0.0040)		
Unbranded Price		−0.0288			−0.0149	
		(0.0049)			(0.0036)	
Branded–Unbranded Price			−0.0268			−0.0330
			(0.0154)			(0.0123)
Number of Refiners	+0.0032	+0.0102		+0.0051	+0.0044	
	(0.0021)	(0.0020)		(0.0018)	(0.0018)	
Distance From Refinery(00 miles)	+0.0132	+0.0098		+0.0155	+0.0123	
	(0.0047)	(0.0050)		(0.0043)	(0.0034)	
Off Pipeline	+0.0697	+0.0538		+0.0525	+0.0410	
	(0.0192)	(0.0213)		(0.0177)	(0.0174)	
PADD I	−0.0631	−0.0437		−0.0712	−0.0136	
	(0.0275)	(0.0327)		(0.0237)	(0.0239)	
PADD II	+0.1977	+0.1498		+0.1714	+0.1129	
	(0.0255)	(0.0278)		(0.0193)	(0.0191)	
Constant	+2.7901	+2.8700	−0.1538	+1.2335	+1.4713	−0.0409
	(0.2884)	(0.2872)	(0.0182)	(0.2535)	(0.2188)	(0.0133)
Observations:	188	188	188	188	188	188
R^2:	0.60	0.55	0.03	0.59	0.37	0.05

White's Heteroskedastic-Consistent Standard Errors in Parentheses

Appendix A:
The Slope of the Net Marginal Convenience Yield Curve

The theory presented in section II relies on the assumption that the net marginal convenience yield (NMCY) declines as inventories rise in comparison to sales. In a competitive commodity market, arbitrage will always force NMCY to be equal to the difference between the current and next-period price. Borrowing from Pindyck (1994) then, we take the NMCY to be equal to the futures price of the nearest unleaded gasoline contract minus the futures price of the next-to-nearest unleaded gasoline contract.[34] We call this variable *SPOT-FUT*. Pindyck (1994) estimates the slope of NMCY curves for three commodities – copper, lumber, and heating oil. In each case, he finds the NMCY to be negative sloped and convex in inventories divided by sales.

Unfortunately, gasoline inventory data are available only for "primary" storage, which includes gasoline stored prior to sale at the terminal. This excludes storage facilities of wholesalers, storage tanks at retail gasoline stations, and private storage in automobiles and elsewhere. We have been told that the exclusion of secondary (jobber) and tertiary (station and individual) level storage from the data make the information on week-to-week inventory changes quite unreliable, but that the data contain some useful indication of the overall level of inventories in the U.S.[35]

With those caveats in mind, we proceed with estimation of the relationship between NMCY and the inventory-to-sales ratio. We use primary level inventories in the east-of-Rockies states (Petroleum Administration for Defense Districts (PADDs) I, II, and III) and divide them by contemporaneous consumption in these states.[36] Because there is extremely high serial correlation for daily observations of the same contract (about 0.97) and

---

[34] Gasoline futures contracts are traded through the end of the month prior to delivery, so the "spot minus future" data for observations in July of 1992 would be the price for August delivery gasoline minus the price for September delivery gasoline.

[35] The inventory data are available from the Oil & Gas Journal Energy Database. The caveats are from a conversation with Thomas Hogarty of American Petroleum Institute.

[36] The consumption data are "first sales" from the Energy Information Administration, which are available monthly. To match these with the weekly inventory data, we linearly interpolated the consumption data, taking the data for a given month to accurately represent the consumption rate for the middle day of the month and taking a weighted average for days between the middle days of each month.

because the week-to-week changes in inventory are very noisy indicators of the true changes, we run monthly regressions. To aggregate to months we use, alternatively, the average *SPOT-FUT* and inventory during a month or the *SPOT-FUT* and inventory on the last trading day of the month.

Table A1 presents descriptive statistics for *SPOT-FUT* and *DAYSINV*, which is the days of inventory on hand at current consumption rates (primary inventory divided by daily consumption). Table A2 presents estimates of linear and quadratic relationships between *SPOT-FUT* and *DAYSINV*. There was serial correlation in the residuals of these regressions, so we correct for an AR(1) process. The results indicate that *SPOT-FUT* is very sensitive to inventories. There is weak evidence that it is more sensitive on the last day of the month, closest to the delivery date of the contract, than averaged over the whole month, but the difference is not statistically significant. The regressions of the linear relationship indicate that an increase in inventories equal to one day of contemporaneous consumption lowers *SPOT-FUT* by about 0.4¢. The relationship, however, does not appear to be linear, but rather convex. The regressions with quadratic terms in *DAYSINV* indicates that the *SPOT-FUT* is steeper at low levels of inventory than at high levels. Using the last-day-of-month observations, the slope is estimated to be -1.05 at 26 days of inventory, -0.59 at 30 days of inventory, and zero at about 35 days of inventory.

One implication of the convex relationship between *DAYSINV* and *SPOT-FUT* is that, according to the discussion in section II, one would expect slower adjustment when inventories are low and the NMCY curve is steep, than when inventories are high and NMCY curve is relatively flat. Translated to the futures price adjustment regressions reported in Table 1, this would mean that the smaller adjustment observed in the last month before a contract expires would be exacerbated when inventories are low or when *SPOT-FUT* is high. However, our attempts to uncover such a relationship in the regressions reported in Table 1 — using interaction terms with *DAYSINV* or *SPOT-FUT*, estimating different last month coefficients depending on whether *DAYSINV* (or *SPOT-FUT*) is above or below its mean, or estimating different coefficients when *DAYSINV* is particularly low or *SPOT-FUT* particularly high — have not revealed any relationship. We cannot report a significant effect of *DAYSINV* or *SPOT-FUT* on the extent of adjustment of the prices of gasoline futures contracts during the last month before delivery.

Table A1: Descriptive Statistics of *SPOT-FUT* and Inventory

(95 observations)

Variable	Mean	Std Dev	Min	Max
Last Trading Day of Month				
SPOT-FUT	1.19	2.37	-5.43	10.14
Days of Inventory	30.26	2.48	25.96	38.66
Average Over Month				
SPOT-FUT	0.97	1.89	-5.65	6.67
Days of Inventory	30.31	2.51	26.37	38.19

Table A2: Estimates of the NMCY Curve

Dependent Variable: *SPOT-FUT*

	Last Day of Month		Average Over Month	
DAYSINV	−0.434	−4.012	−0.402	−2.995
	(0.102)	(1.451)	(0.078)	(0.972)
$DAYSINV^2$		+0.057		+0.041
		(0.023)		(0.015)
Constant	+14.268	+70.251	+13.057	+53.647
	(3.197)	(23.260)	(2.408)	(15.746)
Observations:	95	95	95	95
R^2:	0.14	0.18	0.18	0.22
ρ	0.42	0.42	0.55	0.55
Durbin-Watson:	2.11	2.10	2.18	2.24

White's Heteroskedastic-Consistent Standard Errors in Parentheses