

Sticky prices, inventories, and market power in wholesale gasoline markets

Severin Borenstein*

and

Andrea Shepard**

A model with costly adjustment of production and costly inventories implies that wholesale gasoline prices will respond with a lag to crude oil cost shocks. Unlike explanations that rely upon menu costs, imperfect information, or long-term buyer/seller relationships, this model also predicts that futures prices for gasoline will adjust incompletely to crude oil price shocks that occur close to the expiration date of the futures contract. We test and confirm this implication. Examining wholesale price responses in 188 gasoline markets, we also find that firms with market power adjust prices more slowly than do competitive firms, consistent with the model.

1. Introduction

■ From January 1999 to March 2000, the price of crude oil tripled from about \$10 per barrel to over \$30. Significant crude price volatility continued through 2000. In addition, changes to reformulated gasoline in some parts of the United States have been associated with further volatility in gasoline prices that has continued into 2001. Oil companies argue that gasoline price movements reflect oil prices, costs of adjusting supply, and the normal interaction of supply and demand. Many consumers and politicians have attributed these sharp increases in gasoline prices to refiners seizing an opportunity to exploit their market power. Although there are idiosyncratic aspects to local price fluctuations, recent events have highlighted two of the underlying issues that affect all gasoline markets: the cost of adjusting the supply of gasoline and the relationship between price adjustment and market power.

In this article we examine one piece of this puzzle: the responses of gasoline prices to cost shocks and how those responses might be affected by market power. This research is motivated by a puzzling empirical pattern in the movements of gasoline prices: wholesale gasoline prices do not respond fully and immediately to changes in the price of crude oil. The lag in adjustment occurs even though the price of crude oil is a major determinant of the cost of gasoline and crude price changes are public information. Borenstein, Cameron, and Gilbert (1997), for example,

* University of California, Berkeley, University of California Energy Institute, and NBER; borenste@haas.berkeley.edu.

** Stanford University and NBER; shepard_andrea@gsb.stanford.edu.

For helpful comments we thank Tim Bresnahan, Dennis Carlton, Robert Pindyck, Rob Porter, two anonymous referees, and participants at numerous seminars. Wedad Elmaghraby and Celeste Saravia provided excellent research assistance.

documents that crude oil price shocks are eventually fully passed through to wholesale gasoline prices, but full adjustment takes many weeks. Because wholesale prices are formed in well-organized markets in which one would expect market-clearing prices to incorporate quickly all available information, these lags are surprising. A change in crude prices changes 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 increased by an amount sufficient to cause a five-cent increase in gasoline prices in the long run. If the firm had been producing at a level where short-run marginal cost equalled price, then its marginal cost is now above price and it has an incentive to reduce production. Since all refiners have a similar incentive, each would withhold supply and the price would immediately increase by five cents. Whether the price of crude oil rises or falls, the price of gasoline should adjust immediately. An analogous argument holds if refiners have some market power. The magnitude of the gasoline price adjustment to a given crude oil price change might be different, but the response should still be immediate. The purpose of this research is to explore possible explanations for why immediate adjustment does not happen.

We consider two classes of explanations. One comes from theories of supply adjustment costs. Because adjusting levels of production is costly, firms in these models spread the adjustment over time. A reduction in the price of crude oil, for example, will imply some long-run increase in the supply of gasoline. But if there are adjustment costs that increase with the absolute size of adjustment per period, firms will spread the adjustment over time, gradually achieving the full quantity increase implied by the decline in cost. Models of this type have been proposed by Pindyck (1993, 1994) and Thurman (1988), among others. It is important to note that although it takes some time to achieve the new long-run equilibrium in these models, price adjusts to clear the market at every point in time. There is no nonprice allocation of supply or demand. This explanation is consistent with the arguments made by many refiners over the last few years.

The other class of explanations comes from a large literature that spans industrial organization and macroeconomics and focuses on differences between the market-clearing price and the price at which spot transactions actually occur. (See, for example, Carlton (1986, 1989, 1991), Rotemberg (1982), Mankiw (1985), Ball and Mankiw (1994), and Bénabou and Gertner (1993).) In this literature, the problem is to explain why the transaction prices charged do not change (or do not change quickly) when the market-clearing price changes. The models offer a variety of mechanisms that lead to a spot transaction price that does not clear the market and, thus, to nonprice allocation. Menu costs, for example, that make changing the transaction price costly have been offered as an explanation for infrequent price changes even when market-clearing prices change frequently. Other models point to information imperfections as the source of the divergence. When information is imperfect and demand is relatively elastic, quantity adjustments might be preferable to price adjustment. Or buyers and sellers might minimize transaction costs and reduce price risk by entering long-term agreements that limit price flexibility.

Setting aside the details of individual models, there is an important empirical implication that can distinguish the first class of models from the second. In an efficient auction market—e.g., a well-organized commodity market with many competing, well-informed traders—all transaction prices will be equal to the market-clearing price. There are no menu costs in these markets, no long-term contracts between traders, and no information problems that can produce transaction prices that are different from the market-clearing price. In these markets, price stickiness cannot be explained by models that depend on transaction prices differing from market clearing prices.

Stickiness that results from supply adjustment costs, however, can affect these markets. If supply cannot adjust immediately (at a sufficiently low cost), market-clearing prices may be “sticky” relative to the long-run costs of production. We exploit this difference, using prices from oil and gasoline futures markets, to test whether supply adjustment costs contribute to sticky gasoline prices. Our results indicate that supply adjustment costs induce stickiness in gasoline prices.

The supply adjustment cost models that predict price stickiness also predict that the rate of price adjustment will typically be affected by variation in market power. The price adjustment rate in an imperfectly competitive industry to a given cost shock will differ from the adjustment rate of a competitive industry even if supply adjustment technology is held constant. Unfortunately, the supply adjustment cost models have no robust prediction about the direction of the difference. Without detailed information on the structure of demand and cost, these models cannot tell us whether a monopolist, for example, would adjust more or less quickly than a competitive firm. Nonetheless, if supply adjustment costs are an important determinant of price stickiness, we might expect to find some systematic relationship between market power and price adjustment speeds, *ceteris paribus*.

To explore this possibility, we use data on price adjustments in local wholesale gasoline markets. Arguably, the firms in these markets have similar adjustment technologies and face a similar demand structure. We find evidence consistent with faster price adjustment rates in markets that seem to be more competitive. This result suggests that there might also be some truth to the claim made by many politicians that market power causes slower price adjustment. But market power appears to slow both price decreases and price increases.

The structure of the article is as follows. In Section 2 we review the competing explanations for price stickiness, focusing on the different implications of supply adjustment cost models that predict sticky market-clearing prices and other models that predict a divergence between market-clearing prices and actual transaction prices. In Section 3 we develop and apply a test that distinguishes stickiness caused by supply adjustment costs from price stickiness from other sources. In Section 4 we turn to the relationship between price stickiness and market power. We conclude in Section 5.

2. Sticky wholesale gasoline prices

■ This research is motivated by the observation that wholesale gasoline prices respond slowly to crude oil price shocks. These wholesale prices—called “terminal” prices because they are for transactions occurring at the terminals where refiners sell gasoline to local wholesalers—were formed in well-organized, relatively frictionless markets. Buyers and sellers are well informed about price and quality. Prices are available online from private data firms, and the products’ specifications are well established. The flow of sales is fairly continuous, and prices change frequently. 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 (see, for example, Borenstein, Cameron, and Gilbert (1997)).

Industry participants queried about lagged price adjustment typically cite the difficulty of making supply adjustments. It is well established in the economics literature that production adjustment costs can give rise to lagged adjustment of prices to input cost shocks. This can occur even when firms can carry inventories so long as it is costly to deviate from some otherwise preferred level of inventories. (See, for example, Pindyck (1994) and the references therein.) The Appendix sketches a typical supply adjustment cost model with inventories, but the basic results are quite straightforward and intuitive.¹ A crude oil price shock changes the marginal cost of production. A competitive firm that was in long-run equilibrium will want to adjust production until marginal cost is again equal to the price of gasoline. If marginal production adjustment costs increase in the level of production change, the firm will distribute the adjustment over time and full adjustment will be achieved with a lag. Period by period, the firm will equate the full marginal cost of output (which includes the marginal cost of adjustment) to the market price. If the firm can hold inventories, adjustment will be faster than it would otherwise be. If, however, there are costs of deviating from the otherwise optimal level of inventories, and such costs increase in

¹ We refer to the combination of production adjustment costs and inventory costs as “supply adjustment costs,” since both conditions are necessary to induce sticky prices.

the magnitude of the deviation, the ability to hold inventories will not lead to full, immediate adjustment.

If this explanation is to hold for lagged gasoline price responses, several things must be true. First, there must be significant production adjustment costs. In the gasoline industry, these costs are probably most significant at the refinery. Refinery production schedules are set by solving a complex algorithm that takes into account the (expected) demand for the various refined products and the (expected) cost 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 its current production “run,” as well as the magnitude of the change.

Production adjustment costs alone, however, are not enough to lead to lagged price adjustments when refiners hold inventory. At any point in time, refiners hold gasoline inventories equal to about 25 days of gasoline sales.² When the cost of oil rises, they can increase inventories to reduce the supply of gasoline offered on the market. When the price of oil falls, they can sell out of inventory. If there were no costs to increasing or decreasing the level of inventories held, inventory adjustments would lead to immediate price adjustment.

In reality, large changes in inventory are very costly. Inventory capacity imposes an upper bound on how much inventory can be increased. The distribution system for gasoline poses a lower bound on how much inventory can be reduced. For example, a large share of the gasoline sold in the United States is transported from the refinery to the local wholesale terminals through pipelines. These pipelines must be full of product to function. Furthermore, a refiner typically sells gasoline in substantially more locations than it has refineries. As an example, consider a firm with a single refinery on the Gulf Coast and sales throughout the eastern United States. To have product available for sale in, say, Newark, New Jersey next period, it must have product in transit to Newark this period. 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 via pipeline from the Gulf Coast to Newark (see National Petroleum Council, 1989). In each period, then, the refiner must have many days of supply in transit to Newark.

Within these bounds, there is an optimal inventory level that equates the marginal benefit and cost of holding inventory. In a long-run equilibrium with constant crude oil prices, the cost of carrying inventories would be just the cost of storage services plus the interest cost of forgoing sales this period. The benefit of carrying inventories (above the technical minimum required to maintain the system) is insurance against future stockouts. Refiners selling branded gasoline have a network of wholesalers and retailers committed to carrying the branded product. A refiner unable to supply its distribution network will lose dealers or face a higher cost of maintaining the network. The marginal benefit of carrying inventories minus the marginal storage cost is called the “net marginal convenience yield” (or NMCY) in the adjustment cost literature. NMCY is generally declining in the quantity of inventories. In a long-run steady state, the firm will equate the marginal cost and benefit of inventory so that $NMCY = 0$.

Now, consider what happens to competitive refiners when the cost of crude oil increases.³ In aggregate, refiners reduce production but will not adjust output to the new equilibrium levels immediately because adjusting production is costly. If firms do not adjust their inventory levels, it is clear that output prices in the future (when production has fully adjusted) will be higher than current prices. Expected price changes introduce an additional cost or benefit to holding

² For our sample period, U.S. inventories held by refiners (called “primary inventories”) are estimated to average a bit more than 200 million barrels, and daily sales averaged 8 million barrels per day.

³ We assume here and in the model sketched in the Appendix that the crude oil prices follow a random walk. This is consistent with our empirical model in which we estimate the impact of the nontransitory component of crude oil price changes. See the Appendix for additional discussion.

inventory. If prices will be higher in the future, the firm has an incentive to shift sales from the current period to the future by increasing inventory levels. As inventories rise, NMCY declines. Period by period, the firm will adjust inventories so that NMCY equals the expected change in price ($NMCY = P_t - E[P_{t+1}]$).

Two characteristics of the price adjustment path are worth emphasizing. First, prices change each period. Prices are sticky in the sense that they do not adjust *completely* when the cost shock is realized. This pattern is distinct from the typical menu cost story that attributes stickiness to the cost of changing prices. In those models, the firms change price less frequently than costs change. Second, in adjustment cost models, supply always equals demand. In other models of sticky prices, there is excess demand or supply at the spot transaction price. Put differently, in the supply adjustment cost models, the spot transaction price and the market-clearing price are equal throughout the adjustment process. In other models, the spot transaction price will generally deviate from the market-clearing price. We will exploit this difference to construct a test of the hypothesis that supply adjustment costs lead to sticky gasoline prices.

3. Testing for supply adjustment cost effects

■ The most obvious way to test for supply adjustment cost effects is to examine the behavior of inventories and output as cost shocks are realized. Unfortunately, in our setting, this approach is infeasible because reliable and comprehensive data on output and inventories do not exist.⁴ Given this constraint, we turn to a test that is based on how market-clearing prices adjust. As we argue above, supply adjustment costs models predict that the market-clearing price will adjust slowly to cost shocks. Other explanations of sticky transaction prices predict that the transaction prices will adjust slowly (or infrequently), but the market-clearing price will adjust immediately and fully. If the market-clearing price is sticky it is because supply is sticky due to supply adjustment costs. We can therefore distinguish stickiness caused by supply adjustment costs from price stickiness caused by other factors by testing for stickiness in market-clearing prices. We do this by studying the response of the price of gasoline futures contracts to changes in the price of the futures contracts for crude oil.

U.S. futures markets for oil and gasoline are widely recognized to be quite efficient, with thousands of traders able to exploit even extremely small arbitrage opportunities. Prices change many times within each trading day, so stickiness in transaction price is not a significant factor, and trades are anonymous, so long-term relationships between buyers and sellers are irrelevant. With so little friction in the market mechanism, only stickiness in the market-clearing price can lead to stickiness in futures prices.

Menu costs, long-term relationships between buyers and sellers, or anything else that affects transaction prices but not the market-clearing price will have no effect on futures prices. Suppose, for example, that there were no supply adjustment costs and all spot sales of wholesale gasoline were governed by contracts requiring refiners to adjust wholesale gasoline prices slowly. If the price of crude oil were to decline, refiners' marginal cost would be below the transaction price. Without supply adjustment costs, they would immediately increase their output. Because contracts prevent the transaction price from adjusting immediately, there would be excess supply if there were no futures market. The futures market, however, will absorb the excess supply. Refiners will sell futures contracts and the price of these contracts will fall to the market-clearing price. The price of the futures contract will respond immediately and completely to the cost shock.

When supply adjustment costs are the source of stickiness in the transaction price, however, the price of futures contracts can also be "sticky." To be concrete, consider the futures 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 (during our sample period),

⁴ Inventory data are available only for "primary" storage, which includes only gasoline prior to sale at the terminal. We have been told that the exclusion of secondary and tertiary storage makes the information on week-to-week inventory changes quite noisy. Data on production are also noisy for short-term changes.

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 the price of the July crude oil contract occurring in May or earlier. Whatever gasoline production and inventory adjustments were necessary to reach the new long-run equilibrium could be implemented by July. As a result, the price of the July gasoline contract would adjust fully and rapidly (within a day or two) to the shock. But now consider a shock occurring in June. If supply adjustment costs exist, it may not be optimal for the refiners to complete the necessary adjustments by the contract delivery date. In that case, the price of the July gasoline contract will still adjust immediately (otherwise, weak-form market efficiency would be violated), but it will adjust *incompletely* to the change in the price of the July crude oil contract. The adjustment will be immediate because there is nothing constraining the response of the futures price. The response will not incorporate the full, long-term impact of the cost shock because the anticipated market-clearing price for July will not incorporate the full impact.

In summary, stickiness in the supply response of refiners to cost shocks can lead to incomplete adjustment of the price of gasoline futures contracts to changes in crude oil futures contracts. The other sources of price stickiness commonly suggested in the literature are not present in these futures markets. If we find incomplete adjustment in near-expiration futures prices, the most compelling explanation is supply adjustment costs. Our hypothesis, then, is that the price of gasoline futures contracts will respond incompletely to price changes for contemporaneous crude oil contracts when these contracts are near their expiration dates.

To test this hypothesis, we use a dataset of New York Mercantile Exchange (NYMEX) prices for futures contracts for New York harbor delivery of light sweet crude oil and unleaded regular gasoline. The dataset consists of daily prices for 110 contracts with a delivery month between December 1985 and January 1995, and it includes nearly 21,000 observations. The estimation is based on a simple lagged adjustment model. The dependent variable is the daily change (from the close on trading day $t - 1$ to the close on trading day t) in the price of the gasoline contract for a given month, e.g., the change between May 3, 1991 and May 4, 1991 in the price of the January 1992 gasoline futures contract. 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.⁵ 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.⁶

The estimating equation is

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

where s indexes the month of delivery specified in the contract, t indexes the trading day observed, $\Delta GAS_{s,t} = GAS_{s,t} - GAS_{s,t-1}$, $\Delta CRUDE_{s,t} = CRUDE_{s,t} - CRUDE_{s,t-1}$, and D is a dummy variable that equals one when trades occur in the last month before contract expiration. We hypothesize complete 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 is 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

⁵ Over this period, the average per-gallon price of gasoline on the futures market was about 55 cents with a minimum of 31 cents and a maximum of \$1.01. The average per-gallon price of crude oil on the futures market was about 46 cents, with a minimum of 25 cents and a maximum of 96 cents. The average absolute daily price change per gallon for both crude and gasoline contracts was about .5 cents.

⁶ 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.

calendar days before the crude contract stops trading. Our hypothesis is that the sum of these coefficients is negative. Because supply adjustment costs imply that adjustments will be less complete for shocks that occur later in the last month, we also estimate (1) setting $D = 1$ only for observations after the 10th day of the last month the contract is traded. We hypothesize that the sum of these coefficients is negative and larger in absolute value than the sum of the coefficients for the last 22 calendar days.

Two econometric issues arise in estimating (1). First, crude oil prices might be endogenous. It is possible that demand shocks in the United States, for example, could affect NYMEX futures prices for gasoline and create short-run excess demand or supply in U.S. crude oil markets that affect the NYMEX futures price of crude oil.⁷ To test and correct for endogeneity, we estimate (1) by two-stage least squares. The instruments we use are changes in spot crude oil prices in England⁸ and the change in the six-month-ahead futures price of crude oil on the NYMEX.⁹ These instruments should be immune to transitory, North American demand shocks. They will also remove the effects of transitory, local shocks to crude oil supply. Transitory shocks—unlike the permanent shocks that are the focus of our analysis—might not be fully incorporated into the price of gasoline even in the absence of adjustment costs. In contrast, decision makers would expect the crude oil price shocks incorporated in these instruments to last at least many months, long enough to be fully incorporated into the price of gasoline.

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. We find that we cannot reject nonstationarity of these series, consistent with weak-form market efficiency in these futures contract markets, but we find that the series are cointegrated.¹⁰

The results from estimating equation (1), reported in Table 1, support the supply adjustment cost 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.

Focusing on the 2SLS estimation, the results indicate that the price of the futures contract for a given month goes up by about 1.18 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 .16 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 .22-cent smaller effect on gasoline prices than changes occurring earlier in the contract life, which is statistically different from zero at the 1% level.¹¹

Earlier work has indicated that the speed of response may be asymmetric: decreases in crude prices may be passed along more slowly to downstream price than are increases. If this

⁷ Endogeneity is suggested by occasional reports in the press that movements in the nearest gasoline futures contract affect the associated crude oil contract.

⁸ These are the same-day spot prices for Brent and North Sea Forties crude oil.

⁹ A Hausman test using these instruments confirms the endogeneity.

¹⁰ An augmented Dickey-Fuller test of the crude oil (gasoline) futures price series gives a test statistic of -1.75 (-2.20), well below (in absolute value) the 95% critical value of -3.66 for rejecting a unit root. The Dickey-Fuller test for stationarity of the residual in the cointegrating regression is -4.19 , exceeding (in absolute value) the 95% critical value of -2.23 .

¹¹ The effect in the last 10 days, however, is statistically different from the earlier part of the last trading month, with a p -value of only .13.

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
$\Delta CRUDE_{s,t}$	+0.991 (.021)	+1.034 (.024)	+0.989 (.020)	+1.031 (.023)
$\Delta CRUDE_{s,t-1}$	+0.092 (.024)	+0.097 (.023)	+0.089 (.022)	+0.095 (.022)
$\Delta CRUDE_{s,t-2}$	+0.045 (.017)	+0.044 (.018)	+0.045 (.017)	+0.045 (.018)
$D \cdot \Delta CRUDE_{s,t}$	-.101 (.029)	-.100 (.038)	-.154 (.026)	-.147 (.034)
$D \cdot \Delta CRUDE_{s,t-1}$	-.072 (.032)	-.066 (.031)	-.095 (.035)	-.084 (.041)
$D \cdot \Delta CRUDE_{s,t-2}$	+0.008 (.022)	+0.011 (.023)	+0.017 (.027)	+0.016 (.032)
CONSTANT	+0.000 (.000)	+0.000 (.000)	+0.000 (.000)	+0.000 (.000)
Observations	20,879	20,879	20,879	20,879
R^2	.78		.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.

phenomenon also exists 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.

4. Lagged price adjustment and market power

■ If supply adjustment costs create price stickiness, we would expect the price adjustment rate to be affected by market power. There are several price adjustment cost models that predict response rates will depend on the level of competition in the market. To illustrate this point, we can think about the different optimization problems facing a monopolist and a competitive firm, holding supply adjustment costs and market size constant.¹² A competitive firm is concerned about the expected change in price as the adjustment occurs. It will equate the expected price change to its NMCY and to the change in its marginal cost of production (including the marginal adjustment cost). A monopolist is concerned about the change in marginal revenue and will equate the expected change in marginal revenue to its NMCY and to the change in its marginal cost of

¹² We discuss the monopolist's optimization problem in somewhat more detail in the Appendix.

production. Because marginal revenue will not be equal to price for a monopolist, the adjustment rates will not be the same.

While supply adjustment cost models predict that adjustment speeds will differ between competitive and monopoly firms, they do not provide robust predictions about where prices will adjust more quickly. Instead, adjustment rate differences will be affected not just by market power but by cost and demand as well. For example, all else constant, a monopolist facing a linear demand curve will adjust price less quickly than would a competitive market, but a monopolist facing a constant-elasticity demand curve will adjust price more quickly. The reason is that when demand is linear, the marginal revenue curve is steeper than the demand curve, but in the constant-elasticity case the marginal revenue curve is less steep than the demand curve. More generally, the effect of market power on price adjustment rates is sensitive to demand structure, functional-form assumptions on adjustment costs, and the way competition is modelled. (See, for example, Worthington (1989), Ginsburgh and Michel (1988), and Bedrossian and Moschos (1988).)¹³ While adjustment cost theory does not provide much guidance on the direction of the relationship, it does suggest that market power could affect adjustment rates.

Note that evidence of a link between market power and price adjustment rates—unlike the evidence on incomplete futures price adjustments—cannot be used to distinguish the effect of supply adjustment costs from the effect of other frictions. We discuss frictions from menu costs and long-term relationships after presenting our results. Before proceeding, however, we want to acknowledge that nonprice allocation mechanisms might be important in these markets. There is a long tradition in the industrial organization literature positing that market power will affect the rate of price adjustment in ways that are unrelated to the mechanisms emphasized in the supply adjustment cost literature. Beginning with Stigler and Kindahl (1970) and Carlton (1979), industrial organization economists have searched empirically for a link between market power and price adjustment rates. While the theoretical underpinning for the conjectured link is sometimes not well specified, there has been a widespread acceptance that there might well be a connection between market power and price adjustment patterns that is sustained by more than supply adjustment costs. One interpretation of our work in this section of the article is as an empirical contribution to that literature.

To explore the relationship between lagged responses to cost shocks and market power, we turn to data on wholesale gasoline prices in local markets. There are approximately 370 terminals in the United States where refiners sell 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. The number of refiners supplying product to a given terminal ranges from as few as 4 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 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. A major refiner may sell unbranded as well as branded gasoline at the terminal, e.g., Chevron might sell both “Chevron” gas to be sold at Chevron stations and unbranded gasoline that can be sold through any unbranded station.¹⁴ 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.¹⁵

We describe our empirical model and the data we use in some detail below, but our basic approach is to regress estimated adjustment rates on indicators of market power at the terminal

¹³ A referee has also pointed out that the relationship between adjustment speed and market structure might be nonmonotonic, an issue we address in the empirical work.

¹⁴ Branded gasoline contains the additive package associated with the brand. Unbranded gasoline need not, and often does not, contain the same additive package.

¹⁵ Refiners supply some gasoline stations directly. 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. There are no reliable data on delivered prices.

level. The first step is to produce estimated adjustment rates for each terminal. With these in hand, we implement two tests for a relationship between market power and adjustment rates. First, we test for a systematic relationship across terminals between the adjustment rate and proxies for the price-cost margin at the terminal. If we find evidence that adjustment rates are negatively (positively) correlated with our measure of price-cost margin, we will conclude that greater market power leads to slower (faster) adjustment.

Second, we compare the adjustment rates of branded and unbranded gasoline at the same terminal. We expect demand for branded gasoline to be less elastic than demand for unbranded gasoline. 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, end-use consumers of unbranded product are probably more prone to switching stations in response to price differences. The greater the difference in market power between branded and unbranded sellers, 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 branded-unbranded price-cost margin *difference* for each terminal.

For both tests, we begin by estimating the adjustment speed of branded and unbranded wholesale prices, terminal by terminal. Descriptive statistics on the variables used in the analysis appear in Table 2. The terminal price data are the average prices for 87 octane, unleaded gasoline on each Friday at the given terminal. The data include separate averages for branded and unbranded product.¹⁶ We match these price observations to the Friday Gulf Coast spot price for West Texas Intermediate (WTI) crude oil.¹⁷ The dataset contains weekly observations for the period January 1, 1986 through November 20, 1992 at 188 terminals. We restrict the estimation to terminals located in petroleum administration defense districts (PADDs) I, II, and III, which include all areas of the continental United States east of the Rocky Mountains.¹⁸ These PADDs are most clearly in the same crude and gasoline distribution system.¹⁹ The economic isolation of gasoline and crude markets for PADDs IV and V and the integration of PADDs I, II, and III suggested by the physical distribution systems have been confirmed empirically by Slade (1986). We include only terminals for which we have at least 200 usable observations, i.e., with sufficient lagged price information to include the observation in the data.

We detrend and deseasonalize the price data to remove the effects of inflation and any seasonal pattern. 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. There is a marked seasonal pattern in the demand for refined petroleum products that is reflected in gasoline prices.²⁰ Detrending and deseasonalizing the data are consistent with assuming that prices fully and immediately adjust to these effects. The results are very similar to those reported if the data are not detrended or deseasonalized, though there are small declines in the significance levels of some estimates.²¹

¹⁶ The terminal price data used here are collected and reported by Lundberg Survey, Inc. The branded (unbranded) average is the simple average of all branded (unbranded) terminal prices on Friday.

¹⁷ The spot price of West Texas Intermediate (and the prices of other crudes used as instruments) is constructed from a survey of traders conducted daily by the *Dow Jones International Petroleum Report*.

¹⁸ PADDs I, II, and III are closely linked by major pipelines and water transportation for both crude and gasoline products.

¹⁹ PADD V, the West Coast, refines primarily Alaska North Slope crude oil and does not trade significant quantities of gasoline with other PADDs. PADD IV (Colorado, Idaho, Montana, Utah, and Wyoming) has no significant pipeline links and is largely self-sufficient.

²⁰ 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.

²¹ This robustness isn't surprising, since the right-hand-side variables, crude oil prices, have no significant seasonal pattern, and the time trend reflects only the difference between the beginning and ending price and so is unlikely to affect estimates of short-run price adjustment.

TABLE 2 Descriptive Statistics for Estimating Terminal Price Adjustment Rates (68,320 observations)

Variable	Mean	Standard Deviation	Minimum	Maximum
Branded terminal price (¢/gallon)	62.35	11.04	34.40	106.28
Unbranded terminal price (¢/gallon)	61.39	11.33	33.55	116.46
Crude oil price (¢/gallon)	46.94	10.23	25.83	96.19
Weekly change in branded (¢/gallon)	-.04	2.32	-21.03	15.50
Weekly change in unbranded (¢/gallon)	-.04	2.54	-18.00	26.00
Weekly change in crude (¢/gallon)	-.02	3.07	-15.47	11.55

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 the results from two models: a partial-adjustment model (PAM) and a vector autoregression (VAR) model.²² A partial-adjustment model is attractive because it generates a single parameter estimate of the adjustment rate. The structure imposed by the model also tends to yield reasonably precise estimates. However, the structure constrains price adjustments to be of equal proportions in all adjustment periods, a restriction at odds with prior research (e.g., Borenstein, Cameron, and Gilbert, 1997). 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 functions rather than single parameters.²³

The PAM we estimate for each terminal is

$$\Delta TERM_t = \delta(TERM_t^* - TERM_{t-1}) + \mu_t, \tag{2}$$

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

$$TERM_t = \alpha_0 + \alpha_1 CRUDE_t + \varepsilon_t. \tag{2a}$$

The long-run adjustment proportion—or passthrough rate—is α_1 . Consistent with prior research, our estimates (reported in Table 3) suggest this parameter is near one.²⁴

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 ensure that our estimates of market structure effects are not driven by pooling different price response paths, we also estimate a version of (2) 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

$$\Delta TERM_t = \delta_1^+(TERM_t^* - TERM_{t-1})^+ + \delta_2^-(TERM_t^* - TERM_{t-1})^- + \mu_t. \tag{3}$$

The + (–) superscript indicates that the target minus the lagged terminal price is positive (negative), i.e., that terminal prices are increasing (decreasing).²⁵

²² 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 we report.

²³ For a description of calculating cumulative response functions from a VAR model, see Borenstein, Cameron, and Gilbert (1997).

²⁴ We first estimate the long-run relationship between crude and terminal prices using (2a) and then use the predicted value in the partial-adjustment model (2). We use this two-step procedure for consistency with the asymmetric model discussed below. Estimating the symmetric model in a single step has no substantive effect on the results.

²⁵ To be more precise, $(TERM_t^* - TERM_{t-1})^+ = \max\{0, TERM_t^* - TERM_{t-1}\}$ and $(TERM_t^* - TERM_{t-1})^- = \min\{0, TERM_t^* - TERM_{t-1}\}$.

TABLE 3 Partial Adjustment Estimates (188 observations)

Variable	Mean	Standard Deviation	Minimum	Maximum
Symmetric Adjustment Rates				
Branded adjustment rate	.184	.028	.056	.227
Unbranded adjustment rate	.208	.032	.113	.473
Branded-unbranded adjustment rate	-.024	.027	-.282	.035
Asymmetric Adjustment Rates—UP				
Branded adjustment rate	.185	.040	.052	.253
Unbranded adjustment rate	.218	.054	.072	.318
Branded-unbranded adjustment rate	-.033	.035	-.154	.115
Asymmetric Adjustment Rates—DOWN				
Branded adjustment rate	.184	.037	.034	.275
Unbranded adjustment rate	.202	.046	.113	.556
Branded-unbranded adjustment rate	-.018	.041	-.377	.071
Passthrough Rates				
Branded passthrough rate	.877	.042	.675	.979
Unbranded passthrough rate	.907	.048	.740	1.044

The VAR model we estimate is

$$\Delta TERM_t = \sum_{j=0}^5 \theta_j \Delta CRUDE_{t-j} + \sum_{k=1}^5 \phi_k \Delta TERM_{t-k} + \gamma (TERM_{t-1}^* - TERM_{t-1}) + \mu_t. \quad (4)$$

This VAR model allows us to calculate the cumulative response of terminal prices for the weeks following a cost shock.²⁶ As in 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 regression equation (2a).

The asymmetric version of the VAR uses the same long-run estimates for the error-correction term, but it differentiates between positive and negative price changes. The asymmetric model is²⁷

$$\Delta TERM_t = \sum_{j=0}^5 \left[\theta_j^+ \Delta CRUDE_{t-j}^+ + \theta_j^- \Delta CRUDE_{t-j}^- \right] + \sum_{k=1}^5 \left[\phi_k^+ \Delta TERM_{t-k}^+ + \phi_k^- \Delta TERM_{t-k}^- \right] + \gamma (TERM_{t-1}^* - TERM_{t-1}) + \mu_t, \quad (5)$$

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 3 and 4. All estimates are from a 2SLS procedure in which we instrument for West Texas Intermediate crude

²⁶ We truncate the estimation at six weeks (0 through 5) because coefficients on longer lags were insignificant in all but a few cities.

²⁷ Because the VAR allows so much flexibility in the asymmetric lagged change variables, we restrict the error-correction term to be symmetric.

TABLE 4 Cumulative Adjustment Estimates from Vector Autoregressions (188 observations)

Variable	Mean	Standard Deviation	Minimum	Maximum
Normalized Symmetric Adjustment Rates				
Branded 1-week	.48	.09	.22	.63
Branded 2-week	.89	.14	.45	1.10
Branded 3-week	1.01	.13	.57	1.20
Unbranded 1-week	.65	.11	.17	.83
Unbranded 2-week	1.01	.13	.38	1.16
Unbranded 3-week	1.09	.13	.64	1.26
Branded-unbranded 1-week	-.17	.08	-.37	.12
Branded-unbranded 2-week	-.11	.07	-.32	.19
Branded-unbranded 3-week	-.09	.06	-.25	.16
Normalized Asymmetric Adjustment Rates—UP				
Branded 1-week	.72	.13	.30	.98
Branded 2-week	1.17	.16	.49	1.47
Branded 3-week	1.18	.12	.48	1.40
Unbranded 1-week	.92	.14	.31	1.17
Unbranded 2-week	1.35	.16	.62	1.54
Unbranded 3-week	1.28	.13	.66	1.47
Branded-unbranded 1-week	-.20	.10	-.41	.16
Branded-unbranded 2-week	-.18	.09	-.41	.10
Branded-unbranded 3-week	-.11	.08	-.35	.20
Normalized Asymmetric Adjustment Rates—DOWN				
Branded 1-week	.26	.08	-.07	.40
Branded 2-week	.64	.14	.21	.86
Branded 3-week	.80	.14	.33	1.03
Unbranded 1-week	.42	.10	-.01	.60
Unbranded 2-week	.71	.12	.09	.95
Unbranded 3-week	.88	.13	.47	1.11
Branded-unbranded 1-week	-.16	.09	-.52	.06
Branded-unbranded 2-week	-.07	.09	-.33	.16
Branded-unbranded 3-week	-.07	.07	-.43	.16
Passthrough Rates				
Branded passthrough rate	.877	.042	.675	.979
Unbranded passthrough rate	.907	.048	.740	1.044

prices (*CRUDE*) as described for estimating equation (1).²⁸ Table 3 reports the summary statistics for the terminal-by-terminal adjustment rate estimates from the PAM. Table 4 reports the summary statistics for the terminal-by-terminal 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. For example, these functions record the cumulative movement of terminal prices as of the end of the second week as a proportion of the initial movement in crude prices. To make these comparable to the adjustment rates for the PAM, we divide by the estimated

²⁸ As in the previous section, the use of these instruments controls for endogeneity of the crude oil price and removes transitory, local shocks from the crude oil price movements.

long-run adjustment to the crude shock. This gives us a rate that is the ratio of cumulative terminal price adjustment through the n th week to the full terminal price adjustment.²⁹

As reported in Table 3, 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. The results do not suggest that the response of branded gasoline price is asymmetric, but they give some indication 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 4 for the first three weeks. Like the PAM estimates, the VAR results show faster adjustment for unbranded than for branded prices. The estimates, however, suggest that both adjust at a faster rate than the rates estimated by the PAM. 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 is also 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 estimated 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," which 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). Because these districts are served by different pipeline systems, there may be regional cost differences that would be picked up by these dummy variables.³⁰

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. In other formulations, the number of competitors is used to proxy competition because researchers do not have data on price-cost margins. Since we have data on price but only proxies for marginal cost, the effect of "Number of Refiners" captures unobserved differences in cost. For example, consider two terminals with the same price for branded gasoline and the same values for the cost proxies, but different numbers of competitors. If market power declines in the number 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,

²⁹ Augmented Dickey-Fuller tests for a unit root of the WTI crude oil price series gives a test statistic of -3.64 , just below (in absolute value) the 95% critical value of -3.66 . The branded (unbranded) terminal price series of the 188 terminals give ADF test statistics that average -3.42 (-3.44). Tests for cointegration of branded terminal with WTI or unbranded terminal with WTI all exceed the 99% critical value.

³⁰ We also ran specifications that included a count of the number of "nearby" terminals (defined alternatively as within 50, 100, or 200 miles) in case interterminal competition has an important effect on market power. The coefficients on "nearby terminals" were statistically indistinguishable from zero, and this had no effect on our other coefficient estimates.

TABLE 5 Descriptive Statistics for Estimating Effects of Market Structure on Gasoline Price Adjustment (188 observations)

Variable	Mean	Standard Deviation	Minimum	Maximum
Average branded price ($\text{\$/gallon}$)	62.39	2.07	58.68	69.89
Average unbranded price ($\text{\$/gallon}$)	61.39	2.06	57.02	66.27
Average branded-unbranded ($\text{\$/gallon}$)	1.00	.61	-.54	5.07
Number of refiners	14.32	4.76	4	27
Distance from refinery	197	225	0	1110
Off pipeline	.20	—	0	1
PADD I	.27	—	0	1
PADD II	.54	—	0	1
PADD III	.19	—	0	1

we would expect the number of competitors to have a coefficient of the same sign as the other cost variables. Holding price constant, higher costs imply lower margins and lower market power. Our data on the number of competitors come from a single cross-section in 1990.³¹ Descriptive statistics for the exogenous variables in these regressions are provided in Table 5.

The results using symmetric response models are reported in Table 6 (columns 1 and 2) for the PAM and in Table 7 (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 report results for week two, but 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 symmetric PAM (Table 6), 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 .5-percentage-point decline in the adjustment rate. Holding price constant, higher transportation costs or more competitors reduce the price-cost margin and increase the adjustment rate. Adjustment rates for unbranded products are also 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 those for branded product as well, but the number of competitors and "Off Pipeline" are not significantly different from zero in the unbranded gasoline regressions. 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 7) 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, holding cost constant, lowers the second-week cumulative adjustment (which has a mean of .89) by about .02.

The evidence from the asymmetric models is slightly more mixed, but it generally confirms the implications from the symmetric models. The asymmetric PAM estimates (Table 8) do not support the market-power hypothesis very 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 crude price decline increases with costs, but it also appears to increase with the unbranded

³¹ 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 according to whether or not the fuel was branded did not increase explanatory power.

TABLE 6 Effects of Market Structure on Gasoline Price Adjustment Using Adjustment Rates from Partial-Adjustment Model

Dependent Variable	Branded	Unbranded	Branded– Unbranded
Branded price	-.0051 (.0009)		
Unbranded price		-.0049 (.0013)	
Branded–unbranded price			-.0138 (.0057)
Number of refiners	.0013 (.0004)	.0009 (.0007)	
Distance from refinery (00 miles)	.0029 (.0008)	.0025 (.0009)	
Off pipeline	.0103 (.0030)	.0031 (.0041)	
PADD I	-.0117 (.0039)	-.0071 (.0049)	
PADD II	.0303 (.0031)	.0283 (.0049)	
Constant	.4618 (.0579)	.4768 (.0857)	-.0124 (.0060)
Observations	188	188	188
R ²	.63	.35	.10

White's heteroskedastic-consistent standard errors in parentheses.

terminal price. The regressions using the VAR estimates (Table 9), however, 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 results of the asymmetric adjustment models do not support the assertion that market power leads refiners to raise prices quickly but lower them slowly. Even the VAR estimates, which indicate the strongest effect of market power on adjustment rates, do not indicate that the market power effect is greater for decreases than increases. That is, while market power does appear to lead to slow adjustment of wholesale gasoline prices, it does so for both the increases and the decreases, and to about the same extent.

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 do unbranded sellers, the preceding results suggest that they will want to adjust prices more slowly. Indeed, the estimated response rate suggests that the response rate for branded gasoline is slower on average than the response rate for unbranded gasoline. By regressing the difference between branded and unbranded adjustments rates at the i th terminal on the difference in branded and unbranded price cost margins at the i th terminal, we are identifying the relationship between market power and adjustment rates through variation in the market power of branded gasoline relative to unbranded across terminals. This difference regression has the advantage of eliminating most cost differences. Most of the costs of selling gasoline at the terminal are the same for branded and unbranded gas. Because our cost variables in the "levels"

TABLE 7 Effects of Market Structure on Gasoline Price Adjustment Using Adjustment Rates from VAR Model

Dependent Variable	Branded	Unbranded	Branded– Unbranded
Branded price	–.0216 (.0036)		
Unbranded price		–.0228 (.0035)	
Branded–unbranded price			–.0270 (.0103)
Number of refiners	.0047 (.0017)	.0078 (.0017)	
Distance from refinery (00 miles)	.0132 (.0040)	.0099 (.0035)	
Off pipeline	.0565 (.0155)	.0491 (.0169)	
PADD I	–.0633 (.0224)	–.0285 (.0232)	
PADD II	.1805 (.0189)	.1361 (.0188)	
Constant	2.0561 (.2237)	2.2011 (.2108)	–.0879 (.0101)
Observations	188	188	188
R^2	.66	.55	.05

White's heteroskedastic-consistent standard errors in parentheses.

regressions controlled for cost imperfectly, we potentially measured price-cost margins with error. In the difference specification the effects of omitted cost variables that are the same for branded and unbranded gasoline terminal are differenced out. Differencing also eliminates the impact of the cost variables we had previously included.³²

The results of this approach strongly support the conclusion that the branded adjustment rate is lower relative to the unbranded rate where the branded price is higher 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 the average price differential at the terminal leads to a .0138 decrease in the adjustment speed parameter for branded relative to unbranded. Likewise, the VAR result implies that a one-cent increase in the average price differential leads to the cumulative adjustment of branded being .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.

We have imposed a linear relationship between adjustment rates and our measures of price-cost margins. We use this specification because it is parsimonious and because there is no obvious

³² 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.

TABLE 8 Effects of Market Structure on Gasoline Price Adjustment Using Asymmetric Adjustment Rates from Partial-Adjustment Model

Dependent Variable	Upward Adjustment			Downward Adjustment		
	Branded	Unbranded	Branded– Unbranded	Branded	Branded– Unbranded	Unbranded
Branded price	-.0105 (.0016)			-.0012 (.0015)		
Unbranded price		-.0178 (.0016)			.0039 (.0019)	
Branded–unbranded price			-.0128 (.0054)			-.0134 (.0075)
Number of refiners	.0017 (.0006)	.0012 (.0008)		.0009 (.0006)	.0008 (.0010)	
Distance from refinery (00 miles)	.0023 (.0010)	.0013 (.0015)		.0033 (.0011)	.0031 (.0015)	
Off pipeline	.0059 (.0062)	-.0025 (.0074)		.0137 (.0055)	.0067 (.0073)	
PADD I	-.0130 (.0074)	.0040 (.0091)		-.0109 (.0060)	-.0141 (.0089)	
PADD II	.0122 (.0068)	.0317 (.0078)		.0425 (.0054)	.0253 (.0086)	
Constant	.8053 (.0995)	1.2709 (.0990)	-.0206 (.0051)	.2161 (.0921)	-.0670 (.1288)	-.0043 (.0082)
Observations	188	188	188	188	188	188
R ²	.49	.54	.05	.44	.18	.04

White's heteroskedastic-consistent standard errors in parentheses.

way to estimate a nonlinear relationship in the levels regressions. We have estimated a difference regression (between branded and unbranded at the same terminal) with the price difference and the price difference squared on the right-hand side. While there is some evidence that the quadratic specification improves the fit, the substantive results were unchanged and the estimated relationships are monotonic (a negative relationship between difference in market power and difference in speed of adjustment) over the range of observations in the sample.

The regression results are consistent with the supply adjustment cost theory, but we cannot rule out some other explanation as well. Although the terminals are well-organized markets in which a well-defined product is sold at observable prices, they are probably not as efficient as the futures market. Trades on the futures market, for example, are anonymous whereas the transactions at the terminals are not. Buyers and sellers at terminal markets interact repeatedly. Indeed, wholesalers who distribute branded product buy exclusively and repeatedly from the same branded refiner.³³ Because wholesalers are free to purchase unbranded product from any refiner (and reportedly switch unbranded suppliers frequently), long-term relationships may be less important for unbranded sales. Even so, the parties to unbranded transactions are familiar with one another. There also may be more repeated contact between the same seller and buyer

³³ It is important to note that the branded stations 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 independent wholesalers buy.

TABLE 9 Effects of Market Structure on Gasoline Price Adjustment Using Asymmetric Adjustment Rates from VAR Model

Dependent Variable	Upward Adjustment			Downward Adjustment		
	Branded	Unbranded	Branded– Unbranded	Branded	Branded– Unbranded	Unbranded
Branded price	–.0288 (.0048)			–.0126 (.0041)		
Unbranded price		–.0288 (.0049)			–.0149 (.0036)	
Branded–unbranded price			–.0268 (.0179)			–.0330 (.0123)
Number of refiners	.0031 (.0021)	.0102 (.0020)		.0051 (.0019)	.0044 (.0019)	
Distance from refinery (00 miles)	.0132 (.0048)	.0098 (.0051)		.0155 (.0043)	.0123 (.0034)	
Off pipeline	.0697 (.0195)	.0538 (.0217)		.0525 (.0180)	.0410 (.0177)	
PADD I	–.0631 (.0280)	–.0437 (.0333)		–.0712 (.0242)	–.0136 (.0243)	
PADD II	.1977 (.0260)	.1498 (.0283)		.1714 (.0197)	.1129 (.0194)	
Constant	2.7901 (.2948)	2.8700 (.2927)	–.1538 (.0183)	1.2335 (.2583)	1.4713 (.2230)	–.0409 (.0134)
Observations	188	188	188	188	188	188
R ²	.60	.55	.03	.59	.37	.05

White's heteroskedastic-consistent standard errors in parentheses.

when there are fewer sellers in the market. It is not clear, however, why long-term relationships would lead sellers to respond slowly to public cost shocks.

If menu costs prevented refiners from adjusting prices frequently, one might expect to see some lagged adjustment. Menu costs, however, are not a plausible explanation for the length of lags we observe. Terminal prices are generally reviewed by refiners on a daily basis. Henly, Potter, and Town (1995) report that terminal sellers—of branded or unbranded gasoline—typically change their prices about every other day. Furthermore, many of the observed weekly price changes in our data are quite small. For unbranded gasoline, for example, 25% of the weekly price changes were less than .5 cents in value. The mean absolute value of the unbranded price change was 1.1 cents. If menu costs imply that price changes are delayed until a large adjustment is required, the observed small changes seem inconsistent with the menu cost view. Given that the adjustment lags we see take place over weeks, it is very unlikely that they could be due to menu costs. Furthermore, menu costs could not reasonably be higher for branded than for unbranded product.

Although the analysis based on futures price patterns implies that adjustment costs matter, our results on terminal prices might well be affected by other frictions. This raises a question of whether the adjustment cost story, even if true, is economically important: can adjustment costs plausibly explain the three-week price adjustment lag we find in terminal prices? The industry sources we spoke with claim that inventory adjustment costs are very important, but we would need better data on inventories and production than are available to answer this question systematically. We can, however, make some headway by using the following thought experiment. Suppose crude

oil prices increased by \$1 per barrel, an amount roughly consistent with the average, absolute change in WTI over our sample period. This change implies that the cost of a gallon of gas will increase by 2.4 cents, an increase that will be fully passed through into the price of gasoline. The average retail price of gasoline in our sample period is about \$1 per gallon, so a 2.4-cent increase is a 2.4% change in price. Using the midrange of the short-run elasticity estimates in the literature, .15 (see Dahl and Sterner, 1991), this translates to a .32% decline in gasoline consumption in the short run, or about 25,000 barrels per day. The question, then, is whether it is reasonable to think that an adjustment of 25,000 barrels per day can be costly “enough” to imply that it will not be fully accomplished for three weeks.

Suppose for the moment that production cannot be adjusted at all for three weeks. At the end of three weeks, inventory would then be 525,000 barrels higher than it was before the cost of crude increased. Can inventory absorb this amount without a significant increase in inventory costs? More precisely, can it be absorbed without a substantial increase in the NMCY? One way to think about the effects is to observe what happens when news of inventory changes is revealed to traders in the future markets. Suppose they discover that inventory is higher than they had expected. If the “extra” inventory has little effect on NMCY, it will not change their expectations about prices. If, however, traders respond by bidding down the nearest futures price, we can assume that the increment to inventories has increased NMCY.³⁴

The *Wall Street Journal* reports a weekly gasoline inventory estimate prepared by the American Petroleum Institute. The reports often compare the inventory changes to those expected by analysts as reflected in, for instance, surveys by the *Dow Jones International Petroleum Report*. Deviations of 500,000 barrels between expectations and the API’s estimates are typically noted as the cause for movements in the price of the nearest gasoline futures contract. Indeed, deviations of as little as 200,000 barrels have been described as the cause of a price change. Unfortunately, we do not have a sufficient number of these observations to do a statistical analysis. Still, we take this as at least some evidence that the inventory adjustments of the magnitude implied by the crude oil price changes we study are indeed costly and could lead to adjustment lags.

5. 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. This explanation for lagged price responses is consistent with the behavior of the futures price of gasoline: the futures price of wholesale gasoline responds incompletely (though still very rapidly) to changes in the price of crude oil when futures contracts are near their expiration dates. Because alternative explanations for price stickiness cannot explain such “sticky” futures prices, this is strong evidence that supply adjustment costs create sticky prices.

We have also found that prices adjust more slowly where price-cost margins are higher. These results are consistent with, though not a test of, the supply adjustment models that imply that a firm with market power would generally choose a different adjustment rate than occurs in a perfectly competitive market. Despite the tendency of gas prices to respond to oil price increases faster than decreases, we find no evidence that the effect of market power on price adjustment is asymmetric; where refiners have greater market power they lower prices more slowly, but they also raise prices more slowly.

Supply adjustment costs might be related to the apparent asymmetry in the response of local wholesale prices to cost increases and decreases in all market structures. This asymmetry is consistent with a convex NMCY curve posited by standard adjustment cost models and validated empirically for some commodities (see, for example, Pindyck, 1994). The value of NMCY depends

³⁴ Recall that refiners adjust inventories so that NMCY equals the difference between the current price and the futures price. If traders observe that refiners have increased their inventory holdings more than seems justified by the current price difference, they can infer that refiners will begin to sell out of inventory to increase NMCY, reducing the expected next-period price.

on the ratio of inventory to sales. When this ratio is high (low), the value of holding another unit of inventory is low (high). If NMCY is a convex function of the ratio of inventory to sales, 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 accommodated more quickly. Although the evidence of asymmetry we find is consistent with a convex NMCY, we have no direct evidence that it is this convexity that causes the asymmetry in response rate.

Appendix

■ **A simple adjustment cost model.** In this Appendix we sketch a simple multiperiod model to demonstrate that supply adjustment costs can cause price response lags even when markets are efficient, information is perfect, and trade is anonymous. 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 permanent 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 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.³⁵

For simplicity, we assume that the marginal cost of production is constant in quantity produced 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 a mean zero, i.i.d. shock realized in period t . To capture the idea that production adjustment is costly and increasing in the magnitude of the adjustment, we posit a quadratic adjustment cost function. The total cost of production in period t , therefore, is

$$C_t = (c_{t-1} + \mu_t)y_t + a(\Delta y_t)^2, \quad (A1)$$

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 they include transportation, scheduling, and holding 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 by assuming interest rates are zero. 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 that $s_n > 0$ 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. Because a large amount of gasoline sold in the United States travels by pipeline, there is a minimum level of inventory that must be held within the system. Above this technical minimum, refiners hold inventory to avoid local stockouts at the many locations at which it sells wholesale gasoline. 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.³⁶ The decline in distribution costs as inventory increases is called “convenience yield” in the inventory literature. Convenience yield is a function of inventory levels, n_t , and expected next-period sales, $E_t q_{t+1}$, where q is sales and the subscript on the expectations operator indicates the period in which expectations are formed. Because sales in any period is determined by the inventory level at the beginning of the period, production in the period and inventory carried into the next period (i.e., $q_k = y_k - n_k + n_{k-1}$), the decline in distribution costs through convenience yield can be written as $-\phi(n_t, x_{t+1})$, where $x_{k+1} \equiv E_k(y_{k+1} + n_{k+1})$.³⁷ Holding future production and inventory constant, distribution costs decline as current 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 $\phi_x > 0$ and $\phi_{xx} > 0$. For higher levels of inventory, the same incremental increase in expected sales probably has a smaller effect on distribution costs; we therefore assume

³⁵ Our work owes a clear debt to the macroeconomic study of the relationship between inventories and sticky prices, such as Blinder and Fischer (1981), Blinder (1982, 1994), and West (1986). Because of its business-cycle orientation, however, that line of work tends to focus on demand fluctuations and price adjustment, while we focus on price adjustments to cost fluctuations.

³⁶ 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.

³⁷ We adopt the accounting convention that the firm bears the distribution costs for sales made in $t + 1$ during period t .

$\phi_{nx} < 0$. These assumptions are consistent with prior theoretical and empirical work (see, for example, Pindyck (1994), Thurman (1988) and Fama and French (1987)). To capture the fact that inventories are never drawn down to zero, we assume that $\lim_{n \rightarrow 0} -\phi_n = \infty$.

Total distribution costs, then, are given by $\phi(n_t, x_{t+1}) + s(n_t)$. 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

$$TC_t = (c_{t-1} + \mu_t)y_t + a(\Delta y_t)^2 + \phi(n_t, x_{t+1}) + s(n_t). \quad (A2)$$

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 and increasing its inventory. For a price-taking firm, the optimal adjustment to a cost shock realized in period t is found by maximizing the value function:³⁸

$$V_t = E_t \sum_{j=0}^{\infty} \left[p_{t+j} q_{t+j} - (c_{t+j-1} + \mu_{t+j})y_{t+j} - a(\Delta y_{t+j})^2 - \phi(n_{t+j}, x_{t+j+1}) - s(n_{t+j}) \right]. \quad (A3)$$

Rewriting sales to incorporate the adding up constraint $q_k = y_k - n_k + n_{k-1}$ gives us

$$V_t = E_t \sum_{j=0}^{\infty} \left[p_{t+j}(y_{t+j} - n_{t+j} + n_{t+j-1}) - (c_{t+j-1} + \mu_{t+j})y_{t+j} - a(\Delta y_{t+j})^2 - \phi(n_{t+j}, x_{t+j+1}) - s(n_{t+j}) \right]. \quad (A4)$$

The firm maximizes (A4) with respect to inventory and production.

The first-order condition for y_t is

$$p_t - (c_{t-1} + \mu_t) - 2a(\Delta y_t - E_t \Delta y_{t+1}) = 0. \quad (A5)$$

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 period- t production equals the change in period- t sales. Condition (A5) 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 ($a = 0$), (A5) is 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 n_t is

$$-p_t + E_t p_{t+1} - \phi_n(n_t, x_{t+1}) - s_n(n_t) = 0. \quad (A6)$$

Changing n_t while holding y_t and all future n 's and y 's constant implies an intertemporal shift in the quantity sold. Increasing n_t by one unit, e.g., shifts one unit of sales from t to $t + 1$. Condition (A6), then, defines the optimal tradeoff between the cost and benefit of holding an additional unit of inventory to accomplish an intertemporal shift in sales, taking into account expected 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).

³⁸ To focus on adjustment effects, our formulation of the maximization problem ignores discounting. As written, then, (A3) is unbounded. Including a discount factor is straightforward but increases the complexity of the math without additional economic insight.

Conditions (A5) and (A6) must hold for all periods j . In particular, there is a first-order condition for y_{t+1} analogous to condition (A5):

$$E_t p_{t+1} - (c_t + E_t \mu_{t+1}) - 2a E_t \Delta y_{t+1} + 2a E_t \Delta y_{t+2} = 0. \tag{A7}$$

Jointly, conditions (A5), (A6), and (A7) imply the following equalities:

$$p_t - E_t p_{t+1} = -\phi_n(n_t, x_{t+1}) - s_n(n_t) = 2a(\Delta y_t - 2E_t \Delta y_{t+1} + E_t \Delta y_{t+2}) + E_t \mu_{t+1}, \tag{A8}$$

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 (A8) 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. (A8) 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.³⁹

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. 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 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 price adjustment, i.e., the more elastic the industry demand, the lower the proportion of the eventual passthrough that is accomplished in the first period.

To see why the equilibrium adjustment rate can be affected by market power, we consider how the analysis changes if the firm is a monopoly supplier. The monopoly problem is a simple variation on the competitive problem. The optimization problem is essentially the same as in (A3), except p_t is now a function of q_t , i.e., the firm is not a price taker. The analogs to (A5) and (A6) are given by

$$p_t + \frac{\partial p_t}{\partial q_t} q_t - (c_{t-1} + \mu_t) - 2a(\Delta y_t - E_t \Delta y_{t+1}) = 0 \tag{A9}$$

and

$$-p_t - \frac{\partial p_t}{\partial q_t} q_t + E_t [p_{t+1} + \frac{\partial p_{t+1}}{\partial q_{t+1}} q_{t+1}] - \phi_n(n_t, x_{t+1}) - s_n(n_t) = 0. \tag{A10}$$

The first-order condition with respect to this period's production (A9) differs from (A5) only because marginal revenue is not equal to price for the monopolist. Similarly, the first-order condition for inventories, (A10), 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 firm's NMCY; in a monopoly market the intertemporal difference in marginal revenues is equated with the firm's NMCY. This conclusion assumes that there is little inventory held by other entities. If there are competitive firms (e.g., independent gasoline wholesalers) able to hold gasoline inventories, each will hold inventories at a level that equalizes its NMCY with the expected price change. As the availability of competitive inventory capacity increases, price adjustment will more closely resemble the competitive rate. This potential erosion of the refiners' ability to control the rate of change in gasoline prices may explain why refiners report that they ration sales to independent gasoline wholesalers when there is an arbitrage opportunity these wholesalers could otherwise exploit.

Throughout this modelling exercise, we have assumed that all crude oil price shocks are permanent, but transitory shocks due to capacity and inventory constraints on the production and storage of crude oil also affect crude oil prices. It is straightforward to show that transitory shocks will yield different gasoline price responses. For instance, in the presence of supply adjustment costs, a one-week temporary increase in crude oil price would have almost no effect on gasoline prices. Still, if even a transitory shock lasts substantially longer than the production adjustment period (generally about a month) and the time during which inventories can absorb extreme changes in production (two to three weeks at most),

³⁹ Clearly this is not reasonable, since there are nonnegativity 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.

the price of gasoline will fully adjust to reflect the transitory price. As explained in the text, we avoid the complication of transitory crude oil price changes by using instruments that do not reflect short-term, local crude oil price fluctuations.

References

- BALL, L. AND MANKIW, N.G. "A Sticky-Price Manifesto." *Carnegie-Rochester Conference Series on Public Policy*, Vol. 41 (1994), pp. 127–151.
- BEDROSSIAN, A. AND MOSCHOS, D. "Industrial Structure, Concentration and the Speed of Price Adjustment." *Journal of Industrial Economics*, Vol. 36 (1988), pp. 459–475.
- BÉNABOU, R. AND GERTNER, R. "Search with Learning from Prices: Does Increased Inflationary Uncertainty Lead to Higher Markups?" *Review of Economic Studies*, Vol. 60 (1993), pp. 69–93.
- BLINDER, A.S. "Inventories and Sticky Prices: More on the Microfoundations of Macroeconomics." *American Economic Review*, Vol. 72 (1982), pp. 334–348.
- . "On Sticky Prices: Academic Theories Meet the Real World." In N.G. Mankiw, ed., *Monetary Policy, Studies in Business Cycles*, Vol. 29. Chicago: University of Chicago Press, 1994.
- BLINDER, A.S. AND FISCHER, S. "Inventories, Rational Expectations, and the Business Cycle." *Journal of Monetary Economics*, Vol. 8 (1981), pp. 277–304.
- BORENSTEIN, S., CAMERON, A.C., AND GILBERT, R. "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics*, Vol. 112 (1997), pp. 305–339.
- CARLTON, D.W. "The Rigidity of Prices." *American Economic Review*, Vol. 76 (1986), pp. 637–658.
- . "The Theory of Allocation and Its Implication for Marketing and Industrial Structure: Why Rationing Is Efficient." *Journal of Law and Economics*, Vol. 34 (1991), pp. 231–262.
- DAHL, C. AND STERNER, T. "Analysing Gasoline Demand Elasticities: A Survey." *Energy Economics*, Vol. 13 (1991), pp. 203–210.
- FAMA, E.R. AND FRENCH, K.R. "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage." *Journal of Business*, Vol. 60 (1987), pp. 55–73.
- GINSBURGH, V. AND MICHEL, P. "Adjustment Costs, Concentration and Price Behaviour." *Journal of Industrial Economics*, Vol. 36 (1988), pp. 477–481.
- HENLY, J., POTTER, S., AND TOWN, R. "Price Rigidity, the Firm, and the Market: Evidence from the Wholesale Gasoline Industry During the Iraqi Invasion of Kuwait!" Mimeo, February 1995.
- MANKIW, N.G. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Quarterly Journal of Economics*, Vol. 100 (1985), pp. 529–537.
- NATIONAL PETROLEUM COUNCIL, *Petroleum Liquids Transportation*. Report to the Secretary of Energy, 1989.
- PINDYCK, R.S. "The Present Value Model of Rational Commodity Pricing." *Economic Journal*, Vol. 103 (1993), pp. 511–530.
- . "Inventory and the Short-Run Dynamics of Commodity Prices." *RAND Journal of Economics*, Vol. 25 (1994), pp. 141–159.
- ROTEMBERG, J.J. "Sticky Prices in the United States." *Journal of Political Economy*, Vol. 90 (1982), pp. 1187–1211.
- SLADE, M.E. "Exogeneity Tests of Market Boundaries Applied to Petroleum Products." *Journal of Industrial Economics*, Vol. 34 (1986), pp. 291–303.
- THURMAN, W.N. "Speculative Carryover: An Empirical Examination of the U.S. Refined Copper Market." *RAND Journal of Economics*, Vol. 19 (1988), pp. 420–437.
- WEST, K.D. "A Variance Bounds Test of the Linear Quadratic Inventory Model." *Journal of Political Economy*, Vol. 94 (1986), pp. 374–401.
- WORTHINGTON, P.R. "On the Distinction Between Structure and Conduct: Adjustment Costs, Concentration, and Price Behavior." *Journal of Industrial Economics*, Vol. 38 (1989), pp. 235–238.