

The Incidence of an Oil Glut: Who Benefits from Cheap Crude Oil in the Midwest?

October 2012

Abstract

Beginning in early 2011, crude oil production in the U.S. Midwest and Canada surpassed the pipeline capacity to transport it to the Gulf Coast where it could access the world oil market. As a result, the U.S. “benchmark” crude oil price in Cushing, Oklahoma, declined substantially relative to internationally traded oil. In this paper, we study how this development affected prices for refined products, focusing on the markets for motor gasoline and diesel. We find that the relative decrease in Midwest crude oil prices did not pass through to wholesale gasoline and diesel prices. This result is consistent with evidence that the marginal gallon of fuel in the Midwest is still imported from coastal locations. Our findings imply that investments in new pipeline infrastructure between the Midwest and the Gulf Coast, such as the southern segment of the controversial Keystone XL pipeline, will not raise gasoline prices in the Midwest.

1. Introduction

It is often said that the crude oil market is “global”. That is, the relative ease with which oil can be shipped around the world implies that prices for oil in different locations (but of a similar grade) will be closely tied by an arbitrage condition. Beginning in early 2011, however, a substantial, persistent differential arose between the price of crude oil in the United States’ Midwest and international prices at “on the water” locations. In particular, the two most well-known benchmark crude oil prices—Brent and West Texas Intermediate (WTI)—diverged substantially. Brent crude oil contracts have a delivery point in the North Sea, and most other on the water pricing points, including those on the U.S. Gulf Coast, follow Brent closely. WTI is priced at Cushing, Oklahoma, which is connected to water via pipelines to the Gulf Coast.


Figure 1 shows spot prices for WTI, Brent, and Light Louisiana Sweet (LLS) crude oil from 2006 through 2011. The LLS price, set near the Gulf Coast, closely matches Brent throughout the sample. The WTI price, however, falls significantly below the LLS and Brent prices in 2011, and a gap of 20 to 65 cents per gallon persists to the present (in contrast, LLS and Brent prices have stayed within 7 cents of one another). This gap is generally attributed to substantial increases in crude oil production in North Dakota (driven by technological advances in production of shale oil such as hydraulic fracturing or “fracking”) and oil from the Canadian tar sands, combined with constrained pipeline capacity between Cushing and the Gulf Coast.¹ The resulting excess supply of oil in the Midwest has therefore decreased the WTI price relative to on the water benchmark locations.

This paper studies the implications of this unusual spatial crude oil price divergence for prices of refined products in the U.S., focusing on prices for gasoline and diesel. In theory, a range of outcomes are possible. At one extreme, regional differences in crude prices might pass through completely to regional gas and diesel prices, so that wholesale prices for these refined products in the Midwest are substantially lower than prices along the coast. At the other extreme, the decrease in Midwest crude oil prices may not be passed through at all. As we discuss below, the realized outcome depends primarily on the availability and locations of refining capacity and refined product pipeline capacity.

We find strong evidence that the decrease in the U.S. Midwest crude oil price has not been passed through to the price of gasoline or diesel. This result is shown graphically in figure 2, which plots monthly average wholesale gasoline prices for both the Midwest and Gulf Coast (we will later discuss the source of these data in more detail). These two gasoline price series fol-

¹ See, for instance, “Enbridge, Enterprise to double Seaway pipeline capacity,” *Reuters*, March 27, 2012, <http://www.reuters.com/article/2012/03/27/us-enbridge-idUSBRE82Q03520120327>.

Figure 1: Brent, West Texas Intermediate (WTI), and Light Louisiana Sweet (LLS) Spot Prices for Crude Oil




Notes: Data plotted are monthly averages of daily spot prices. Delivery points are Cushing, Oklahoma for WTI, North Sea for Brent, and St. James, Louisiana for LLS. Source: *Bloomberg*.

low each other closely, and we will later show that, per this figure, there is no evidence that depressed oil prices in the Midwest caused gasoline or diesel prices to decline in the Midwest relative to other U.S. locations. We show that this lack of pass through is closely related to several other empirical findings: (1) Midwest refiners are operating at or near their capacity; (2) there has been no build-up of refined product inventory in the Midwest despite a build-up of crude oil inventory; and (3) the Midwest continues to be a net importer of refined product from coastal areas. These facts together are consistent with a model in which Midwest refiners are consuming as much low-priced Midwest crude oil as they have capacity to handle, yet imports of gasoline and diesel from the coast are still necessary to satisfy demand for refined products in the Midwest.

Our results have two main implications. First, the lack of pass-through implies that refiners, not consumers, are receiving the rents generated by depressed crude oil prices in the Midwest.² This outcome does not imply that Midwest refiners are exerting market power, but rather that they are operating at a near-vertical part of their supply curve given by their capacity constraint. Second, concern has been expressed that investments in pipeline capacity intended to

² This rent allocation is consistent with press reports. See, for instance, Albert (2012).

Figure 2: PADD 2 (Midwest) and PADD 3 (Gulf Coast) wholesale gasoline prices




Notes: Data are PADD-month level prices for “gasoline for resale”. PADD denotes Petroleum Administration for Defense District. Source: Energy Information Administration (EIA).

allow Midwest crude oil to reach the Gulf Coast will increase not only crude oil prices in the Midwest but gasoline and diesel prices as well (Verleger 2011). Our analysis indicates that, because Midwest crude oil price shocks are not passed through to refined product markets, increasing the price of crude oil in the Midwest will not increase gasoline prices there. Instead, the increase in crude prices will be borne by Midwest refiners.

The remainder of the paper proceeds as follows. Section 2 provides institutional background and lays out a simple theory of pass through in local crude oil and refined product markets. Section 3 presents our main evidence that the locally depressed crude oil price in the Midwest has not been passed through to refined product prices. Section 4 presents supporting evidence that: (a) the marginal gallon of refined product in the Midwest is imported; (b) refinery utilization in the Midwest is very high and likely constrained; and (c) while inventories of crude oil the Midwest have risen, inventories of refined product have not. Section 5 concludes by discussing implications of our results.

Figure 3: Map of Petroleum Administration for Defense Districts (PADDs)



Source: EIA

2. Institutions and Model

For most of the last century, the most cost-effective way to transport crude oil and refined petroleum products to and from landlocked parts of the U.S. has been in pipelines. A complex web of petroleum pipelines covers the country. For technical reasons, crude oil and refined products must travel in separate pipelines and those pipelines can flow only in one direction. Both of these attributes can be changed, but doing so requires months or years and costs millions of dollars. Where water borne transport is available, it is quite cost effective, but that is not an option for transport to landlocked locations. Crude and refined products can also be transported by rail and truck, but these are much more expensive options and are generally used only for short distances and in areas with sparse demand.

During World War II, the U.S. government designated Petroleum Administration for Defense Districts (PADDs) to aid in planning and allocation of oil and refined products. The country was divided up into five PADDs, of which the Northeastern PADD has since been divided into three sub-PADDs. The result is shown in figure 3. The PADDs correspond roughly to areas within which transportation is relatively unconstrained and between which bottlenecks or transport barriers potentially exist. Until recently, however, PADDs 1, 2, and 3 were considered to be quite well integrated with one another at nearly all times. PADD 3 (New Mexico, Texas and

most of the Gulf Coast) has been the primary oil production and refining area of the U.S. as well as the primary receiving point for imported crude oil, receiving about half of all crude imported to the U.S. Historically, PADD 3 has exported both crude oil and refined products to PADDs 1 and 2, which have had very limited oil production and modest refining capacity. The prices of crude oil at the major pricing locations across these PADDs have until recently stayed very much in sync with one another, as shown in figure 1. In fact, the primary crude oil pricing point in the U.S. and the delivery point for futures contracts traded on the New York Mercantile Exchange (NYMEX) has been Cushing, Oklahoma, in PADD 2, because many pipelines that carry crude converge at Cushing.³ This location was therefore thought to minimize basis risk due to its minimal transportation constraints to major oil markets.

Beginning around January 2011, the tight link between crude oil prices in PADDs 2 and 3 changed with greatly increased oil production from the Bakken oil shale formation in North Dakota and the tar sands area of Alberta, Canada. Figure 4 illustrates this production increase and shows that it accelerated, particularly in North Dakota, in early 2011. The primary transport route for oil produced in these regions is through pipelines that carry it to Cushing and from there to the Gulf Coast. The increased production has created a glut of supply at Cushing that exceeds the pipeline capacity from there to the Gulf Coast.⁴ As a result, the price of crude oil in Cushing has declined relative to the Gulf Coast (and North Sea) price, as shown in figure 1.⁵


We study the economic incidence of this bottleneck, which has caused virtually identical quality crude oil to be 20 to 65 cents per gallon less expensive at Cushing and elsewhere in PADD 2 than along the Gulf Coast since the beginning of 2011. While the costs of this bottleneck are borne by oil producers, it is not clear *a priori* who benefits from the reduced price of

³ See <http://www.neb.gc.ca/clf-nsi/rnrgynfmtn/nrgyrprt/lsnd/pprntnsndchllngs20152006/mg/cndntdsttlppln-eng.jpg> for a map of North America pipelines.

⁴ An alternative story behind the oil glut would be that there was a demand decrease in PADD 2 rather than a supply increase; however, the data strongly support the latter explanation. Crude production from the Alberta tar sands and PADD 2 increased by about 60% from January 2006 to December 2011. Consumption of petroleum products in PADD 2, in contrast, has decreased by only 8% (to January 2012), and most of this decrease occurred during the financial crisis before the glut began in early 2011 (consumption data are from the EIA's "product supplied" data series). Moreover, PADD 2 refinery utilization has not changed substantially relative to the rest of the U.S. (as shown in figure 8 below), nor has refinery throughput, and PADD 2 has continued to be a net importer of refined product (as shown in figure 6). Thus, it seems clear that the supply glut and price differential are due to a supply increase in PADD 2 (and Alberta), not to a demand shock in PADD 2.

⁵ Other modes of transportation—particularly rail transport—that are more expensive than pipeline transportation may be mitigating the impact of the pipeline bottleneck on the PADD 2 to PADD 3 price differential. Indeed, Peters and Lefebvre (2011) reports that investment in Midwestern railroad tank car and depot capacity has substantially increased since early 2011. Because estimates of the cost of railroad transportation (\$0.12 to \$0.24 per gallon (Peters and Lefebvre 2011)) are smaller than the observed PADD 2 to PADD 3 price differential, it appears that either railroad transportation is capacity constrained (on the track or at loading/unloading terminals) or that railroad owners are exerting market power.

Figure 4: Midwest and Canadian tar sands oil production



Notes: U.S. data were sourced from the Energy Information Administration (EIA). Canadian data were sourced from the National Energy Board of Canada and include upgraded bitumen and heavy crude oil produced from Alberta.

crude oil in the Midwest. To the extent that this local price reduction is passed through to prices for refined products—primarily gasoline and diesel fuel—then Midwest consumers benefit. Otherwise, the benefits are captured by Midwest refiners.

A simple model of refined product supply and demand

To see how a broad range of pass-through rates are theoretically possible and to understand how pass-through relates to other observables such as PADD-level trade flows, refinery throughput, and inventories, we consider a simple two-region model in which production, transportation and refining are all supplied competitively. We call the two regions PADD 2 and ROW, for rest of world. We assume that the ROW market is very large, so that production and refining shocks in PADD 2 do not significantly affect ROW prices, even if transportation capacity is unconstrained. This assumption reflects the facts that the crude oil production of PADD 2

constitutes only a small share of global crude production and that refining capacity in PADD 2 is an even smaller share of global refining capacity.⁶


We assume that transportation has a simple technology of sunk pipeline costs and zero marginal costs, though the qualitative analysis isn't changed if there are small marginal transport costs (relative to the value of the product). For a given cost of crude, refineries have essentially constant marginal cost (MC) for quantities well below their capacity constraint, strictly upward sloping MC as quantity approaches capacity, and then vertical MC at capacity. For the purpose of this theoretical discussion, we refer to all refined product output as "gasoline," though the term is meant to include diesel and other refined products.

In figure 5, D_2 is the native demand for gasoline in PADD 2. There is pipeline export capacity for gasoline and a different import capacity. These capacity constraints are shown as horizontal segments in D_{net} , the net demand for gasoline in PADD 2, accounting for imports and exports. These segments are horizontal because if the price in PADD 2 is less than the price in ROW, P_{ROW} , the export pipeline will be used to capacity and the net demand for gasoline in PADD 2 will be the native demand shifted out by the capacity of the export pipeline. If, on the other hand, $P_2 > P_{ROW}$, the import pipeline will be fully utilized, and net demand in PADD 2 will be shifted in by the import pipeline capacity.

We consider four distinct possible supply situations, labeled a, b, c, and d in figure 5. For all four cases, we assume that the constant MC for PADD 2 production quantities well below capacity is at a level below the ROW price of gasoline. This assumption is likely to hold for two reasons. First, we are considering only cases in which crude in PADD 2 is as cheap or cheaper than in ROW. Second, refineries are very capital intensive so that equilibrium prices will typically include some quasi-rents to capital owners in the short run.

⁶ During 2011, the bottleneck between PADD 2 and the Gulf Coast likely reduced supply outside PADD 2 by an average of less than 500,000 barrels per day (this value is the total increase in North Dakota's and Canada's daily crude oil production between January and December 2011, according to the EIA and Canada's National Energy Board) in a global market of about 90 million barrels per day. The small impact of this supply restriction is unlikely to be detectable in world oil prices (such as that for Brent crude). Relieving the bottleneck would almost surely reduce Gulf Coast and rest-of-world crude prices very slightly and increase the Cushing crude price substantially, until it was once again very close to the Gulf Coast price. For similar reasons, relieving the bottleneck is also unlikely to raise overall U.S. (i.e., not just PADD 2) and world gasoline prices. A recent report (Swift 2012) claims that the increase in the PADD 2 oil price following a debottlenecking would reduce PADD 2 refinery utilization, and that because PADD 2 refineries are particularly tuned to produce large amounts of gasoline (rather than diesel), the resulting decrease in PADD 2 gasoline production will increase overall U.S. gasoline prices by decreasing the world supply of gasoline. However, as with the crude oil market, the large size of the world gasoline market implies that the price effect of a decrease in PADD 2 refinery utilization would be very small in magnitude. Moreover, our findings suggest that the change in PADD 2 refinery utilization following debottlenecking would itself be very small, as utilization increased only slightly when the Midwest oil glut began in 2011 (see figure 8).

Figure 5: Simple model of four cases for the PADD 2 refined product market



As shown in figure 5, the equilibrium PADD 2 gasoline price and trade flows depend on where the PADD 2 gasoline supply curve is relative to net demand. In case (a), the gasoline refining capacity in PADD 2 is small relative to demand so that, even utilizing the full import capacity for gasoline, the gasoline price in PADD 2 is still higher than in ROW. In case (b), PADD 2 refining capacity is larger so that, even though PADD 2 must still import gasoline, the import pipeline capacity is unconstrained. In this case, arbitrage implies that the PADD 2 gasoline price must equal the ROW price. Case (c) is similar to case (b) in that pipelines are unconstrained, except now refining capacity is sufficiently large that PADD 2 is a net exporter of gasoline rather than a net importer. Finally, in case (d) the PADD 2 refining capacity is so large that the export pipeline is at its capacity. In this case, the PADD 2 gasoline price will be below the ROW price.

To see how the different cases have different implications for pass-through, consider an exogenous shock (such as an improvement in crude oil extraction technology) that creates a large

increase in oil production in PADD 2 so that the PADD 2 crude price falls relative to ROW.⁷ Will this price decrease be passed through to the PADD 2 gasoline price? The decrease in the PADD 2 crude price will shift down the refiners' supply curve. Thus, if the market is initially in case (b) or (c) and remains in one of these two cases following the crude oil shock, then pass-through will be zero because the PADD 2 gasoline price will remain tied by arbitrage to the ROW price. PADD 2 refinery throughput may, however, increase to the extent that refiners are on the upward sloping rather than vertical part of their supply curve.

If, on the other hand, the market is initially in case (d), there will be at least partial pass-through: the export pipeline constraint prevents gasoline market arbitrage so that the PADD 2 supply curve shifts down along the downward sloping PADD 2 demand curve, reducing the PADD 2 gasoline price and increasing PADD 2 refinery throughput. Pass-through will be 100% if both the initial and final equilibria are such that refiners are operating on the flat part of their supply curve (the specific case shown in the figure). Partial pass-through occurs when refiners are on the upward sloping part of their supply curve (or if the market is initially in case (c) but is pushed by the supply shock into case (d)).

If the market is initially in case (a), in which the gasoline import pipeline is capacity constrained, there are two possibilities. If, as drawn in Figure 5, PADD 2 refineries are also initially capacity constrained, the decrease in the PADD 2 crude price will not pass through to the PADD 2 gasoline price, nor will refinery throughput be affected. Partial pass-through is possible, however, to the extent that PADD 2 refineries are initially on the upward sloping portions of their supply curves. In this case, refinery throughput will increase.

To summarize, so long as there is unconstrained gasoline pipeline capacity between PADD 2 and ROW, we would expect to see essentially no pass-through of PADD 2's decrease in crude oil prices. If pipeline export of gasoline from PADD 2 is (or becomes) constrained, some or all of the crude price drop will pass through to PADD 2 gasoline prices. In the final case (a) in which pipeline imports of gasoline to PADD 2 are constrained, zero or partial pass-through may occur. In each of these cases, PADD 2 refinery utilization may increase. Thus, an examination of trade flows will have more power to discern which case is in effect than will an examination of refinery throughput.

Finally, the different cases have implications for how PADD 2 gasoline inventories may be expected to change following the decrease in PADD 2 crude prices. The incentive to inventory

⁷ We recognize that the realized decrease in the PADD 2 oil price may be partially mitigated by downstream responses, particularly increases in PADD 2 refinery throughput. That response would not change the qualitative effect of the increased crude oil supply, however.

gasoline arises from the difference between the expected future gasoline price and the current price in PADD 2. If import and export capacities are unconstrained (cases (b) and (c)), then these prices will be the expected future and current prices in ROW. Thus, in these cases we should expect that changes in PADD 2 inventories will be similar to changes in ROW. Further, to the extent that firms hold no-change forecasts of future ROW prices, we should see essentially no change in inventories.⁸

In case (d), however, the decrease in the PADD 2 crude price reduces the PADD 2 gasoline price relative to ROW. In this case, the change in PADD 2 inventories will be determined by beliefs about the future gasoline price in PADD 2. Given a belief that investments in crude export pipeline capacity will increase the future PADD 2 oil price back toward parity with the ROW price—increasing the future PADD 2 gasoline price as well—firms will have an incentive to increase their gasoline inventories more in PADD 2 than in ROW. A similar incentive will exist in case (a) to the extent that the initial decrease in the PADD 2 oil price is passed through to the PADD 2 gasoline price.

3. Data and Analysis of Pass-through

In this section, we test directly whether, and to what extent, the decrease in PADD 2 oil prices was passed through to PADD 2 prices for refined products. Section 4 then examines PADD 2 trade flows, refinery throughput, and inventories to assess the extent to which they change in a way that is consistent with the theoretical model discussed above.

Most of the data for this study come from the U.S. Energy Information Administration (EIA). Data on wholesale prices for refined products (called “sales for resale” in the EIA data) are available monthly at both the state and PADD level.⁹ Spot prices for crude oil were obtained from Bloomberg; we calculate monthly spot crude oil prices by taking unweighted averages of the daily raw data. All data cover the period 2006 through 2011.

Inspection of figures 1 and 2 suggests that pass-through of the post-2010 decrease in PADD 2 crude oil prices has been very limited. Here, we study pass-through more formally with two types of regression-based tests. Our first test is that given by specification (1) below, which examines the extent to which changes in the difference between PADD 2 and PADD 3 crude oil

⁸ Alquist, Kilian, and Vigfusson (2010) shows that no-change forecasts for real crude oil prices generally outperform all other forecasts, at least at horizons greater than one year, and Anderson, Kellogg, and Sallee (2011) shows that consumers generally have no-change forecasts of the future price of gasoline.

⁹ “Sales for resale” denotes prices for sales on wholesale markets to “purchasers who are other-than-ultimate consumers” according to the EIA. We use the “motor gasoline” price series for gasoline and the “no. 2 distillate” series for diesel. Monthly average prices are volume weighted.

prices correlate with changes in the difference between PADD 2 and PADD 3 refined product prices.

$$G_{2t} - G_{3t} = \beta_0 + \beta_1(C_{2t} - C_{3t}) + \varepsilon_t \quad (1)$$

In (1), G_{2t} and G_{3t} denote the average prices of gasoline in PADDs 2 and 3 in month t , while C_{2t} and C_{3t} denote the WTI and LLS prices, respectively, for crude oil in month t . The parameter β_1 denotes the extent to which changes in the PADD 2 to PADD 3 crude price differential are passed through to the PADD 2 to PADD 3 gasoline price differential.

Estimating equation (1) by ordinary least squares raises the issue of the endogeneity of the crude price differential. If there is a shock to local demand for refined product in PADD 2 relative to PADD 3 that affects the relative refined product prices, then that shock also may change the relative crude prices between the PADDs. Such endogeneity, however, would clearly bias estimates in the positive direction. It is hard to see how endogenous crude price could incorrectly lead to the conclusion that the pass through is zero. Moreover, given the substantial increase in crude oil production in PADD 2 during the time period under study, it seems likely that regional shocks to crude oil markets are much more substantial than regional gasoline market shocks during this time. We therefore treat these regressions and those that follow as measuring the impact of shocks to crude oil differentials on refined product differentials.

Estimates of specification (1) are given in table 1, column 1. The estimate of β_1 is small in magnitude at -0.003, statistically insignificant, and precisely estimated, with a standard error of 0.026.¹⁰ The 95% confidence interval bounds the impact of a \$1 per gallon increase in the WTI – LLS crude price differential on the PADD 2 – PADD 3 gasoline price differential at between -\$0.05 and +\$0.05. This result accords with figures 1 and 2 and supports the inference that the decrease in Midwest crude oil prices that began in 2011 has not passed through to gasoline markets.

One potential concern with the estimate of specification (1) is that it matches PADD-level averages of gasoline prices to prices for crude oil that are for delivery to very specific locations: Cushing, Oklahoma and St. James, Louisiana. We address this concern by re-estimating specification (1) using only gasoline prices for Oklahoma and Louisiana. The results of this regression are given in column II of table 1. Again, the estimate of β_1 is small in magnitude, statistically insignificant, and precisely estimated, consistent with no pass-through.

¹⁰ Standard errors for specification (1) are Newey-West with 12 lags to account for seasonal correlation in the residuals. These standard errors are sometimes smaller and sometimes larger than when only Eicker-White standard errors are used. For example, the Eicker-White robust standard error for β_1 in both columns I and II is 0.031.

Table 1: OLS regressions of PADD 2 to PADD 3 refined product price differentials on PADD 2 to PADD 3 crude oil price differentials

	I	II	III	IV
	Gasoline price differences		Diesel price differences	
Coefficient on covariate:	PADD 2 minus PADD 3	Oklahoma minus Louisiana	PADD 2 minus PADD 3	Oklahoma minus Louisiana
WTI crude price minus LLS crude price	-0.003 (0.026)	-0.047 (0.041)	0.027 (0.026)	0.048 (0.057)
Constant	0.043 (0.010)	0.012 (0.017)	0.063 (0.011)	0.069 (0.026)
R ²	0.0001	0.012	0.009	0.011
N	72	72	72	72

Notes : Crude prices are monthly averages of daily spot prices, all in \$/gallon. Delivery points are Cushing, Oklahoma for WTI and St. James, Louisiana for LLS. Gasoline and diesel prices are month-level prices for "sales for resale" obtained from the EIA. Data span 2006-2011. Standard errors are Newey-West with 12 lags.

Columns III and IV repeat columns I and II but use prices for diesel rather than prices for gasoline. In line with our gasoline price results, we find no evidence that low Midwest crude oil prices have passed through to markets for diesel.

Our second specification is given by equation (2) below, which studies how PADD 2 gasoline prices are determined by PADD 3 oil prices and the difference between PADD 2 and PADD 3 oil prices. We estimate this equation in first differences both because we are interested in the relationship between gasoline and oil prices in the short-run rather than long-run (when pipeline construction will tend to arbitrage spatial price differences away) and because oil and gasoline prices are generally found to be very persistent, so that a unit root cannot be rejected (see, for example, Borenstein, Cameron, and Gilbert 1997 and Chen, Finney, and Lai 2005).

$$\Delta G_{2t} = \beta_0 + \beta_1 \Delta C_{3t} + \beta_2 \Delta(C_{2t} - C_{3t}) + \varepsilon_{2t} \quad (2)$$

As with specification (1), C_{2t} is given by the WTI crude oil price, and C_{3t} is given by the LLS crude oil price. We estimate equation (2) using state-by-month level gasoline prices as the dependent variable so that the regression uses a panel across the 15 states in PADD 2 over 2006-2011. We cluster standard errors on month to account for cross-sectional correlation of residuals.

Results from estimating equation (2) are given in table 2. Column I uses gasoline prices as the dependent variable, while column II uses diesel prices. While the estimates are somewhat imprecise, we find in both regressions that changes in the PADD 3 crude oil price pass through

Table 2: First-differenced OLS regressions of PADD 2 refined product prices on crude oil prices and differentials

	I	II
Coefficient on covariate:	$\Delta(\text{Gasoline price})$	$\Delta(\text{Diesel price})$
$\Delta(\text{LLS crude price})$	1.074 (0.106)	1.048 (0.058)
$\Delta(\text{WTI crude price minus LLS crude price})$	-0.178 (0.204)	0.180 (0.209)
R^2	0.714	0.820
F	51.04	167.9
N	1063	1065


Notes : Crude prices are monthly averages of daily spot prices, all in \$/gallon. Delivery points are Cushing, Oklahoma for WTI and St. James, Louisiana for LLS. Gasoline and diesel prices are state-month level prices in PADD 2 for "sales for resale" obtained from the EIA. Data span 2006-2011. The difference in the number of observations between columns I and II is due to a missing gasoline observation for Kansas in 2006. Standard errors are clustered on month-of-sample.

essentially completely to PADD 2 refined products prices (the estimate of β_1 is statistically indistinguishable from one) but that pass-through of the difference between PADD 2 and PADD 3 crude prices, β_2 , is statistically indistinguishable from zero.

These results are consistent with cases (b) and (c) in section 2, in which Midwest refined product prices are tightly linked to the Gulf Coast, and therefore to world crude oil prices, because the pipelines carrying refined products between these areas are not constrained. As a result, changes in Midwest crude oil prices do not move Midwest refined product prices. Section 4 below provides additional evidence in support of these cases by studying data on product flows, refinery throughput, and inventories.

In the appendix, we present additional evidence showing that the pass-through results we obtained for Midwest refined product prices also hold for PADD 4, the Rocky Mountains area. Because PADD 4 is connected to the Gulf Coast only through PADD 2, PADD 4 has also experienced a relative decrease in crude oil prices, though not as large a relative decline because PADD 4 crude prices have been lower than elsewhere in the U.S. throughout our sample period. Figures A1 and A2 and tables A1 and A2 show that this decrease in PADD 4 crude oil prices has not passed through to PADD 4 prices for refined products.

Figure 6: PADD 2 net flows of gasoline and distillate (diesel)




Notes: Data are PADD-month level net imports (by tanker, pipeline, and/or barge) from other PADDs. The data do not include international trade. Source: EIA.

4. Why Hasn't Cheap Crude in the Midwest Lowered Gas Prices?

Why did the decrease in PADD 2 crude oil prices fail to pass through to PADD 2 gasoline and diesel prices? To some, this outcome may raise concerns of nefarious behavior on the part of PADD 2 refiners. The economic evidence, however, supports a benign explanation: figure 6 shows that PADD 2 has been a net importer of refined product for the entire period of our sample.¹¹ Moreover, net imports to PADD 2 were lower in 2011 than in the previous 5 years, suggesting that import pipelines were not constrained. This outcome corresponds to case (b) from section 2: the marginal gallon of refined product sold in PADD 2, both before and after the PADD 2 oil price decrease, has been produced outside of PADD 2 and therefore reflects the mar-

¹¹ Flow data for crude oil show that PADD 2 is also a net importer of crude. At first glance, these data appear to contradict the premise that PADD 2 is export-constrained in the crude oil market. However, the PADD 2 crude imports come from the Permian Basin in West Texas. While this area is technically part of PADD 3, it is economically part of PADD 2 in that there exists substantial pipeline capacity to carry oil from the Permian Basin to Cushing, Oklahoma but very little direct capacity to the Gulf Coast. Indeed, the industry broadly considers the price for Permian crude to be the WTI price at Cushing (see Alpert 2012, for instance). In contrast, the Permian Basin does not, to the best of our knowledge, export significant quantities of refined product to PADD 2.

Figure 7: Inventories of gasoline and distillate (diesel) in PADDs 2 and 3




Notes: Data are PADD -month level stocks held at, or in transit to, refineries and bulk terminals, and stocks in pipelines. Source: EIA.

ginal cost of more expensive crude that is processed elsewhere. The crude oil market has become separated between the PADDs, but the refined product markets have remained integrated.

This conclusion is also supported by data on inventories of refined product. Cases (b) and (c) predict no change in PADD 2 refined product inventories relative to other locations outside the Midwest, while case (d)—which implies substantial price pass-through—is likely to be accompanied by an increase in PADD 2 refined product inventories. Figure 7 shows that both gasoline and diesel inventories were fairly constant in PADD 2 during 2011, while they rose somewhat in PADD 3. There is no evidence of a buildup of PADD 2 product inventories in response to temporarily depressed product prices, as suggested by case (d).

The discussion in section 2 indicated that refinery utilization data would not help to distinguish among the possible cases because utilization could increase in any of them. Figure 8 does not reveal an obvious pattern in PADD 2 utilization versus PADD 3. At the end of 2011, however, PADD 2 utilization was nearly 5 percentage points higher than in PADD 3, and PADD 2 utilization was higher than it had been in the previous 5 years. This increase in PADD 2 utilization is consistent with the increase in PADD 2 refining margins since early 2011.

Figure 8: Operable refinery utilization in PADDs 2 and 3



Notes: Data are PADD-month level operable refinery utilization. Source: EIA.

These observations fit together in a unified picture. The increased production of crude oil in the upper Midwest and Canada over the past two years has led to a bottleneck in shipping crude out of PADD 2. PADD 2 has reduced net imports of refined products, but it has not become a net exporter, let alone a large enough net exporter to congest transport of refined product out of PADD 2. The ability to arbitrage refined product prices, but not oil prices, between PADDs has resulted in essentially no pass-through of the reduction in the PADD 2 oil price to PADD 2 product prices. This situation has created very large refining margins in PADD 2 and also appears to be causing higher utilization of refineries in PADD 2 relative to other parts of the U.S. This does not imply that refineries in PADD 2 have exercised market power; in fact, the high utilization of operable capacity suggests that they have responded to the higher margins by trying to squeeze out more output.¹²

¹² The story for PADD 4 is similar, except PADD 4 has varied between being a small net importer and a small net exporter of refined product. In 2011, PADD 4 net flows were close to zero. As in PADD 2, refined product inventories were flat in PADD 4 during 2011.

5. Conclusions and Implications

Beginning in early 2011, increases in crude oil production from North Dakota's shale resource and Canada's tar sands created a transportation bottleneck as the pipelines capable of carrying oil from the Midwest to the Gulf Coast reached full capacity. This constraint caused the benchmark Midwest crude oil price to fall substantially below the "world" oil price on the Gulf Coast, despite the fact that these two prices have historically been very close to one another. We show in this paper that this relative price change has not passed through to markets for refined products: Midwest wholesale prices for gasoline and diesel have not fallen relative to those along the Gulf Coast. This result is explained by the fact that the marginal gallon of gasoline (and diesel) in the Midwest is being imported from the Gulf Coast, where it is refined using relatively expensive crude oil. In other words, while trade in crude oil between the Midwest and Gulf Coast is capacity constrained, trade in refined products is not, and in fact the Midwest is actually importing rather than exporting gasoline and diesel.

Our results imply that the primary beneficiaries of depressed Midwest crude oil prices have been Midwest refiners rather than Midwest consumers. We emphasize that this outcome does not imply that Midwest refiners are exerting market power. While our analysis doesn't rule out that possibility, the empirical findings are equally consistent with refiners operating at or near their production capacity while benefitting from the fact that the marginal refined product suppliers in PADD 2—refineries on the Gulf Coast—are producing from more-expensive crude oil.

A simple "back of the envelope" calculation suggests that the rents accruing to Midwest refiners are substantial, relative to a counterfactual in which crude oil pipeline capacity were sufficient to equalize Midwest and Gulf Coast crude prices. The average Midwest to Gulf Coast crude price differential during 2011 was \$17.45. Given the average 2011 Midwest (PADD 2) refinery throughput of 3.39 million barrels per day, Midwest refiners therefore earned rents of \$59.2 million per day. Rents also accrued to holders of crude oil transportation rights: given average exports from the Midwest to the Gulf Coast of 195 thousand barrels per day, these rents are \$3.4 million per day. The countervailing losses of \$62.6 million per day have of course been borne by Midwest and Canadian crude oil producers.¹³

The substantial rents accruing to Midwest refiners and to holders of the limited Midwest crude oil export capacity strongly suggest that the present situation is not a long-run equilibrium.

¹³ To the extent that refinery demand and Canadian and PADD 2 crude oil supply are elastic, refineries' gains and producers' losses will be lower and greater, respectively, than these estimates. Given that most produced and refined crude is inframarginal and that demand and supply are likely to be quite inelastic, at least in the short to medium run, accounting for these elasticities will have only a very small effect on our calculations.

In fact, several investment projects have already been announced or are underway that would increase Midwest crude oil export capacity,¹⁴ including construction of the southern segment of the controversial Keystone XL pipeline (the northern segment would expand capacity from the Canadian tar sands to the Midwest). These projects will relieve the Midwest crude oil export bottleneck as they come on-line, bringing the Midwest oil price closer to, if not ultimately back into equality with, the Gulf Coast price. This re-equilibration will primarily increase the Midwest crude oil price rather than decrease the Gulf Coast price because the Gulf Coast is tied to the very large world oil market, of which the Midwest is only a small part.

The merits of these capacity expansions—particularly the Keystone XL project—have been a matter of public debate on both environmental grounds and the extent to which it will impact U.S. gasoline prices. While this paper is silent on environmental impacts, it does imply that the impacts on gasoline prices will be extremely limited. Because expanding Midwest crude oil export capacity will have only a minimal impact on Gulf Coast and world oil prices, U.S. consumers outside the Midwest will not experience a decline in gasoline prices. As for Midwest consumers, our results imply that capacity expansions that increase the Midwest crude oil price will not increase the Midwest gasoline price. This price is already being set by gasoline refined using Gulf Coast rather than Midwest oil, despite the depressed Midwest oil price. Resolving the Midwest crude oil transportation bottleneck will not affect this situation, thereby leaving Midwest gasoline prices unaffected as well.

In May 2012, the Seaway pipeline began carrying about 150,000 barrels per day of crude oil from Cushing to the Gulf Coast. By July, however, the crude price differential remained at about the same level as prior to that capacity addition – about \$15 per barrel – as production in the Midwest and Canada continued to increase. Refined product prices available through April 2012 show gasoline and diesel price differentials between PADD 2 and PADD 3 remained around their historical average, about five cents per gallon, unaffected by the difference in crude prices.

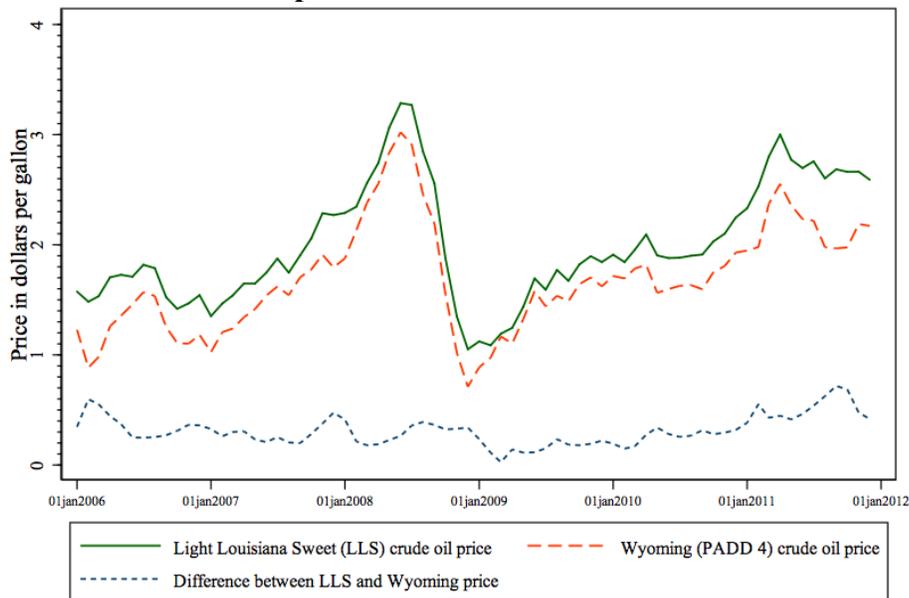
¹⁴ Another type of investment that would dissipate these rents would be additions to Midwest refinery capacity. The fact that investors are sponsoring crude oil pipeline projects rather than refinery projects suggests that the latter option is relatively expensive and/or time consuming.

References

- Albert, Bill, "Going with the Flow," *Barron's*, 31 March, 2012.
- Alquist, Ron, Lutz Kilian, and Robert J. Vigfusson, "Forecasting the Price of Oil," in Graham Elliot and Allen Timmerman, eds., *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland, forthcoming.
- Anderson, Soren T., Ryan Kellogg, and James M. Sallee, "What do Consumers Believe about Future Gasoline Prices?" NBER working paper #16974 (2011).
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert, "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics* 112 (1997), 305-339.
- Chen, Li-Hsueh, Miles Finney, and Kon S. Lai, "A Threshold Cointegration Analysis of Asymmetric Price Transmission from Crude Oil to Gasoline Prices," *Economics Letters* 89 (2005), 233-239.
- Peters, Mark and Ben Lefebvre, "Riding the Dakota Oil Boom," *Wall Street Journal*, 2 November, 2011.
- Swift, Anthony, "Keystone XL: A Tar Sands Pipeline to Increase Oil Prices," Natural Resources Defense Council report (2012).
- Verleger, Philip, "If gas prices go up further, blame Canada," *Star Tribune*, 13 March, 2011.

Appendix to “The Incidence of an Oil Glut: Who Benefits from Cheap Crude Oil in the Midwest?”

Figure A1: Wyoming (PADD 4) and Light Louisiana Sweet (LLS) Spot Prices for Crude Oil



Notes: Data plotted are monthly averages of daily spot prices. Delivery points are Guernsey, WY for Wyoming and St. James, Louisiana for LLS. Source: *Bloomberg*.

Figure A2: PADD 4 (Rocky Mountains) and PADD 3 (Gulf Coast) wholesale gasoline prices



Notes: Data sourced from the Energy Information Administration (EIA) and are PADD-month level prices for “gasoline for resale”. PADD denotes Petroleum Administration for Defense District.

Table A1: OLS regressions of PADD 4 to PADD 3 refined product price differentials on PADD 4 to PADD 3 crude oil price differentials

	I	II	III	IV
	Gasoline price differences		Diesel price differences	
	PADD 4 minus	Wyoming minus	PADD 4 minus	Wyoming minus
Coefficient on covariate:	PADD 3	Louisiana	PADD 3	Louisiana
Wyoming crude price minus LLS crude price	0.057 (0.112)	0.056 (0.136)	-0.040 (0.092)	-0.114 (0.138)
Constant	0.089 (0.043)	0.090 (0.055)	0.147 (0.043)	0.147 (0.069)
R ²	0.006	0.004	0.004	0.018
N	72	72	72	72

Notes : Crude prices are monthly averages of daily spot prices, all in \$/gallon. Delivery points are Guernsey, WY for Wyoming and St. James, Louisiana for LLS. Gasoline and diesel prices are month-level prices for "sales for resale" obtained from the EIA. Data span 2006-2011. Standard errors are Newey-West with 12 lags.

Table A2: First-differenced OLS regressions of PADD 4 refined product prices on crude oil prices and differentials

	III	IV
Coefficient on covariate:	$\Delta(\text{Gasoline price})$	$\Delta(\text{Diesel price})$
$\Delta(\text{LLS crude price})$	1.044 (0.124)	1.071 (0.086)
$\Delta(\text{Wyoming crude price}$ $\text{minus LLS crude price})$	0.096 (0.187)	-0.136 (0.201)
R ²	0.629	0.671
N	353	353

Notes : Crude prices are monthly averages of daily spot prices, all in \$/gallon. Delivery points are Guernsey, WY for Wyoming and St. James, Louisiana for LLS. Gasoline and diesel prices are state-month level prices in PADD 4 for "sales for resale" obtained from the EIA. Data span 2006-2011. Standard errors are clustered on month-of-sample.