Competition and pass-through of input cost shocks: Evidence from the U.S. fracking boom*

Erich Muehlegger †
Richard L. Sweeney ‡

May 2017

Preliminary and incomplete – Do not cite or circulate

Abstract

The advent of hydraulic fracturing lead to a dramatic increase in US oil production. Due to regulatory, shipping and processing constraints, this sudden surge in domestic drilling caused an unprecedented divergence in crude acquisition costs across US refineries. We take advantage of this exogenous shock to input costs to study the nature of competition and the incidence of cost changes in this important industry. We begin by estimating the extent to which US refining's divergence from global crude markets was passed on to consumers. Using rich microdata, we are able to decompose the effects of firm-specific, market-specific and industry-wide cost shocks on refined product prices. We show that this distinction has important economic and econometric significance, and discuss the implications for prospective policy which would put a price on carbon emissions. The implications of these results for perennial questions about competition in the refining industry are also discussed.

JEL Codes: H22, H23, Q40, Q54

1 Introduction

Imperfect competition plays a central role in determining the incidence of tax and input cost shocks. Yet, despite the importance of pass-through as a parameter of primary economic and policy importance, the empirical literature examining incidence and competition is still at an early stage. Two challenges have hindered empirical examination of this topic to date. First, to study how competition affects pass-through, researchers need access to detailed data on both prices and costs for all firms active in a market. Second, to make any causal claims based on the co-movement between these prices and costs, the econometrician needs exogenous variation in input costs both

*Both authors declare they have no interests, financial or otherwise, that relate to the research describe in this paper, nor do they have any current ties, directly or indirectly to the energy industry. This work has been supported by the Sloan Foundation and NBER Hydrocarbon Infrastructure and Transportation workshop. Assistance with the data from the Energy Information Administration, especially Joseph Conklin and Lawrence Stroud, is gratefully acknowledged.

†University of California - Davis and NBER, emuehlegger@ucdavis.edu.
‡Boston College, sweeneri@bc.edu.

within markets and across different market structures. In this paper, we address both of these concerns by focusing on petroleum refining, and provide new evidence on the relationship between competition and the transmission of cost shocks from one of the largest manufacturing industries in the United States.

The emergence of hydraulic fracturing (or "fracking") precipitated a nearly doubling of oil production in the United States between 2005 and 2015, abruptly reversing a decades-long trend towards reliance on foreign oil.[1] Before this new oil could be consumed, it had to first be processed into usable end products, like gasoline and diesel, at a refinery. The extent to which consumers benefited from the shale boom thus depends on how much of the resulting price declines were passed on to end users.

While oil markets were generally considered to be well integrated, a series of constraints limited the set of refineries positioned to benefit from the shale boom in the short run. First, a long-standing (and recently revoked) ban on the export of domestically produced crude oil caused US oil prices to diverge from the rest of the world. Second, shipping and logistical constraints temporarily trapped new production in producing regions, causing an unprecedented divergence in crude prices across regions withing the U.S. Third, the highly tailored nature of refinery operations meant that not every refinery could process these new crude streams. The net result of these three factors was a dramatic reduction in the input costs, often tens of dollars per barrel below world market, for some firms while other competing firms costs remained largely unchanged.

We exploit this asymmetry of the input cost shock to pass-through using detailed micro data on the universe of oil refiners in the United States from the the Energy Information Administration (EIA). For every firm that owns a refinery in United States, we observe monthly crude procurement costs (by region). Moving downstream, we also observe detailed production decisions for each refinery that firm owns, and realized prices and quantities for all refined products at the state level each month. Using this detailed data, we first document the importance of accounting for rival cost shocks and competition when estimating pass-through. When fine time-space dummies are included, we find that refiners are essentially unable to pass-through idiosyncratic cost shocks. However, when these controls are relaxed, we find that 13% of market specific costs and 98% of industry-wide cost shocks are passed through.

Motivated by these results, we estimate a reduced-form pass-through equation which models firm prices as an explicit function of own and rival costs. We find strong evidence that the nature of the cost shock and the degree of competition within a market are central to pass-through. Unlike the prior literature on the transmission of state-wide fuel tax shocks (e.g., Marion and Muehlegger (2011)), which finds evidence of full pass-through of a common tax shock, we find strong evidence that competition limits the ability of firms to pass localized input cost shocks onto consumers. When we instrument for costs, we find that firms are only able to pass-through 20-40% of firm-specific cost shocks (consistent with recent estimates from Ganapati et al. (2017)).

[1]What is commonly referred to simply as fracking is really the concurrent widespread commercialization of two related technologies, horizontal drilling and hydraulic fracturing. Additional background information provided in Section 3.

We then illustrate the implications of these results for the incidence of a prospective carbon tax on the industry. Refineries are the second largest industrial point sources of greenhouse gases in the U.S (behind the electricity sector) . Using data from the EPA Greenhouse Gas Inventory, we demonstrate that the carbon intensity of transportation fuels produced at domestic refineries varies by as much as 200% between the 25th and 75th percentiles of the carbon-intensity distribution. Since our estimates distinguish between firm-specific shocks distinctly from regional shocks, we can estimate the effect of a tax on carbon much more directly than other work relying on aggregate data and regional (at best) cost shocks. Applying such an analysis, we find that while a carbon tax on refineries would be almost fully passed through to consumers on average, the impact across refineries substantially.

This paper builds on two strands of the existing literature: a long literature that studies the transmission of cost shocks through the supply chain and a more recent one that studies the impacts of fracking on the oil and gas industries. Input cost pass-through has long been of interest of researchers in industrial organization and public finance. The majority of the empirical evidence suggest that industry-wide input price shocks or tax shocks are heavily passed onto consumers.Marion and Muehlegger (2011) finds that the incidence of per-gallon excise taxes fall fully on consumers and that changes in the spot price of oil prices are transmitted one-for-one through the supply chain.[2]

Specific to fracking the most closely related paper is Borenstein and Kellogg (2014)that studies the whether the decline in oil prices in the upper Midwest induced a contemporaneous decline in the retail price of refined fuels. It finds little evidence that lower oil prices in PADD 2 translated into lower product prices and interprets the results as evidence that Gulf Coast refineries supply the marginal gallon of fuel in the Midwest. The authors conclude that refiners on average capture the majority of the benefits of the glut of tight oil.

The rest of the paper proceeds as follows. In Section 2, we discuss how competition and the nature of input cost shocks are related to pass-through. We then provide background detail on the U.S. fracking boom and how the fracking boom affected crude oil input prices at refineries (Section 3). Section 4 discusses the data and empirical strategy. Section 5 presents the results, and Section 7 discusses and concludes.

2 Pass-through and imperfect competition

Weyl and Fabinger (2013) (WF) formalize a general model capturing the incidence of a tax on asymmetric, imperfect competitors. We first restate their main set of results and then adapt them to the case in which we are interested, imperfectly competitive refineries receiving firm-specific shocks to input costs. Following their exposition, let the vectorσ to denote a single dimensional strategic

[2] A second strand of the pass-through literature examines the speed at which the supply chain transmits input cost shocks. In the context of refined products, a number of papers examine the asymmetric speed at which the supply chain transmits oil price increases and decreases to wholesale and retail prices (see, e.g., Bacon (1991), Borenstein et al. (1997),Bachmeier and Griffin (2003)) . Coined "rockets and feathers," the papers document that retail prices respond quickly to oil price increases, but slowly to oil price decreases.

variable chosen by firms satisfying $q(\sigma) = q(p(\sigma))$ with firms maximizing $\Pi_i = p(\sigma_i)q(\sigma_i) - c(q(\sigma_i)$. WF represent the first order condition for the firm i as

$$\theta_i = \frac{m \frac{dq}{d\sigma_i}}{-q \frac{dp}{d\sigma_i}}$$

.

Where the vector m denotes the difference between each firm's price and marginal cost, and θ denotes the vector of conduct parameters from Genesove and Mullin (1998), varying from 0 (perfect competition) to 1 (monopoly). The numerator corresponds to the marginal non-pecuniary externality imposed by i's strategic variable and the denominator corresponds to the marginal pecuniary externality imposed by i's strategic variable. As an example, for differentiated products in which firms set Nash prices, the conduct parameter for firm i is a function of the margin-weighted diversion ratios, $\theta_i = 1 - \sum_{j \neq i} d_i^j \frac{m_j}{m_i}$where $d_i^j = \frac{dq_j}{dp_i} / \frac{dq_i}{dp_i}$. Intuitively, if a firm lowers it's price and all of the increase in sales comes at the expense of other firms, $\theta_i = 0$. Alternatively, if demand at other firms is unaffected by a price decrease, $\theta_i = 1$.

Of primary interest in our context is how input cost pass-through relates to: (1) the degree to which the cost shock is idiosyncratic to the firm and (2) whether the firms exposed to the input cost shock are differentially exposed to market competition. We define $c_i = \bar{\alpha} + \alpha_i$ as the marginal cost of production faced by firm i, which is the sum of a shared (market-wide) cost component ($\bar{\alpha}$) and a firm-specific component (α_i), and represent the optimal choice of the strategic variable as $\sigma_i^*(c_i, \sigma_j(c_j))$. The profit maximizing strategy chosen by firm i depends directly on its own marginal cost and indirectly on the marginal cost of its competitor, firm j. Following the notation in WF, we use $\boldsymbol{\rho_\alpha}$ to denote the pass-through of a shock α onto the vector of firm-specific prices. Formally, firm-specific and common cost shocks impact prices both directly and indirectly.

$$\boldsymbol{\rho_{\alpha_i}} = [\frac{d\boldsymbol{p}}{d\sigma_i} + \sum_{j \neq i} \frac{d\boldsymbol{p}}{d\sigma_j} \frac{d\sigma_j}{d\sigma_i}] \frac{d\sigma_i}{d\alpha_i}$$

$$\boldsymbol{\rho_{\bar{\alpha}}} = \sum_k [\frac{d\boldsymbol{p}}{d\sigma_k} + \sum_{j \neq k} \frac{d\boldsymbol{p}}{d\sigma_j} \frac{d\sigma_j}{d\sigma_k}] \frac{d\sigma_k}{d\bar{\alpha}}$$

Examining the equations for pass-through provide some guidance as to how the nature of the cost shock and the degree of competition relate to pass-through. A firm-specific shock α_i impacts the vector of prices directly through i's choice of strategy and indirectly through j's strategic response to i.

Of particular interest is how pass-through changes with respect to the features of the competitive environment in which a firm operates. Under standard models of competition, the signs of $\frac{d\boldsymbol{p}}{d\sigma_i}$ and $\frac{d\sigma_i}{d\alpha_i}$ will be identical. In addition, $|\frac{d\sigma_j}{d\sigma_i}| \in [0, 1]$ with values closer to 0 representing representing firms that are not close competitors and values closer to 1 representing more aggressive competition between the firms. These two points imply that an indirect shock magnifies pass-through in markets in which the strategies of i and j are strategic complements. The firm-specific cost shock spills over to firm i's competitor, causing them to adjust their strategic variable and raise prices. In markets in where the strategies are strategic substitutes, pass-through of an indirect cost shock is attenuated , as changes in j's strategic variable offsets the direct response of firm i.

In contrast, a common shock,$\bar{\alpha}$, is transmitted more fully onto retail prices as it affects both i and j's strategy directly. The first term of the expression is identical to the pass-through of the firm-specific shock to i. The second term reflects the incremental impact of the common shock on the input costs of firm j. As in the case above, the strategic compelementarity and substitutability either increase or attenuate the pass-through of an input cost shock.

2.1 Differentiated Nash in Prices

As an illustration, we consider a specific case directly relevant to the industry we study - differentiated Nash-in-price, where differentiation reflects the geographic nature of competition within the refining industry. In this context, the strategic variable σ_i represents the price set by a particular refiner. The only non-zero term in $\frac{d\boldsymbol{p}}{d\sigma_i}$ is the term corresponding to a firm's own price, simplifying equations for $\boldsymbol{\rho_{\alpha_i}}$ and $\boldsymbol{\rho_{\bar{\alpha}}}$ considerably. Denoting ρ^i_α to be the pass-through of shock α onto the price charged by firm i, we can examine represent the pass-through of firm-specific and market-wide shocks on a firm's own prices and the prices of its competitors as:

$$\rho^i_{\alpha_i} = \frac{d\sigma_i}{d\alpha_i};\ \rho^j_{\alpha_i} = \frac{d\sigma_j}{d\sigma_i}\frac{d\sigma_i}{d\alpha_i}$$

$$\rho^k_{\bar{\alpha}} = \frac{d\sigma_k}{d\bar{\alpha}} + \sum_{j\neq k}\frac{d\sigma_k}{d\sigma_j}\frac{d\sigma_j}{d\bar{\alpha}}\ \forall k$$

All the terms in the expressions above are positive: strategic complementarity implies that $\frac{d\sigma_j}{d\sigma_i} > 0$ for all pairs of firms and the comparative statics imply $\frac{d\sigma_i}{d\alpha_i} > 0$.

To relate pass-through of a firm-specific shock to competition, note that under a model of differentiated Nash-in-prices, we can simplify the solution to the firm's objective function to $-q_i\theta_i = m\frac{dq}{d\sigma_i}$. Differentiating with respect to α_i, we can represent the pass-through of a firm-specific shock onto its own price:

$$\frac{d\sigma_i}{d\alpha_i} = \frac{\frac{dq_i}{dp_i}}{\frac{\theta_i q_i}{p_i}[\epsilon_\theta + \epsilon_q] + \sum_j \frac{dp_j}{dp_i}\frac{dq_j}{dp_i} + \sum_j (p_j - c_j)\frac{d^2 q_j}{dp_i^2}}$$

which under assumptions of linearity and constant degree of competition, reduces to

$$\frac{d\sigma_i}{d\alpha_i} = \frac{\frac{dq_i}{dp_i}}{\frac{\theta_i q_i}{p_i}[\epsilon_q] + \sum_j \frac{dp_j}{dp_i}\frac{dq_j}{dp_i}}$$

The first term in the denominator includes both the degree of competition with in the industry (θ_i) and the elasticity of residual demand faced by firm i. All else equal, pass-through falls as either the elasticity of residual demand rises or as the market becomes less competitive. The second term is negative - strategic complements imply $\frac{dp_j}{dp_i} > 0$, but the amount firm i's sales fall when it increases its own price are greater than any increase in sales to competitors, implying $\sum_j \frac{dq_j}{dp_i} < 0$

3 Input cost shocks and the U.S. fracking boom

The previous section provided a theoretical framework for pass-through in the presence of imperfect competition, and highlighted the distinction between firm-specific and market-wide cost shocks. In this section, we provide background on the recent oil boom in the United States, and illustrate that the resulting crude price changes were in fact large and heterogeneous.

3.1 The fracking boom lowered U.S. crude prices relative to world crude prices

Hydraulic fracturing injects a mixture of sand, water and chemicals at high pressure into horizontally drilled shale formations. The pressure cracks the shale formation and releases previously unrecoverable natural gas and "tight oil" from the newly-created fissures in the shale. The rapid maturity of this technology in beginning around 2005 unlocked billions of barrels of previously uneconomical crude oil reserves. As a result, U.S. oil production nearly doubled during the ensuing decade.

This rapid reversal in U.S. crude production caused US prices to diverge from global prices.Prior to 2015, the United States prohibited the export of the vast majority of domestically produced crude oils in the name of energy security.[3] While this measure had been in place since the 1970's, equilibrium import and production patterns were such that domestic crude prices moved in lockstep with foreign prices for most of this period. Figure 1 graphs the West Texas Intermediate spot price and the Brent spot price, which are the benchmark crude prices for the United States and Europe, along with the spread between the two spot prices. From 2000 to 2010, the WTI spot price was only $1.41 per barrel more expensive on average. After the tight oil boom, the WTI spot price diverged from its historical position relative to Brent crude, trading at a $12 per barrel reduction on average between 2011 and 2015.

3.2 Initially, the fracking boom primarily affected refineries close to shale deposits

The extent of this divergence from global prices varied considerably within the United States, due to the highly uneven geographic nature of the fracking boom. In locations with oil-bearing shale deposits, oil production has increased tremendously.[4] In contrast, production continued to decline

[3] A handful of exceptions are allowed: (1) export of crude to U.S. territories, (2) export of North Shore crude, (3) export of California heavy oil, amongst others.

[4] In the past five years, oil production in North Dakota rose 400 percent due from fracking the Bakken shale play. The state now accounts for approximately 15 percent of U.S. oil production. Production increased almost three-fold in Texas (from 400 million barrels annually in 2009 to 1.15 billion in 2015) and a similar fraction in Colorado (from 30 million barrels to 86 million barrels). In Oklahoma and New Mexico, production nearly doubled over the five years to 128 million barrels and 121 million barrels in 2014, respectively.

Figure 1: Crude Oil Spot Prices ($/bbl)

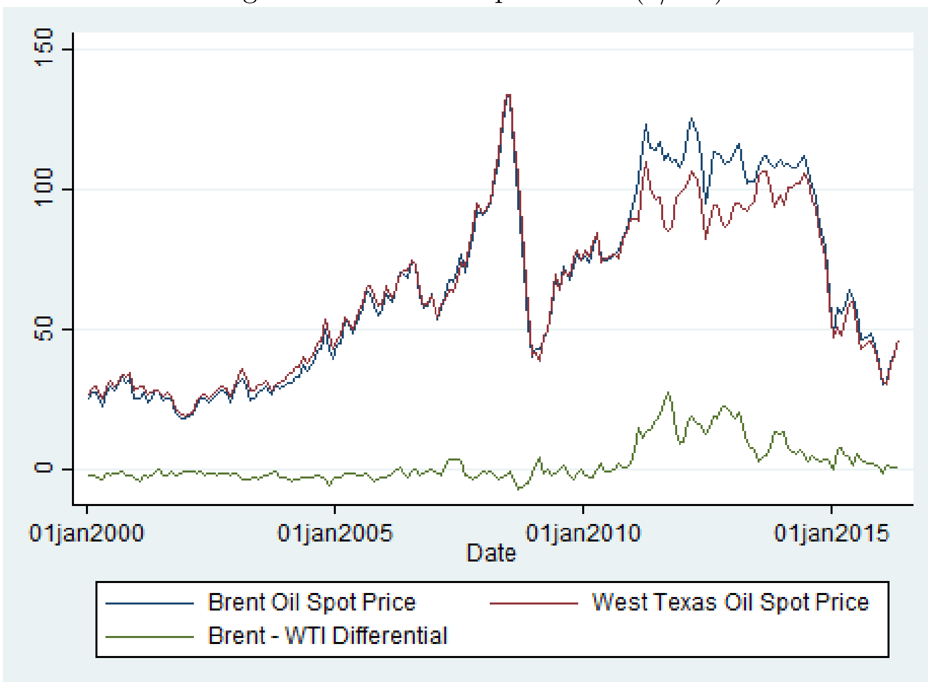

in other traditional oil deposits.[5] Crude oil transportation infrastructure has been slow to adjust to this this dramatic shift in domestic production, limiting the ability of oil producers to move product to locations with greater refining capacity (Figure 2).

The result was an unprecedented divergence in crude acquisition costs across refining regions within the United States. In Figure 3, we plot the average crude acquisition price discount (relative to the Brent spot) at refineries by Petroleum Administration Defense District (PADD).[6] Prior to 2010, crude acquisition prices in all five regions were reasonably close to the Brent spot price. After 2011, though, refinery acquisition costs in PADDs 2 (Midwest) and 4 (Rocky Mountain) begin trading at a deep discount, consistent with production of crude exceeding refinery and transportation capacity in these PADDs.

3.3 Even within region, realized cost reductions varied across refineries

Even within region, crude acquisition costs vary across firms. This is because crude oil is differentiated, primarily based on how dense it is and how much sulfur it contains. The characteristics in turn define what refining capacity is necessary to convert crude into usable refined products. Denser products contain smaller natural occurring shares of valuable end products (like gasoline), and therefore require additional processing. Due to it's corrosive properties and, more recently, harmful health effects post-combustion, sulfur must be largely removed from refined products be-

[5]Production from the Alaska and federal offshore deposits fell approximately 20 percent over same period.

[6]Petroleum Administration Defense Districts (PADDs), are a commonly used geographic aggregation for the industry dating back to World War II. The regions correspond to: (1) East Coast, (2) Midwest, (2) South, (4) Rockies, and (5) West Coast. A map of these regions is provided in Figure A.1.

Figure 2: Crude Oil Shipments from PADD 2

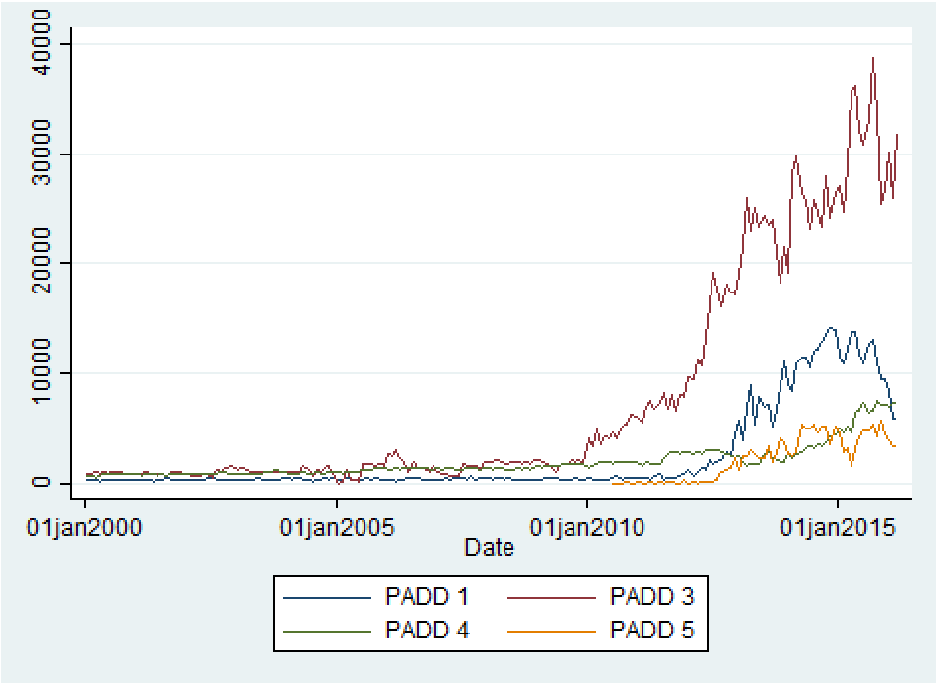

fore shipment, making “sour” (high-sulfur) crudes less valuable, all else equal. Figure 4 summarizes the distribution of domestic crude oil acquisition price over time within each PADD. We calculate crude oil acquisition costs relative to the Brent spot price - thus, declining values correspond to domestic acquisition costs at a greater discount relative to the the Brent crude spot price.

Refineries are highly tailored to process specific crudes, and substantial changes to the crude slate requires either months of reconfiguration or costly capital changes. This is important, because the fracking boom has largely increased domestic supply of “light” (low density) crudes. Figure 5 shows the price domestic oil producers received relative to Brent by tercile. Historically, lighter crudes traded at a slight premium, since they have larger naturally occurring shares of valuable end products. However, from 2010-2015, this long standing ordering was reversed, with the most valuable input trading at a substantial discount.

An additional source of within-region variation comes from the fact the some refineries get their crude from outside of the United States. This crudes often have subtle differences that make substitution away from them costly. Furthermore, many domestic refineries processing foreign crude are part of a vertically integrate international oil company like Citgo (Venezuela) or Aramco (Saudi Arabia). Figure 6 shows that despite the deep domestic discounts during this time, many firms continued to use imported oil.

Figure 3: Average Refinery Crude Price Minus Brent Spot by PADD

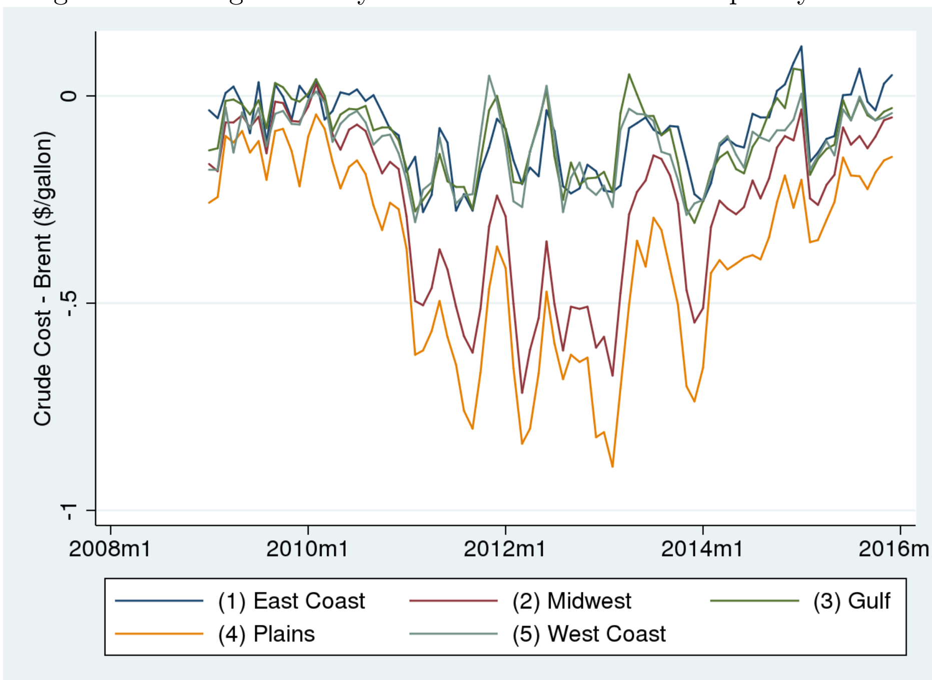

4 Data and Empirical Strategy

4.1 Data

Through a confidential data request we obtained several data sets on refinery operations from the Energy Information Administration (EIA), each described in appendix Table B.1. Beginning in 2004, every firm that owns a refinery in the United States reports total volume and total cost of crude oil acquired both domestically and abroad in each month (Survey EIA-14). On a separate survey, these firms report detailed input and output volume data, including information on the quality of crude used and source (domestic or foreign) each month (Survey EIA-810). This monthly refinery data is supplemented with an annual refinery survey which records detailed information about the capacity and technology installed at each refinery each year (Survey EIA -820).

This production data is then combined with a census of monthly state-level sales by every firm which owns a refinery in the United States (Survey EIA-782A). Refiners report sales in the state where the transfer of title occurred, regardless of where that product is ultimately consumed. Both the volume sold and the price are reported, broken out by sales to end users (retail) and sales for resale (wholesale).[7]

Despite the richness of the data, the different levels of spatial aggregation for reporting purposes necessitate additional assumptions to relate input costs to product prices. The primary challenge

[7] A related survey, EIA-782C, provides a census of all "prime suppliers", which includes firms that own refineries as well as large importers and marketers. 782C asks respondents to only report sales for which they are the final supplier into a state and the fuel is going to be consumed within state. The 782C data does not contain price and does not break volume sold into retail and wholesale.

Figure 4: Average Refinery Crude Price by PADD

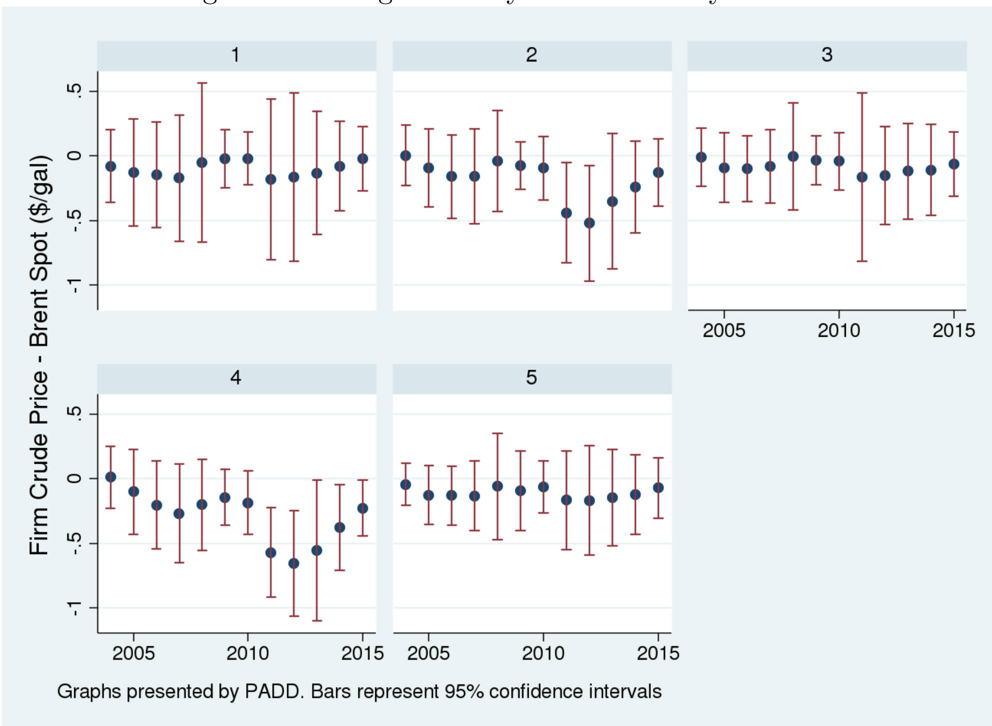

Figure 5: Domestic Crude Price Differentials by Density

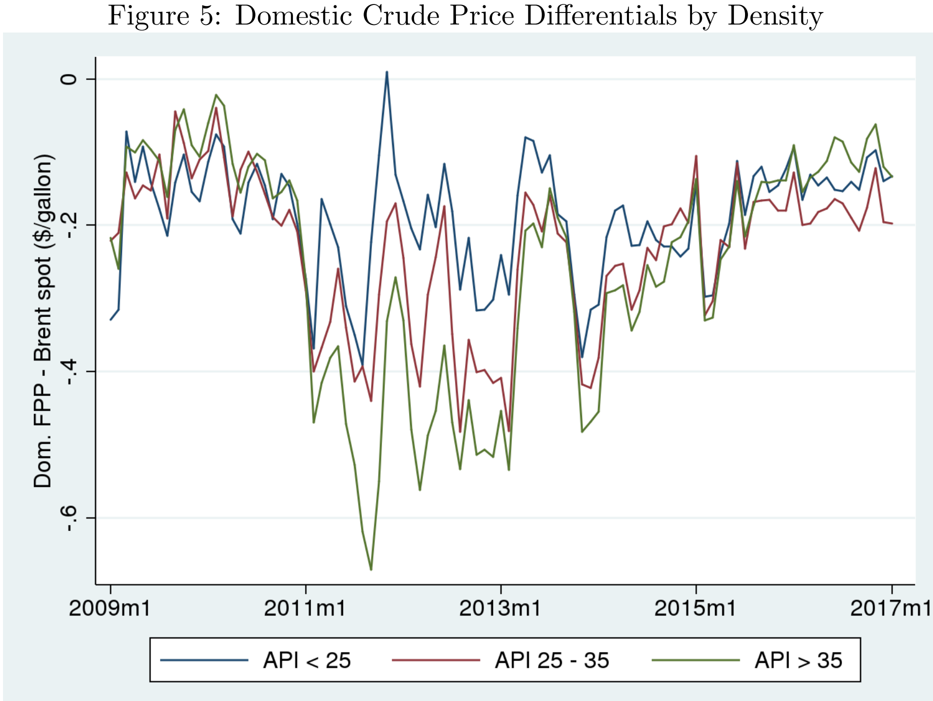

Figure 6: Domestic Crude Acquisition Shares by PADD

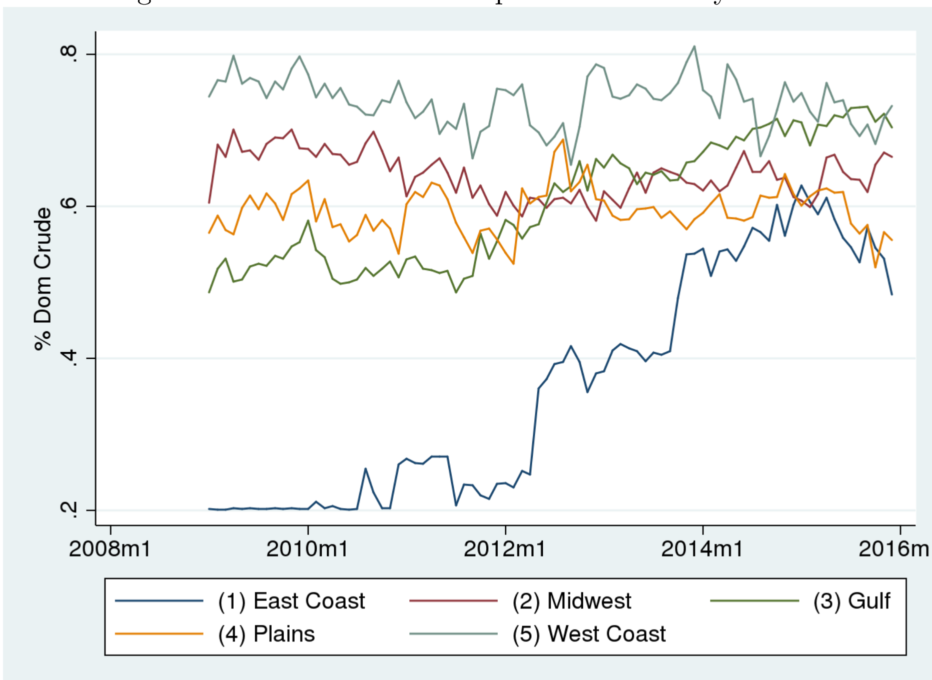

stems from the fact that firms own multiple refineries. Since crude costs are only reported at the firm-PADD level, we only observe the average input costs (and characteristics) for refineries owned by the same firm in the same PADD. A larger challenge stems from the fact that sales are reported at the firm level, not the refinery level. Thus, if firms own refineries in multiple regions, and we observe sales into a third region, we do not know which refinery's input costs should be associated with that sale. In these cases, we proceed by matching sales to the closest refining region reported in the crude price data each month.

Table 1 presents summary statistics from this combined data at the firm-PADD level.

Table 1: Summary statistics

	mean	sd
Crude cost (2013 $/gal)	1.875	0.577
Crude - Brent	-0.150	0.218
% Domestic	0.595	0.376
Price Gas	2.395	0.599
Price Diesel	2.512	0.676
Price Total	2.386	0.627
Resale Price Total	2.372	0.627
% Gas	0.466	0.215
% Diesel	0.345	0.205
% Resale	0.825	0.208
N	10229	

4.1.1 Choice of dependent variable

All refineries are multi-product firms. Table 1 reports the average monthly shares and prices for the two largest product, gasoline and diesel, and the total price, which includes jet fuel, other distillates and residual fuel as well. Thus, before studying pass-through, we first need to specify the refined product market of interest. To date the literature has typically focused on gasoline prices. However, given that all products are produced *jointly,* single product markups are misleading: simply subtracting gasoline prices from crude prices overstates the true markup because it doesn't take into account that less than 50% of the barrel could be converted to gas. The rest is sold at lower markups, but this does not mean refiners have less market power in those markets (the structure is the same). An alternative approach would be to sum the revenue across all products and use the average price as the dependent variable.

Table 2 explores the impact of this assumption on own-cost pass through. Monthly prices are projected onto crude costs, time dummies and firm-state fixed effects. In column (1) the dependent variable is the total revenue per gallon (including) all products, controlling for the fraction coming from gas and diesel. Column (2) and (3) use gasoline price and diesel price as the dependent variable respectively. Observed product-specific pass through rates are lower, likely reflecting profitable substitution between gas and diesel. Finally, column (4) repeats the regression in column (1) using wholesale revenue per gallon (i.e. excluding retail sales). Motivated by this table, the remained of this paper includes uses wholesale revenue per gallon as dependent variable of interest.

Table 2: Total vs product specific pass-through

	(1)	(2)	(3)	(4)
Crude Price	0.0523***	0.0329***	0.0313***	0.0576***
	(0.0126)	(0.00650)	(0.00577)	(0.0126)
Dep.Var.(Price)	Total	Gas	Diesel	Total
Firm-State FE	Y	Y	Y	Y
Controls	Y	N	N	Y
Resale only				Y
N	59002	57262	57057	59002
r2	0.968	0.980	0.987	0.967

The data include all firm-month-states with more than 1% of wholesale sales between 2004 and 2015. The dependent variable in each model is the price ($/gal), where price in models (1) and (4) is the weighted average revenue per gallon sold. All models include year-month and firm-state dummies. Models (1) and (4) include the fraction of sales coming from gasoline and diesel in the month. Standard errors are clustered at the firm-state level and presented in parentheses.

4.1.2 Choice of time fixed-effects

Empirical pass-through papers face an inherent omitted variables problem. The aim is to study how prices change in response to cost changes. However, costs could be correlated with unobserved demand shifters, confounding the preferred interpretation. To address this, most empirical pass-through studies include fine time-space dummies, effectively restricting cost variation to shocks which differentially affect close competitors. While this focus on firm-specific cost variation alleviates concerns about common unobserved factors, such as macro trends, it also returns correspondingly limiting estimates. The recovered pass-through parameters will only be suitable for analyzing prospective shocks that affect some firms but not others. However, if the goal is to estimate the incidence of a market- or industry-wide cost change, the firm-specific estimates could substantially understate the true parameter of interest.

To illustrate this point, Table 3 presents the results of several regressions of wholesale price on firm crude acquisition costs. Column (1) includes state-year-month fixed effects. The results suggest that firms have essentially zero ability to increase price within a state when their costs increase but their in-state rivals' don't. Column (2) includes as an independent variable the average cost of firms serving the state, and replaces the fixed effects with just year-month dummies to allow for this variation. The coefficient suggests that while idiosyncratic state supplier costs are not passed through, 13% of cost shocks that affect all suppliers in the state are. Column (3) includes national average crude costs, and relaxes the fixed effects to simply include year and month dummies. Importantly, the coefficients on own and state-specific cost shocks are unaffected by this change, suggesting that theoretically reasonable concerns about the correlation between macroeconomic unobservables and crude costs may be unwarranted.

Table 3: Pass-through Sensitivity to Fixed Effects

	(1)	(2)	(3)
Crude Price	0.0151	0.0156	0.0213
	(0.0120)	(0.0144)	(0.0147)
State Avg Crude Price		0.129***	0.124***
		(0.0246)	(0.0249)
US Avg Crude Price			0.842***
			(0.0224)
Firm-State FE	Y	Y	Y
Time FEs	St-Y-M	Y-M	Y,M
N	59002	59002	59002
r2	0.978	0.967	0.945

The dependent variable in each regression is the average (total) wholesale revenue per gallon received by firm in the state. The data include all firm-month-states with more than 1% of wholesale sales between 2004 and 2015. All models include firm-state dummies. Standard errors are clustered at the firm-state level and presented in parentheses.

The implications of this exercise for analyzing policies which broadly affect refineries are quite stark. Model (1) implies that the incidence of a cost-change is entirely borne by refiners, while summing the coefficients in model (3) reveals that on average any crude cost will be pass only almost entirely to consumers. Motivated by these results, we include only year and month fixed effects in the analysis that follows.

4.2 Empirical Strategy

We are interested in understanding how refiner f sets it's price p in market m at time t.[8] Specifically, if the price is equal to marginal cost plus markup, we want to know how markups change when costs change.

As demonstrated in section 2, realized prices will be a function of both own costs and competitors costs. The corresponding reduced-form pass through equation is,

$$p_{fmt} = \sum_{k \in F} \rho_{fkmt} c_{kt} + X'_{fmt}\delta + \mu_{fm} + \eta_{year} + \eta_{month} + \epsilon_{fmt} \tag{1}$$

where c are observable costs, X contains demand and supply side factors.[9] Miller et al. (2016) state that this linear functional form is exact with Bertrand-Nash and linear demand, and a first-order approximate to many other models.

As written, equation (1) is not estimable, since there are $F \times F \times M \times T$ pass-through parameters. To reduce the number of parameters, we first drop the t subscripts.[10] Next, again following Miller

[8]Note that refiners could own multiple refineries.

[9]In theory c should include all marginal costs of supplying market m at time t. In practice, we only observe average crude costs, which makeup over 90% of variable refinery costs.

[10]Note that this rules out complicated strategies that may vary over time.

et al. (2016), we further reduce parameter space by imposing that firm $f's$ pass-through of firm $k's$ costs in market m is a function of observables. Specifically, we put structure on ρ_{fkt} as follows:

$$\rho_{fkmt} = \begin{cases} \beta^0 & \text{if } k = f \\ \beta^1 & \text{else if firm } k \text{ sells into state } m \\ \beta^2 & \text{else if firm } k \text{ is on the fringe of } m \\ \beta^3 & \text{else} \end{cases}$$

So β^0 captures own-cost pass-through, β^1 is the rate at which the costs of other refiner's serving the same state are passed through, β^2 reflects pass through of refiners serving nearby states, and β^3 captures pass-through of far away competitors.[11] In some specifications we allow pass-through to vary within groups: $\beta^g = \sum \beta_l^g W_{lfkt}$.

5 Results

Table 4 presents the results from several specifications of equation (1). In column 1, pass-through is constant within each β group. The first coefficient has the interpretation that if a firm's crude costs increase by \$1 per gallon, the average price it receives in states it serves only increases by 6.6 cents if all other refiners' costs are held constant. Then, the coefficient on β^1 implies that if that firm's rivals serving the same state see their costs increase by \$1, firm f will increase it's price by 3 cpg. So direct rival costs are passed through at about half the rate of own costs on average. Turning to the estimates for other refiners, the rival cost pass-through rate declines by around 15% for refiners with the capacity to serve a given state (but not active in it that month), and another 15% for refiners outside market m's vicinity entirely.

[11] Specifically, the firms are included in the fringe (group 2) if they sell into a neighboring state at time t, or if they sell into market m or it's neighbor more than 50% of the active sample months.

Table 4: OLS Pass-Through Estimates

	(1)	(2)	(3)
β^0	0.0661***	0.146***	0.151***
	(0.0121)	(0.0223)	(0.0243)
β^0 Nfirms		-0.0113***	-0.0102***
		(0.00266)	(0.00258)
β^0 log(cap)			-0.00112*
			(0.000669)
β^1	0.0314***	0.0364***	0.0304***
	(0.00107)	(0.00365)	(0.00486)
β^1 Nfirms		0.000102	0.000110
		(0.000305)	(0.000309)
β^1 log(cap)			0.000443**
			(0.000180)
β^2	0.0269***	0.0251***	0.0214***
	(0.000782)	(0.00183)	(0.00231)
β^2 Nfirms		-0.0000794	-0.0000732
		(0.0000808)	(0.0000812)
β^2 log(cap)			0.000361***
			(0.000123)
β^3	0.0233***	0.0248***	0.0240***
	(0.000664)	(0.00153)	(0.00164)
β^3 Nfirms		-0.000112**	-0.000109*
		(0.0000565)	(0.0000566)
β^3 log(cap)			0.0000831
			(0.0000553)
Firm PT	0.0661	0.0538	0.0559
State PT	0.323	0.353	0.343
Fringe PT	0.513	0.530	0.519
US PT	1.002	1.047	1.036
N	59002	59002	59002
r2	0.944	0.944	0.944

Results OLS estimation of Equation (1) under three separate specifications of β^g. The dependent variable in each regression is the average wholesale price received by firm f in a given month. The data include all firm-month-states with more than 1% of wholesale sales between 2004 and 2015. All models include firm-state, year and month dummies, as well as controls for weather, income and population. Standard errors are clustered at the firm-state level and presented in parentheses.

While nearby rival cost pass-through is substantial relative to own-cost pass-through, the fact that they are not equal suggests some degree of product differentiation. If firms serving the same state were competing a la Cournot, own and rival cost shocks would have the same effect on price (to a first order approximation).

Although model (1) allows pass-through to vary by the proximity nests listed above, it imposes that pass-through within these nests be constant across markets. However, as can be seen in figure (7), the competitiveness of these markets varies considerably. Columns (2) and (3) allow pass-through to vary with the total number of firms in the nest and with the size of the firm experiencing the cost shock. Consistent with intuition, own-cost pass through declines with competition, and with firm size, while firms pass through a greater amount of larger competitors' costs.

Figure 7: Number of refiners serving state

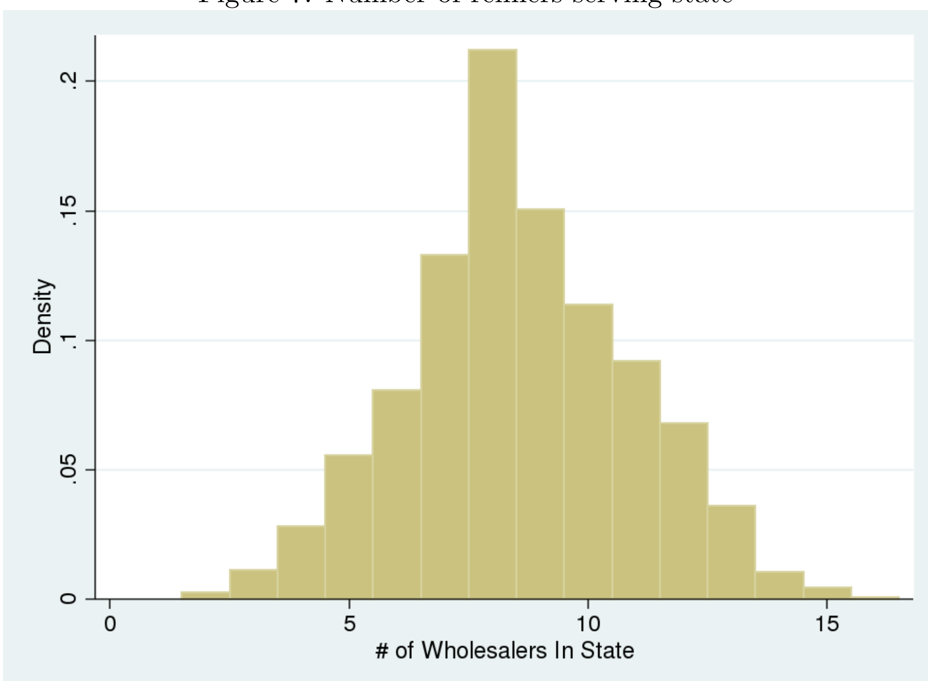

Based on annual data (2011-2015) in EPA GHGRP

Table 5 summarizes the heterogeneity from model 3 by region. The estimates are used to predict the impact of a \$1 cost increase that affects only firm f , only firms serving state m, all firms capable of serving state m, and the entire industry. These average of these predicted pass-through rates is then averaged by PADD for the entire sample.

Table 5: Pass-Through Estimates by PADD

	1	2	3	4	5
Firm PT	0.063	0.055	0.030	0.075	0.065
State PT	0.33	0.35	0.41	0.30	0.32
Fringe PT	0.44	0.55	0.64	0.54	0.46
US PT	1.03	1.04	1.04	1.04	1.04

The results from model 3 in Table 4 are used to predict the impact of a $1 cost increase that affects only firm f , only firms serving state m, all firms capable of serving state m, and the entire industry. The average of these predicted pass-through rates is then averaged by PADD for the entire sample.

5.1 Instrumental variables

One concern with the using OLS to estimate equation (1) is endogeneity. Observed prices are a function of both supply and demand. While all models include controls for demand shifters, firm-market and time fixed effects, if unobserved demand changes in a state also affect the input cost of refineries serving that state, then the results presented in Table 4 may be biased.

To address these concerns, we construct several instrumental variables building off the discussion in Section 3. Our approach is based on Bartik (1991). As a proxy for the divergence from world crude markets, we first calculate the WTI-Brent price differential each month (Figure 1) . We then interact this measure of the shale boom with one year lagged domestic crude share, and one year lagged API gravity. We then sum these instruments and interact with pass-through shifters to mirror the modeled structure on ρ_{jk}. Table 6 presents the results from re-estimating model (1) in table 4 using two-staged least squares. In column (1), only own costs excluded from the second stage. In column (2), the costs for refiners serving the same state are excluded as well, and in column (3) the costs of all firms capable of serving each market are instrumented for. In all three models, the firm-specific cost shock pass through are significantly larger than would be estimated using OLS. The results are also centered around the pass-through estimates from refineries obtained by Ganapati et al. (2017).

Table 6: IV Pass-Through Estimates

	(1)	(2)	(3)
β^0	0.416***	0.189***	0.248***
	(0.0761)	(0.0561)	(0.0603)
β^1	0.0170***	0.0314***	0.0295***
	(0.00318)	(0.00373)	(0.00372)
β^2	0.0188***	0.0226***	0.0313***
	(0.00191)	(0.00150)	(0.00257)
β^3	0.0146***	0.0190***	0.0155***
	(0.00204)	(0.00147)	(0.00201)
IV	Own	State	Padd
Firm PT	0.416	0.189	0.248
State PT	0.555	0.446	0.490
Fringe PT	0.688	0.606	0.711
US PT	0.993	1.004	1.036
N	58533	58533	58533
r2	0.940	0.943	0.941

Results from estimating equation (1) using 2-stage least squares. The dependent variable in each regression is the average wholesale price received by firm f in a given month. The data include all firm-month-states with more than 1% of wholesale sales between 2004 and 2015. All models include firm-state, year and month dummies, as well as controls for weather, income and population. Standard errors are clustered at the firm-state level and presented in parentheses.

6 Welfare and policy discussion

We now consider the welfare impacts of the asymmetric input cost shock caused by fracking and relate our analysis to a hypothetical tax on the carbon intensity of transportation fuels levied at the refinery level.

6.1 Welfare and fracking-induced input cost shocks

Firm-level pass-through depends on both the direct effect on the taxed firms and the indirect strategic effect on untaxed firms. Thus, incidence of a tax is specific to a tax τ. Weyl and Fabinger characterize the tax-specific effect on consumer and producer surplus as:

$$\frac{dCS}{dt_\tau} = -\rho_\tau Q, \tag{2}$$

$$\frac{dPS}{dt_\tau} = -Q\left[1 - (1-\theta)\rho_\tau + cov(\lambda_i^\tau \rho_{\tau_i}, \theta_i)\right] \tag{3}$$

where ρ_τ and θ reflect quantity-weighted average pass-through and conduct parameters, and $\lambda_i^\tau, \rho_{\tau_i}$ and θ_i correspond to the economic incidence of the tax on firm i, firm-specific pass-through of firm i and the conduct parameter of firm i, respectively. The marginal effect on consumers surplus is equal to the weighted-average pass-through rate multiplied by aggregate quantity. With symmetric firms, the marginal effect on producer surplus is a weighted average of the pass-through rates under perfect competition ($-Q[1-\rho]$) and monopoly ($-Q$), weighting by former by $1-\theta$ and the latter by θ. If firms differ with respect to θ_i, the covariance term in non-zero - the marginal effect on producer surplus increases if the tax is borne by firms with relative high markups (i.e., conduct parameters close to 1).

6.2 Implications for a carbon tax on transportation fuels

Approximately 20 percent of the lifecycle emissions from gasoline occurs prior to the pump (IHS 2012). A long literature documents that fuel excise taxes levied on a per-gallon basis and world oil price shocks are fully transmitted to retail prices (e.g., Marion and Muehlegger (2011), Borenstein et al. (1997)). Conventional wisdom often assumes that a similar, full and uniform, pass-through would result from a carbon tax levied on transportation fuels. Conversely, using input cost rather than tax variation, Ganapati et al. (2017) find that refineries would only pass-on 24 - 33% of a prospective carbon tax. Using the framework presented above, we are able do reconcile these two findings as pertaining to an industry average vs firm-specific carbon tax. Using detailed emissions data, we are also able to show that these average predictions mask considerable variation in incidence across firms, due to both heterogeneity in emissions intensity and competition.

In 2010, the EPA began reporting facility-level annual greenhouse gas emissions for the largest emitting industries. Oil refineries are the second highest ranking sector in terms of GHG per facility, behind the electric power sector. With average emissions of 1.22 MMT CO2e, these 145 facilities make up approximately 3% of total annual GHG emissions in the United States. Although the carbon released from combustion is identical regardless of the refinery at which the fuel was produced, the carbon intensity of the refining process actually varies substantially. Figure (8) documents this heterogeneity in refining emissions across U.S. refineries. As discussed in Section 3, some refineries subject crude inputs to considerably more processing than others. Table (7) documents that this additional processing implies a substantial increase in refining emissions per barrel.

Figure 8: GHG heterogeneity

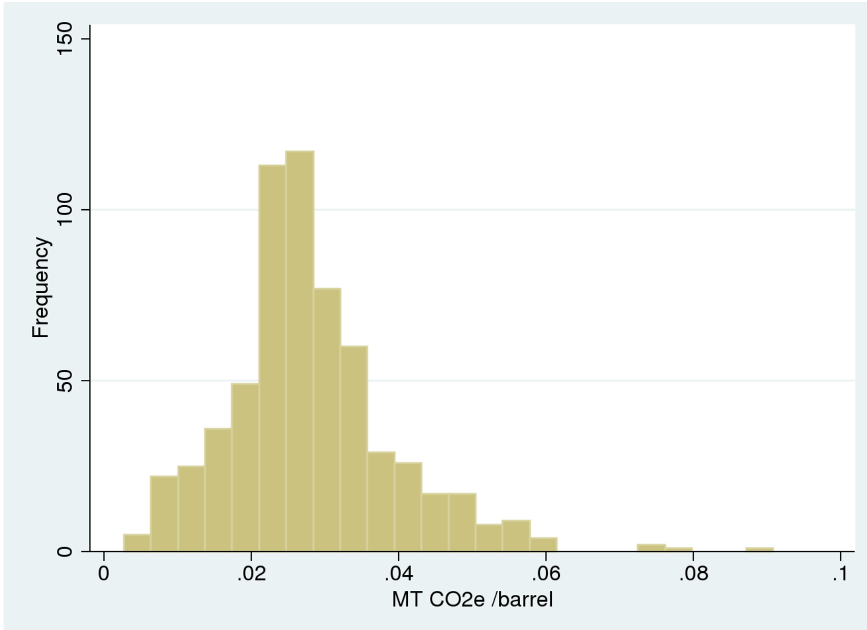

Based on annual data (2011-2015) in EPA GHGRP

Table 7: Determinants of CO2 heterogeneity

API Gravity	0.000193***
	(0.0000544)
log(Capacity)	-0.000197
	(0.000441)
% Coking	0.0240***
	(0.00345)
% Cracking	0.0207***
	(0.00189)
Constant	0.0229***
	(0.00528)
mean(Y)	.027
N	614
r2	0.370

Dependent variable: Metric tons of CO_2 equivalent per barrel of inputs processed. Data from EPA GHGRP (2011-2015). Model includes PADD and year dummies.

This heterogeneity across competing refineries have important implications for predictions about the incidence of a carbon tax. Figure (9) presents the distribution of markup changes (price change less tax change) under a $40 carbon tax, based on the parameter estimates in column (3) of Table 4. The red line plots the distribution when only when only firm-specific pass-through is considered.

All refiners see marginal profits decline. The blue line presents the distribution using the full-pass through estimates across all four nests. 42% of refineries actually see markups increase under the tax, with a difference in markup changes of 1.7 cents per gallon between the top and bottom deciles.

Figure 9: Full vs firm-specific carbon tax incidence

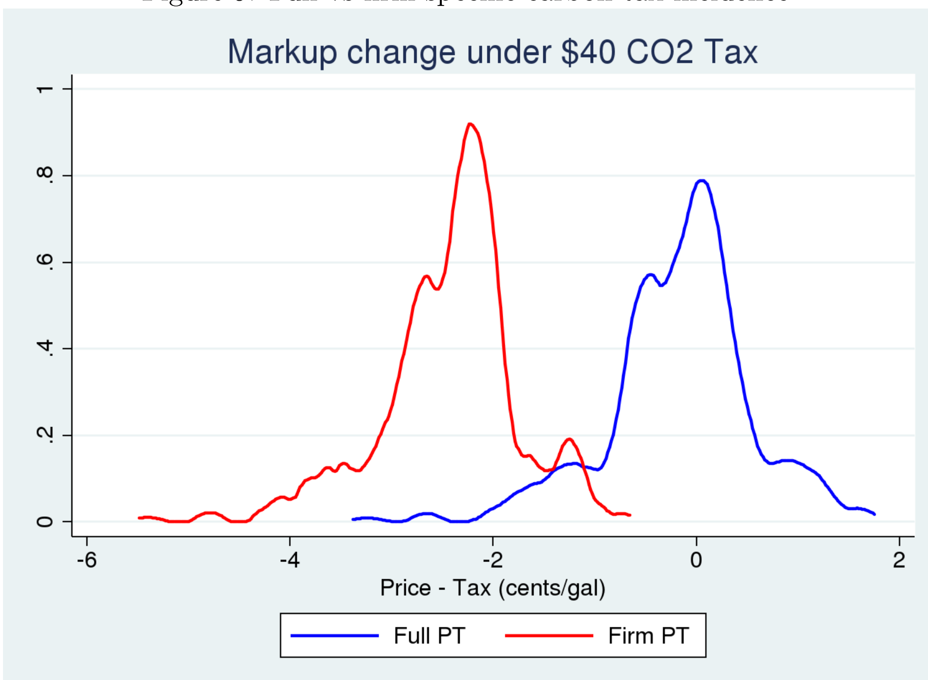

Table 8: Carbon tax results by PADD

	1	2	3	4	5
Firm PT	0.063	0.058	0.027	0.078	0.068
State PT	0.32	0.34	0.41	0.29	0.31
Fringe PT	0.47	0.56	0.66	0.54	0.46
US PT	1.11	1.11	1.11	1.11	1.11
Tax (cpg)	2.70	2.45	2.36	2.56	2.88
Markup change - Firm (cpg)	-2.53	-2.32	-2.30	-2.36	-2.69
Markup change - Full (cpg)	-0.29	-0.13	0.038	-0.17	-0.35

The results from model 3 in Table 4 are used to predict the impact of a $40/ton carbon tax. Results presented are average predictions by firm-month from 2011-2015.

7 Conclusion

The advent of fracking has lead to a remarkable increase in domestic crude oil production. At the same time, the presence of regulatory and infrastructure constraints, combined with private constraints and incentives among refiners suggests that the social benefits of this boom may not have been fully realized. We provide evidence that crude oil prices have diverged both internationally and domestically as a result of these constraints. We then use this variation to estimate pass-through in this industry of perennial economic and political import. We find that while refineries have little ability to pass on idiosyncratic cost shocks, shared cost changes have increasingly larger impacts, culminating in slightly greater than full pass-through for an industry-wide shock.

These findings have important implications. While we find that own cost pass through significantly exceeds pass-through of nearby rivals' costs, countering conventional wisdom that refined product markets are well approximated by Cournot competition, when combined rivals' costs have a large effect on realized prices. We demonstrate the important economic and econometric implications of this by evaluating the incidence of carbon tax on the industry. While a idiosyncratic, firm-specific levy would fall largely entirely on the refiner, the more relevant industry-wide measure would be largely pass-through to consumers. Nevertheless, heterogeneity in emissions intensity interacts with market structure to create winners and losers within the policy.

References

Bachmeier, Lance J and James M Griffin, "New evidence on asymmetric gasoline price responses," *Review of Economics and Statistics*, 2003, *85* (3), 772–776.

Bacon, Robert W, "Rockets and feathers: the asymmetric speed of adjustment of UK retail gasoline prices to cost changes," *Energy economics*, 1991, *13* (3), 211–218.

Borenstein, Severin, A Colin Cameron, and Richard Gilbert, "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?," *The Quarterly Journal of Economics*, 1997.

_ **and Ryan Kellogg**, "The incidence of an oil glut: Who benefits from cheap crude oil in the midwest?," *The Energy Journal*, January 2014, *35* (1).

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker, "The Incidence of Carbon Taxes in U.S. Manufacturing: Lessons from Energy Cost Pass-Through," April 2017.

Genesove, David and Wallace P. Mullin, "Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890-1914," *The RAND Journal of Economics*, 1998, *29* (2), 355–377.

Marion, Justin and Erich Muehlegger, "Fuel tax incidence and supply conditions," *Journal of Public Economics*, 2011, *95* (9), 1202–1212.

Miller, Nathan H., Matthew Osborne, and Gloria Sheu, "Pass-Through in a Concentrated Industry: Empirical Evidence and Regulatory Implications," *RAND Journal of Economics*, 2016.

Weyl, E. Glen and Michal Fabinger, "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, June 2013, *121* (3), 528–583.

Appendix A: Additional figures

Figure A.1: Refinery locations and PADD Map

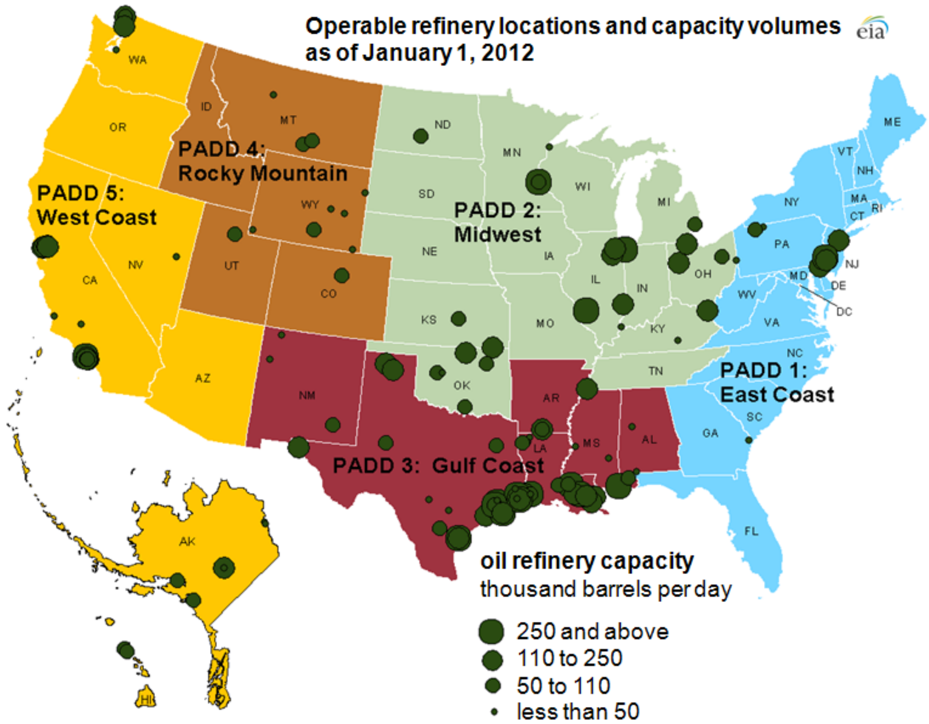

Appendix B: Data appendix

Table B.1: Description of EIA Data

Survey	**Dates**	**Description**
Monthly Refinery Report (EIA-810)	1986-2015	Collects information regarding the balance between the supply (beginning stocks, receipts, and production) and disposition (inputs, shipments, fuel use and losses, and ending stocks) of crude oil and refined products located at refineries.
Annual Refinery Report (EIA-820)	1986-1995 1997 1999-2015	Collects data on: fuel, electricity, and steam purchased for consumption at the refinery; refinery receipts of crude oil by method of transportation; current and projected capacities for atmospheric crude oil distillation, downstream charge, and production capacities.
Refiners' Monthly Cost Report (EIA-14)	2002-2015	Collects data on the weighted cost of crude oil at the regional Petroleum for Administration Defense District (PADD) level at which the crude oil is booked into a refinery.
Refiners'/Gas Plant Operators' Monthly Petroleum Product Sales Report (EIA-782A)	1986-2015	Price and volume data at the State level for 14 petroleum products for various retail and wholesale marketing categories are reported by the universe of refiners and gas plant operators
Monthly Report of Prime Supplier Sales of Petroleum Products Sold for Local Consumption (EIA-782C)	1986-1990 1992-2015	Prime supplier sales of selected petroleum products into the local markets of ultimate consumption are reported by refiners, gas plant operators, importers, petroleum product resellers, and petroleum product retailers that produce, import, or transport product across State boundaries and local marketing areas and sell the product to local distributors, local retailers, or end users.

Notes: Additional information as well as the survey forms for each dataset available at http://www.eia.gov/survey/.