

NBER WORKING PAPER SERIES

PASS-THROUGH OF INPUT COST SHOCKS UNDER IMPERFECT COMPETITION:
EVIDENCE FROM THE U.S. FRACKING BOOM

Erich Muehlegger
Richard L. Sweeney

Working Paper 24025
<http://www.nber.org/papers/w24025>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2017

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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Erich Muehlegger and Richard L. Sweeney. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Pass-Through of Input Cost Shocks Under Imperfect Competition: Evidence from the U.S.
Fracking Boom

Erich Muehlegger and Richard L. Sweeney

NBER Working Paper No. 24025

November 2017

JEL No. H22,H23,Q40,Q54

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.

Erich Muehlegger
University of California - Davis
Economics Department
One Shields Ave
Davis, CA 95616
and NBER
emuehlegger@ucdavis.edu

Richard L. Sweeney
Economics Department
Boston College
Chestnut Hill, MA 02467
sweeneri@bc.edu

Pass-through of input cost shocks under imperfect competition: Evidence from the U.S. fracking boom*

Erich Muehlegger [†]
Richard L. Sweeney [‡]

November 2017

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

The extent to which input cost or tax changes are passed through to market prices has wide ranging economic and policy implications. It is used to estimate the incidence of actual taxes (e.g., [Marion and Muehlegger \(2011\)](#); [Fabra and Reguant \(2014\)](#); [Stolper \(2016\)](#), among many others) and to forecast the incidence of hypothetical taxes of particular interest (e.g., [Ganapati et al. \(2017\)](#); [Miller et al. \(2017\)](#)). In public finance, it is often used as a sufficient statistics for welfare analysis generally (e.g., [Chetty \(2009\)](#)). More broadly, estimates of pass-through can be used to recover other objects of policy interest, such as trade costs ([Atkin and Donaldson \(2015\)](#)) or demand elasticities ([Miller et al. \(2013\)](#)).

*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, sweeneyr@bc.edu.

Yet, [Weyl and Fabinger \(2013\)](#) notes that the use of pass-through as a tool for economic analysis is complicated by imperfect competition. Even simple oligopoly models illustrate that changes in a firm’s costs will be transmitted to realized prices both directly through it’s own decisions, but also indirectly through the strategic response by the firm’s competitors. The magnitude of that strategic response will depend importantly on the nature of competition and the strategic closeness of the firms in question. Furthermore, indirect shocks affecting a subset of firms may result in very different observed price responses than shocks affecting an entire industry. Thus, estimates of pass-through relying on firm-specific cost shocks for identification may prove poor guides for forecasting the impact of common, industry-wide changes in taxes or regulation.

To date, three empirical challenges have further hindered the estimation of pass-through in imperfectly competitive settings. First, to directly observe pass-through, researchers need access to detailed data on both prices and costs for all firms active in a market. Firm-level prices and particular input costs are often difficult to observe directly, with [Fabra and Reguant \(2014\)](#) as a notable exception. Second, it is often difficult to directly observe which pairs of firms are close competitors and which are not. Third, to make any causal claims based on the co-movement between these prices and costs, the econometrician needs exogenous variation in input shocks both within markets and across different market structures. In this paper, we address these concerns using a novel data set, and provide new evidence on the relationship between pass-through and competition from one of the largest manufacturing industries in the United States, petroleum refining.

The emergence of hydraulic fracturing (or “fracking”) precipitated a near 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 an oil refinery. While oil markets are generally considered to be well integrated, a series of constraints limited the set of refineries that were able to benefit from this production 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 within the U.S. Third, the highly tailored nature of refinery operations, resulting from capital investments made long before the advent of fracking, meant that not every refinery could process these new crude streams. The net result of these three factors was a dramatic reduction in the primary input cost of some refiners, often tens of dollars below prevailing world market prices, while the costs for other firms, often located in similar areas and producing identical products, remained largely unchanged.

To examine this exogenous shock to the input costs of a subset of refineries, we employ detailed micro data on the universe of oil refiners in the United States from the the Energy Information

¹What 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.

Administration (EIA). For every firm that owns a refinery in United States, we observe regional monthly crude procurement costs. 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. Thus, the data allow us to circumvent two challenges often confronted when estimating pass-through: we directly observe firm-level prices and costs, and because we observe the states into which firms sell product, we can distinguish between cost shocks affecting close rivals and cost shocks that affecting firms that do not directly compete.

Using this detailed data, we document the importance of accounting for rival cost shocks when estimating pass-through. When fine time-space dummies are included, we find that refiners are essentially unable to pass-through idiosyncratic input cost shocks. However, when these controls are relaxed, we find that 20% of market specific costs, 35% of national cost shocks, and 95% of global cost shocks are passed through. These results reconcile estimates from the existing literature. For common, industry-wide cost shocks, we estimate pass-through close to unity, in line with the prior literature on the transmission of state-wide fuel tax shocks (e.g., [Doyle Jr. and Samphantharak \(2008\)](#), [Marion and Muehlegger \(2011\)](#), and [Stolper \(2016\)](#)) and estimates of pass-through of renewable fuel credits (e.g., [Knittel et al. \(2015\)](#) and [Lade and Bushnell \(2016\)](#)) . For shocks affecting single firms, for which we expect competition to limit pricing power, we find much lower rates of pass-through, at the lower end of recent estimates from [Ganapati et al. \(2017\)](#).

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 document 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 and 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 differs substantially, with 15% of firms actually experiencing price increases beyond their change in costs.

This paper contributes to 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. [Doyle Jr. and Samphantharak \(2008\)](#) also find evidence of high rates of pass-through of sales tax moratoria in midwestern states as does [Stolper \(2016\)](#), who examines pass-through of regional fuel taxes in Spain. [Knittel et al. \(2015\)](#) and [Lade and Bushnell \(2016\)](#) find that the implicit

tax created by renewable fuel credits are fully passed onto consumer prices. Yet, a separate literature (e.g., [Ganapati et al. \(2017\)](#)) find evidence that regional input costs shocks are incompletely passed onto consumers.² [Miller et al. \(2017\)](#) study pass-through in another concentrated industry, cement, finding that costs are more than fully passed onto consumers.

Relative to the previous literature, - our data allow us to directly observe both the input and output prices of firms as well as the markets into which the firms sell. Furthermore, the cost variation exploited is particularly rich - in some areas, the shock associated with fracking affects virtually all firm, where in other locations, very few firms are impacted. In contrast, much of the existing literature exploits tax shocks that, by their nature, affect all the firms in a market.³

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 relate to price changes in theory. We then provide background detail on the U.S. fracking boom and how it affected crude oil input costs at refineries (Section 3). Section 4 discusses the data and empirical strategy, and Section 5 presents the results. In Section 6 we discuss the policy implications with an application to a prospective carbon tax before concluding in Section 7.

2 Pass-through and imperfect competition

A long-standing theoretical literature on the relationship between pass-through and competition dating back to [Bulow and Pfleiderer \(1983\)](#) has recently become the subject of renewed interest. Before we move to our empirical application, we expand on this work to illustrate how pass-through relates to the nature of the cost-shock and degree of competition in imperfectly competitive markets, using several standard models of competition to motivate and guide our empirical analysis.

[Weyl and Fabinger \(2013\)](#) (henceforth, WF) formalize a general model capturing the incidence of a tax on asymmetric, imperfect competitors. As a starting point, we first restate their main set of results and then adapt them to the case in which we are interested, imperfectly competitive firms receiving either firm-specific or common shocks to input costs. Following their exposition, let the

²In addition to the work on fuels, another equally developed literature considers the pass-through of sales and excise taxes on a range of consumer good, from cigarettes (e.g., [Harding et al. \(2012\)](#); [Chiou and Muehlegger \(2008\)](#); [Rozema \(2017\)](#)) to Soda Taxes (e.g., [Berardi et al. \(2016\)](#)).

³A 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.

vector σ denote a single dimensional strategic 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

$$\theta_i = \frac{m \frac{dq}{d\sigma_i}}{-q \frac{dp}{d\sigma_i}} \quad (1)$$

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.⁴

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). Following the notation in WF, we use ρ_α to denote the pass-through of a shock α onto the vector of firm-specific prices.

$$\rho_\alpha = \sum_i \frac{\partial p}{\partial \sigma_i} \left(\frac{\partial \sigma_i}{\partial \alpha_i} + \sum_{j \neq i} \frac{\partial \sigma_i}{\partial \sigma_j} \frac{d\sigma_j}{d\alpha} \right)$$

As this equation illustrates, an input cost shock affects a firm's choice of strategy both directly through the firm's own costs and indirectly, in response to the strategic choice of competitors. Of primary interest in our empirical exercise is how input cost pass-through relates to the degree the cost shock is idiosyncratic to the firm. Specifically, we contrast the pass-through of a firm-specific shock, ρ_{α_i} , and the pass-through of a market-wide shock, $\rho_{\bar{\alpha}}$, using three different models of competition as examples. Intuitively, 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 . In contrast, a common shock, $\bar{\alpha}$, affects both i and j 's strategy directly.

2.1 Cournot competition

In the Cournot model, the strategic variable σ_i is simply the quantity produced by a particular firm. Under constant marginal costs, the market price is determined by the sum of the marginal costs of the market participants. Following the n -firm asymmetric case examined in [Fevrier and Linnemer \(2004\)](#), we sum across the n first-order conditions and take the derivative with respect to α_i and $\bar{\alpha}$ respectively, to obtain:

$$\frac{dQ}{d\alpha_i} = \frac{1}{(n+1)P'(Q) + QP''(Q)}; \quad \frac{dQ}{d\bar{\alpha}} = \frac{n}{(n+1)P'(Q) + QP''(Q)}. \quad (2)$$

The pair of equations have several important implications for input cost pass-through in a Cournot model. First, a firm-specific shock is passed along at a rate of $1/n$ relative to commensurate market-

⁴As 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 its 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, $\theta_i = 1$.

wide shock. As the number of competitors increases, a change in the affected firm's production is offset by an increase in production by an increasing number of firms, and the pass-through of a firm-specific shock declines. In contrast, a common shock causes all firms to lower production, and pass-through to increase with the number of competitors - asymptotically approaching full pass-through in the case of linear demand. Second, a firm-specific shock to any single market participant has a similar effect on the market price, regardless of firm's initial market share and marginal cost. Finally, in the case of monopoly, both expressions are identical and reduce to the standard expression for the pass-through of a cost-shock under monopoly, $\frac{dP}{dc} = \frac{P'(Q)}{2P'(Q)+QP''(Q)}$, highlighted by [Bulow and Pfleiderer \(1983\)](#).

2.2 Differentiated Nash-in-Prices

Next consider the differentiated Nash-in-price model. In the context of refining, differentiation could reflect the geographic nature of delivery and competition. In this model, the strategic variable σ_i represents the price set by a firm. As in the previous case, Nash competition implies that the only non-zero term in $\frac{dp}{d\sigma_i}$ is the one corresponding to a firm's own price, simplifying the expression for ρ_α . For expositional simplicity, we focus on the two firm case, for which we can express the pass-through of a cost shock α onto firm i's price as:

$$\rho_\alpha = \frac{\frac{\partial \sigma_i}{\partial \alpha} + \frac{\partial \sigma_i}{\partial \sigma_j} \frac{\partial \sigma_j}{\partial \alpha}}{1 - \frac{\partial \sigma_i}{\partial \sigma_j} \frac{\partial \sigma_j}{\partial \sigma_i}}. \quad (3)$$

The pass-through of a cost shock depends on three terms - the direct effect of the cost shock on firm i's strategy, the indirect response to a competitor's shock and the strength of strategic complementarity. We can express the ratio of pass-through for a firm-specific shock to a common shock as:

$$\frac{\rho_{\alpha_i}}{\rho_{\bar{\alpha}}} = \frac{\frac{\partial \sigma_i}{\partial \alpha_i}}{\frac{\partial \sigma_i}{\partial \alpha_i} + \frac{\partial \sigma_i}{\partial \sigma_j} \frac{\partial \sigma_j}{\partial \alpha_j}}. \quad (4)$$

The pass-through predictions under this model differ from those under Cournot. Although in both cases, the pass-through of a firm-specific shock is lower than that for a common shock affecting both firms, under differentiated Nash, the degree to which the two types of shocks differ in magnitude depends on the degree of competition between the two firms. As the products become closer substitutes (reflected as an increase in $\frac{\partial \sigma_i}{\partial \sigma_j}$), a competitor's cost shock exerts an increasingly large impact on the firm i's optimal price and the pass-through rates of the two different shocks diverge. Second, in the Cournot model, the pass-through of a cost shock was independent of the identity of the affected party. This is not the case with differentiated products. The pass-through of a competitor's cost onto firm i's price depends on the degree of strategic complementarity. Furthermore, a cost shock to the firm itself will, under typical circumstances in which $\frac{\partial \sigma_i}{\partial \alpha_i} > \frac{\partial \sigma_i}{\partial \sigma_j} \frac{\partial \sigma_j}{\partial \alpha_j}$, be passed on more fully than a commensurate cost shock to the firm's competitor.

2.3 General oligopoly with price competition

As a final benchmark for consideration, we extend the differentiated products model above to the case in which firms choose prices, but do not necessarily play Nash strategies. Following the setup in [Weyl and Fabinger \(2013\)](#) and [Genesove and Mullin \(1998\)](#), we denote θ_i as the firm i 's conduct parameter, such that firm i 's first-order condition satisfies:

$$- \mathbf{q} \frac{d\mathbf{p}}{d\sigma_i} \theta_i = \mathbf{m} \frac{d\mathbf{q}}{d\sigma_i}. \quad (5)$$

As firms choose prices as their strategic variable, the vector of the derivative of prices with respect to the strategy of i , $\frac{d\mathbf{p}}{d\sigma_i}$, is a vector of zeroes with the exception of a one in the i th position. Simplifying the left-hand side of (5) and taking the derivative with respect to α_i , we can represent the pass-through of a firm-specific shock onto its own price:

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

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 more 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 $\sum_j \frac{dq_j}{dp_i} < 0$.

Denoting ρ_α^i to be the pass-through of shock α onto the price charged by firm i , the pass-through of firm-specific and market-wide shocks on a firm's own prices are:

$$\rho_{\alpha_i}^i = \frac{d\sigma_i}{d\alpha_i}; \quad \rho_\alpha^i = \frac{d\sigma_i}{d\alpha_i} + \sum_{i \neq k} \frac{\partial \sigma_i}{\partial \sigma_j} \frac{d\sigma_j}{d\alpha_j}$$

In each of the three cases above, pass-through depends on the nature of the cost shock. Firm-specific cost shocks have lower rates of pass-through than cost shocks that are common to all firms. Moreover, the patterns of pass-through distinguish the nature of competition. Under Cournot, cost shocks have identical effects, regardless of the affected party, whereas under models with differentiated products, the degree of competition between two parties plays a central role in determining the pass-through of own and competitor cost shocks.

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. Specifically, we highlight

that: (1) U.S. crude oil prices fell relative to the prevailing world price, (2) prices fell more in regions proximate to fracking, (3) even within these regions, input prices of refiners fell idiosyncratically, and (4) the degree to which input prices fell was correlated with exogenous factors, driven by capital decisions made by firms many years earlier.

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.⁵ 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, fracking 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.⁶ In contrast, production continued to decline in other traditional oil deposits.⁷ Crude oil transportation infrastructure has been slow to adjust to 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

⁵A 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.

⁶In 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.

⁷Production from the Alaska and federal offshore deposits fell approximately 20 percent over same period.

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



to the Brent spot) at refineries by Petroleum Administration Defense District (PADD).⁸ 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, the impact of the fracking boom on refinery crude acquisition costs varied across firms. 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. Consistent with figure 3, the average acquisition price in PADDs 2 and 4 fell during the early years of the fracking boom. But, the fracking boom also increased the spread of oil acquisition prices for firms, even within particular PADDs. Fortunately, a significant fraction of the within PADD variation relates to the capital decisions made by refineries substantially before the fracking boom, and thus provides an potential instrument for cost-shocks.

Crude oil is a highly differentiated input, primarily based on how dense it is and how much sulfur it contains. These characteristics define what refining equipment is necessary to convert crude into

⁸Petroleum 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




usable refined products. Denser products (those with low API gravity) contain smaller naturally occurring shares of valuable end products (like gasoline). Refineries must further process these crude oils to yield high shares of gasoline and diesel fuel. Similarly, “sour” (high-sulfur) crudes require additional processing to low sulfur levels, reducing the corrosiveness of refined products and reducing harmful health effects post-combustion. Under typical market conditions, denser, higher-sulfur crudes sell at a discount to crude oils that are more easily processed into valuable end products.

As a result of this input differentiation and decades of costly investment, refineries are highly tailored to process specific crudes. Substantial changes to the crude slate requires either months of reconfiguration or large 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. These 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, the share of crude coming from imports was surprisingly stable, with many firms continuing to use

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



imported oil. The one major exception was the East Coast, which shifted substantially over the period due to the expansion of crude by rail.

The net effect of these different capital decisions can be seen in a simple regression (Table 1). The dependent variable is the change in average crude price discount (firm price minus Brent spot) between a “post” period (2010 - 2013) during which domestic crude markets diverged considerably, and a “pre” period (2005-2008) during which fracking was still in its infancy. This average fracking boom cost shock is then projected onto “pre” period measures of which firms were well situated to benefit. As was already discussed, the deepest discounts occurred in the upper Midwest and Plains states (PADD’s 2 and 4). However, even within region, larger discounts went to firms tailored to use lighter crude and domestic crude in the period before the fracking boom. In contrast, crude oil acquisition costs fell much less for firms that specialized in heavier and international crude oil streams in the pre-period.

Figure 4: Average Refinery Crude Price by PADD

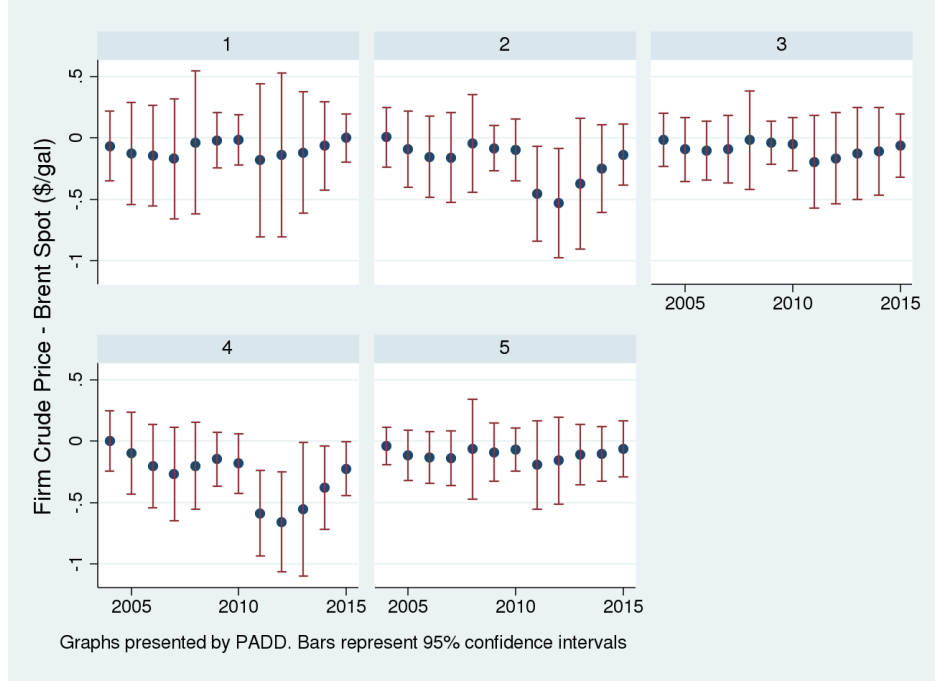



Table 1: Within-PADD Fracking Boom Beneficiaries

padd=2	-0.108* (0.0583)
padd=3	0.0737 (0.0559)
padd=4	-0.168*** (0.0585)
padd=5	0.115 (0.0713)
API Gravity	-0.00593* (0.00334)
% Domestic Crude	-0.121** (0.0593)
Log(Capacity)	0.0108 (0.0164)
Downstream Capacity	-0.0650 (0.0543)
mean(Y)	-.13
N	50
r2	0.716

The dependent variable is the change in crude price discount (relative to Brent) between the pre period (2005-2008) and the post period (2010-2013). The explanatory variables are the pre period firm characteristics. The sample is restricted to firms with at least 24 months of observations in both periods.

Figure 5: Domestic Crude Price Differentials by Density



4 Data and Empirical Strategy

4.1 Data


Through a confidential data request we obtained data on refinery operations from the Energy Information Administration (EIA), each described in detail 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).⁹

Our data provide two advantages that lend it to the study of competition and pass-through. First, we observe firm-level crude oil input costs (which comprise 90% of refiners input costs) and

⁹A 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 6: Domestic Crude Acquisition Shares by PADD



output prices, allowing us to directly observe margins for each firm. Second, the data provides us with relatively direct evidence as to the degree to which firms compete. We observe the location into which each firm sold refined products. This allows us to distinguish firms that are close rivals, serving similar markets, from more distant competitors, serving different markets. These features are particularly valuable in the context of the refining industry, a major industry for which previous research has been stymied by a lack of publicly-available, disaggregated data.

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

Table 2: Summary statistics

	mean	sd
Crude cost (2013 \$/gal)	1.858	0.572
Crude - Brent	-0.158	0.220
% Domestic	0.562	0.368
Price Gas	2.376	0.584
Price Diesel	2.493	0.676
Price Total	2.390	0.613
Resale Price Total	2.378	0.611
% Gas	0.528	0.184
% Diesel	0.399	0.171
% Resale	0.862	0.142
N	9011	

4.2 Empirical Strategy

Our primary goal in this paper is to examine the relationship between cost changes and prices in the refining industry. But, the richness of our data and the nature of the cost shock under study also allow us to discuss and empirically document several potential pitfalls in estimating pass-through in imperfectly competitive markets generally. Consider the canonical pass-through regression that projects the price a firm f receives in a given market m at time t onto its own costs,

$$Price_{fmt} = \alpha Cost_{ft} + X'_{fmt}\delta + \epsilon_{fmt} \quad (6)$$

where X contains other demand and supply side factors which may shift the level of prices.¹⁰ The theory model in Section 2 suggests several challenges when taking this regression to data in imperfectly competitive settings.

First, in oligopolistic markets, realized prices will be a function of both own costs and competitors' costs. Thus, as written above, equation (6) suffers from omitted variable bias, the magnitude of which is the product of the covariance of a firm's own cost shocks with those of its competitors and the degree to which the competitors cost shocks are passed onto a firm's own prices. As a simple example, in the asymmetric Cournot model with N firms in which a cost shock to any competitor is equally passed-onto prices, we can precisely quantify the degree to which α will be biased upwards. Denoting the average cost of a firm f 's rivals as \bar{Cost}_{ft} , omitting rivals costs for the equation above implies the estimate of the pass-through of a firm's own costs will be biased upwards by a factor of $1 + (N - 1)cov(Cost_{ft}, \bar{Cost}_{ft})$. In a more general case, we would expect the first term to depend on the nature of the cost shocks being used to estimate pass-through and the latter to depend on the degree of competition between a firm and its rivals.

This suggests an augmented specification that allows rivals' costs to enter into equation (6)

¹⁰In theory $Cost$ should include all marginal costs of supplying market m at time t . In practice, we only observe average crude costs, which make up over 90% of variable refinery costs.

directly. Denoting $f(RivalCost_{-f,mt})$ as an average a f 's competitors' costs weighted to reflect the nature of competition in m at time t , we propose to estimate:

$$Price_{fmt} = \alpha Cost_{ft} + \beta f(RivalCost_{-f,mt}) + X'_{fmt}\delta + \epsilon_{fmt} \quad (7)$$

Continuing the example of the Cournot case above, in which all firms' costs have identical effects on pass-through, $f(RivalCost_{-f,mt})$ would consist of a simple average and $E[\hat{\beta}] = (N - 1)E[\hat{\alpha}]$.

In many markets, it is less clear if two firms are direct competitors. For example, in the refining industry transportation constraints may limit the degree of competition, while other, seemingly more distant refineries may be in close competition. Even if the identity of competitors is known, the exact relationship between own price and competitor costs depends on the nature of competition in the market, functional form assumptions on the log convexity of supply and demand (as well as implied timing assumptions). This suggests a second pitfall with estimating pass-through in imperfectly competitive markets: Weighting rivals' costs inappropriately introduces measurement error that would tend to attenuate estimates of the pass-through of competitors' costs. As with omitting rivals' costs completely, mis-weighting rivals costs will bias the estimate of own-cost pass-through. Viewed from this perspective, we can frame omitting rivals' costs entirely as simply assigning a weight of zero to the costs of all rivals. In our case, we exploit the richness of the EIA data, as we are able to observe the frequency with which two rival firms serve the same market (state). In lieu of making assumptions about these structural primitives, we instead estimate the a reduced form pass-through equation whereby we group close rivals (that serve the same location) and rivals that do not.

This observation suggests a third potential pitfall when estimating pass-through in imperfectly competitive markets - the inclusion of fixed effects changes the source of cost-shock variation and meaningfully affects the interpretation of the estimated coefficients on pass-through. A natural approach to the omitted variable problems in this setting might be to consider the inclusion of increasingly fine fixed effects. For example, if unobserved demand shocks in market m caused crude suppliers to increase the price of firms supplying this market, both α and β would be biased.¹¹ Or alternatively, if it is difficult to assess the appropriate weighting of distant competitors, we might consider fixed effects that flexibly subsume the cost-shocks of distant competitors. But, the inclusion of highly disaggregated fixed effects introduces a tradeoff between reducing the potential bias in own cost estimates and using the full suite of price variation in the data. For example, if year-month fixed effects are included, the econometrician cannot recover the pass through rate of an industry wide average shock. As we will illustrate in our empirical results below, this is directly relevant for policy analysis if one goal of the pass-through calculation is to extrapolate to the welfare or distributional consequences of an industry-wide tax or regulation.

¹¹Clearly, fixed effects are insufficient to adequately address cost endogeneity concerns. As described below, we also instrument for own- and competitors' costs.

5 Results

Before estimating equation (7), we must first decide what the appropriate regressand is. All refiners are multi-product firms, and some sell into both wholesale and retail markets.¹² 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 the fact that less than 50% of the barrel could be converted to gasoline. 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 3 explores the impact of this assumption on own-cost pass-through by displaying the estimated α from equation (6). In column (1) the dependent variable is the total revenue per gallon from all products and sales channels. The coefficient on own-cost indicates that a \$1 increase in a firm's crude costs leads to an average increase in price of \$0.047. Column (2) repeats this regression replacing the dependent variable with the average wholesale price (i.e. excluding retail sales). The estimated own cost pass-through rate is slightly higher, although the difference is not statistically significant. Columns (2)-(4) estimate the pass-through into the wholesale price of gasoline, distillate and other products, restricting the sample to months where firms had positive wholesale sales of each product. Observed product-specific pass-through rates are lower for gasoline and distillate, and significantly higher for other products. Given the small share of these products, the relatively high rate of total pass-through compared to light products likely reflects both composition and substitution effects. Motivated by this table, the remainder of this paper uses wholesale revenue per gallon as dependent variable of interest.

Table 3: Total vs product-specific pass-through

	(1)	(2)	(3)	(4)	(5)
Own	0.0474*** (0.0109)	0.0554*** (0.0131)	0.0174*** (0.00647)	0.0314*** (0.00579)	0.374*** (0.110)
Price	Total	Total	Gas	Dist	Other
Resale only	N	Y	Y	Y	Y
N	55568	55568	54372	54428	29376
r ²	0.977	0.975	0.981	0.988	0.750

Equation (6) is estimated using different average prices at the firm-state-month level as the dependent variable. Distillate sales include both jet fuel and diesel, and the other category contains all other products excluding gasoline and diesel. All models include firm-state and time dummies. Standard errors are clustered at the firm-state level and presented in parentheses.

Table 4 presents the results from several specifications of equation (7). In panel (a), the unit

¹²Table 2 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.

of observation is a firm-state-month. Model 1 is identical to that of column 2 in Table 3, with the inclusion of the average acquisition cost of other refineries in the same region. First focusing on the coefficient of a firm’s own-costs, we see that omitting nearby firms’ costs biases the coefficient on own-cost upwards, as we would expect from the discussion in section 4.2. Although in both specifications, our estimates of own-cost pass-through are relatively close to zero, the coefficient on own-costs is biased upwards by about 70 percent when we exclude other firms’ costs. This suggests that after conditioning on time fixed effects, much of the residual variation in own-costs is correlated with regional variation in input prices.

In column 2, we restrict the set of rivals to the firms that sell into state m each month. In essence, this re-weights the importance of potential competitors - reducing the weight of firms that are located in the same region, but not currently selling into a particular state. The PADD average in column 1 might include some refineries that are unlikely to supply a specific state, despite being located within the same PADD.¹³ We find that re-weighting the competitors costs in this way causes a modest increase in the coefficient on rival cost, suggesting that cost shocks to immediate competitors are more important than cost shocks to more distant competitors. The results in column 2 also provide suggestive evidence that Cournot assumptions might not be appropriate for the refining industry. Under Cournot with N competitors, we would expect that the coefficient on rivals’ costs would be $N - 1$ times the coefficient on own-costs. On average around ten “rivals” sell into each state. The coefficient on average rivals’ cost is substantially below what we would expect under Cournot competition.

¹³For example, the St. Paul Park refinery located in Minneapolis might be unable to directly supply Ohio, also located in PADD 2. Although the a cost shock at St. Paul might have an indirect effect on firms selling into Ohio, theory might lead us to expect much lower rates of pass-through.

Table 4: Pass-Through Regression Results

(a) State Level Results							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own	0.0343*** (0.0115)	0.0441*** (0.0121)	0.0433*** (0.0120)	0.0497*** (0.0127)	0.0420*** (0.0122)	-0.00599 (0.0266)	0.00576 (0.0269)
Padd Avg	0.0632*** (0.0186)						
Rival		0.0788*** (0.0182)	0.0758*** (0.0180)	0.131*** (0.0197)	0.0612*** (0.0159)	0.0928** (0.0416)	0.0687* (0.0352)
NonRival			0.168*** (0.0428)	0.811*** (0.0231)	0.245*** (0.0198)	0.000348 (0.106)	0.173*** (0.0334)
Brent Spot					0.611*** (0.00971)		0.700*** (0.00946)
Time FEs	Y-M	Y-M	Y-M	Y,M	Y,M	Y-M	Y,M
IV						Yes	Yes
first F						900	963
N	55568	55568	55568	55568	55568	55568	55568
r2	0.975	0.975	0.975	0.951	0.956	0.975	0.956

(b) PADD Level Results							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own	0.0400 (0.0245)	0.0355 (0.0237)	0.0353 (0.0238)	0.0277 (0.0276)	0.0292 (0.0256)	-0.0671 (0.0771)	-0.0646 (0.0766)
Padd Avg	0.0714* (0.0369)						
Rival		0.145*** (0.0466)	0.156*** (0.0465)	0.319*** (0.0538)	0.169*** (0.0377)	0.243 (0.148)	0.179 (0.112)
NonRival			0.0823 (0.0549)	0.635*** (0.0498)	0.150*** (0.0379)	0.0358 (0.126)	0.123* (0.0713)
Brent Spot					0.602*** (0.0158)		0.701*** (0.0196)
Time FEs	Y-M	Y-M	Y-M	Y,M	Y,M	Y-M	Y,M
IV						Yes	Yes
first F						108	142
N	8007	8007	8007	8007	8007	8007	8007
r2	0.984	0.984	0.984	0.961	0.966	0.983	0.965

This table presents the results of estimating Equation (7) using total average wholesale prices as the dependent variable. Panel (a) is estimated at the firm-state-month level, and includes firm-state fixed effects; Panel (b) is estimated at the firm level and includes firm fixed effects. Time FEs “Y-M” reflect year-month dummies, while “Y,M” implies year and month dummies. Rival costs include the average crude price of other firms selling into the same market each month, and non-rival costs are the average cost of all other firms, weighted by the inverse shipping cost of supplying the market. Standard errors are presented in parentheses, clustered at the firm-state level in panel (a) and the firm level in panel (b). All models include demand shifters (state population, income, heating and cooling degree days) and supply shifters (diesel and gasoline shares, proportion of retail sales, API gravity, and operating refinery capacity).

Extending this logic, model 3 also includes the average cost of firms on the competitive fringe of each market. This average is constructed by weighting the costs of all other firms not serving

state m by the inverse shipping cost each firm would face if it were to enter. Taken together, model 3 suggests that were all the firms directly and indirectly serving a given state to experience a cost increase of \$1 *conditional on national average costs*, prices would increase by almost 30% in that state.

Next we illustrate the impact of conditioning on national average costs on estimated pass-through. In model 3, time fixed effects subsume market-wide shocks arising, for example, from shocks to world crude prices. In model 4, we replace these time fixed effects with more coarse year and month fixed effects which capture broad time trends and seasonality, but now partially identify the coefficients in the model off of these more aggregate shocks. Comparing the coefficients in the two columns, we see that the coefficient on non-rival cost pass-through increases significantly, leading to the conclusion that a cost shock increasing the costs of all firms, directly and indirectly serving a given state would be fully passed onto consumers. The coefficient on own cost remains essentially unchanged - a firm's own-costs are uncorrelated with world crude price shocks after conditioning on the input costs of all of the firms that do not serve a particular state.

While the ordering of the direct and indirect rival coefficients makes sense given the fact that US petroleum product markets are relatively integrated and states are only served by a small fraction of refineries, it is possible that the industry wide average may be picking up within-year variation in global petroleum markets than far away domestic competitor costs. If we further condition on the Brent spot price (the benchmark crude price outside of the United States), as we do in Column 5, we find that the coefficients on own-, rival- and non-rival costs return to similar magnitudes as in the case we included time fixed effects. One interpretation of these results is that the coefficient on the Brent spot price reflects the potential impact of foreign competition on U.S. refiners ability to pass-along cost-shocks. The estimates imply that a cost shock affecting all firms would be passed fully along to consumers, whereas a cost shock affecting only U.S. firms would be passed along at a rate closer to 35%.

One concern with the using OLS to estimate equation (7) is endogeneity. Observed prices are a function of both supply and demand. While all models include controls for demand shifters, and 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. The two most significant determinants of crude prices during our sample are the density of the crude (API gravity) and whether it was domestic or imported. For each firm, first calculate the average density and share of domestic crude during the sample, excluding the current period. We then construct two instruments, using these leave-one-out averages and contemporaneous oil prices: a weighted average price using domestic and international spot prices, and the average acquisition cost of crude (at the point of first purchase) by density. The logic of these instruments is similar to Bartik (1991). Changes in the price of domestic oil arose due to the rapid maturation of hydraulic fracturing. The oil export ban and pipeline constraints temporarily drove a wedge between domestic and foreign crude prices, which could only be captured by firms with access

to domestic crude. Similarly, this new found oil was largely “light” (low density), temporarily reversing a long standing premium paid for such crudes. Despite this, many refineries were not able to switch to lighter crudes due to technology constraints.

The endogeneity we are concerned with is that unobserved demand factors in a given state within a particular period may be correlated with cost shocks to suppliers of those states. As a result, we instrument for both own and competitor costs, where the latter instruments are constructed by averaging the instruments of rival and non-rival firms accordingly. Column 6 presents the results with time fixed effects. The coefficient on own cost is now zero, while the estimated impact of rival costs has increased slightly. In column 7, when time fixed effects are replaced with year and month dummies, the results are similar. In sum, these results suggest that own costs are likely endogenous at the firm-state level, while average costs are not.

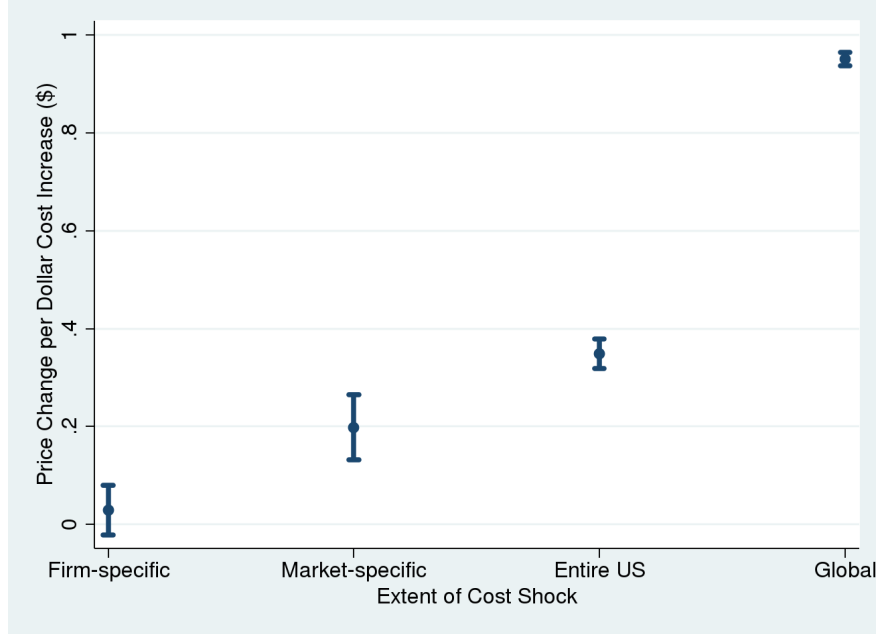
Panel (b) repeats the previous regressions after aggregating the data to the PADD level. In general, the own-cost coefficients are roughly similar to the state level regressions. Estimates of own cost pass through are again close to zero, relatively precisely estimated, and quite consistent across the different specifications, although once we aggregate to the PADD-level, the coefficients are no longer statistically significant. Two ways in which the results change deserve particular notice. First, with the exception of the model 4, the magnitudes of rival and non-rival average costs have been ordinally reversed. In the state-level regressions, the group of firms classified as non-rivals potentially included firms that sold into nearby states, or firms that could reallocate production to a state at low cost. As such, many of these non-rivals may have been substantially close competitors and exerted a greater influence on a firm’s production and pricing decisions. Aggregating to the PADD-level, non-rivals are now classified as firms that did not serve any of the same states as firm f in a given month, as such, are likely to consist of more distant competitors. Second, although the coefficients in the IV specifications are largely similar to those in the OLS specifications, the IV estimates are less precise than the estimates at the state-level. Intuitively, aggregating to the PADD-level removes much of the within-PADD variation in instrumented costs that was previously present when the regressions were run at the state-level.

6 Discussion

Empirical pass-through papers face an inherent omitted variables problem. The aim is to study how prices change in response to cost changes. However, observed costs could be correlated with un-modeled demand or supply changes, 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. This is directly relevant for welfare analysis. As [Weyl and Fabinger \(2013\)](#) notes, the welfare impacts and distri-

butional consequences of a tax change depend on the pass-through rate.¹⁴ Extrapolating from a firm-specific shock, estimated from a model with fine temporal fixed effects, may therefore substantially understate the consumer burden of a market-wide tax.

Figure 7: Implied Pass-Through by Shock Type



Estimates from model (5) of panel (b) in Table 4.

6.1 Implications for a carbon tax on transportation fuels

We now consider the implication of these points for a hypothetical tax on the carbon intensity of transportation fuels levied at the refinery level. In particular, we consider the case of a national carbon tax coupled with a commensurate border tax to tax the carbon content of imported fuels.¹⁵ 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., Borenstein et al. (1997); Marion and Muehlegger (2011)). 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 to reconcile these two findings as pertaining to an industry average vs firm-specific carbon tax. Using detailed


¹⁴In the setting of a general oligopoly, Weyl and Fabinger (2013) characterize the tax-specific effect on consumer and producer surplus as: $\frac{dCS}{dt_\tau} = -\rho_\tau Q$; $\frac{dPS}{dt_\tau} = -Q[1 - (1 - \theta)\rho_\tau + cov(\lambda_i^\tau \rho_{\tau_i}, \theta_i)]$, where ρ_τ and θ reflect quantity-weighted average pass-through for a particular shock and conduct parameters, and λ_i^τ , ρ_{τ_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.

¹⁵From the perspective of U.S. firms and consumers, this would be equivalent to a world-wide tax on the refining industry.

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 the eventual combustion of petroleum is identical regardless of the refinery at which the fuel was processed, 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 (5) 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 5: 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 CO2 equivalent per barrel of inputs processed. Data from EPA GHGRP (2011-2015). Model includes PADD and year dummies.

This heterogeneity across competing refineries has 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 (5) of panel (b) in 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 an industry wide carbon tax.¹⁶ Interestingly, 15% of refineries experience an increase in markups under the tax. The vertical blue line is positioned at the implied average full pass-through rate of 95%. The dispersion in the blue density plot around this average highlights the distinction between estimating full pass through and imposing in uniformly on each firm's cost increase, and instead predicting full pass through based on the full set of industry cost increases.

¹⁶We interpret the coefficient on the Brent spot price in this model as capturing the average cost of far away (foreign) rivals. As such, in the policy considered, we need to adjust importer costs by the tax. As we do not observe emissions for foreign refineries, we apply the average in the US GHGRP data.

Figure 9: Full vs Firm-specific Carbon Tax Incidence



7 Conclusion

The advent of fracking has led 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 for both the estimation of pass-through and for policy analysis in the refining industry. While small, 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.¹⁷

We further demonstrate the important economic and econometric implications of this by evaluating the incidence of carbon tax on the refining industry. While an idiosyncratic, firm-specific levy would fall largely entirely on the refiner, the more relevant industry-wide measure (reflecting perhaps a U.S. carbon tax coupled with a border tax on imported fuels) would be largely pass-through to consumers. Nevertheless, heterogeneity in emissions intensity interacts with market structure to create winners and losers within the policy.

¹⁷However, Cournot cannot be ruled out in our IV results.

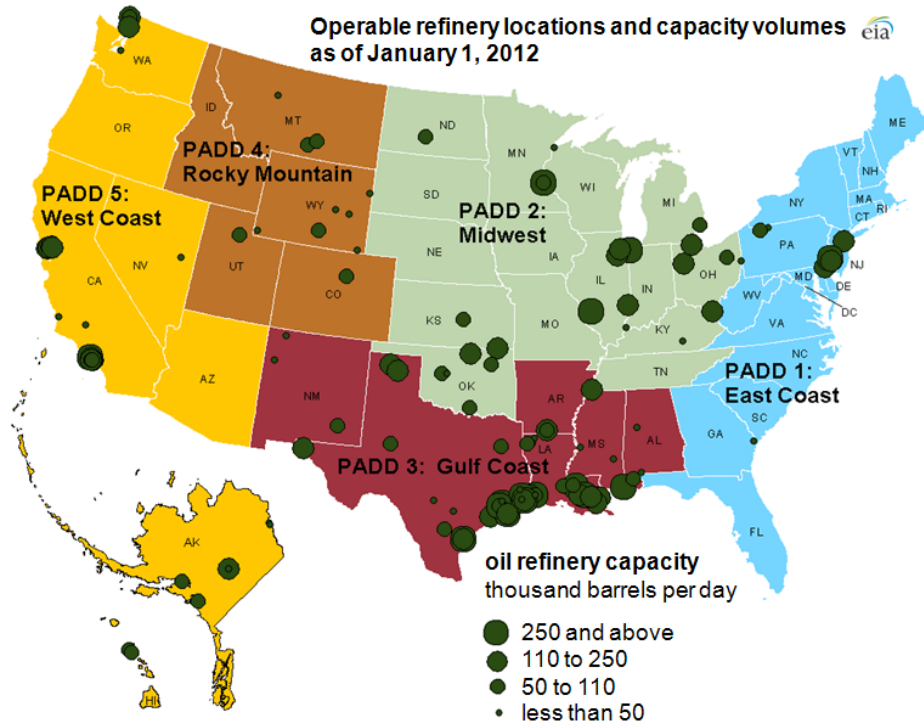
References

- Atkin, David and Dave Donaldson**, “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs,” Working Paper 21439, National Bureau of Economic Research July 2015. DOI: 10.3386/w21439.
- Bachmeier, Lance J. and James M. Griffin**, “New evidence on asymmetric gasoline price responses,” *Review of Economics and Statistics*, August 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*, July 1991, 13 (3), 211–218.
- Berardi, Nicoletta, Patrick Sevestre, Marine Tepaut, and Alexandre Vigneron**, “The impact of a soda tax on prices: evidence from French micro data,” *Applied Economics*, September 2016, 48 (41), 3976–3994.
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert**, “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?,” *The Quarterly Journal of Economics*, February 1997, 112 (1), 305–339.
- **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).
- Bulow, Jeremy I. and Paul Pfleiderer**, “A Note on the Effect of Cost Changes on Prices,” *Journal of Political Economy*, February 1983, 91 (1), 182–185.
- Chetty, Raj**, “Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods,” *Annual Review of Economics*, 2009, 1 (1), 451–488.
- Chiou, Lesley and Erich Muehlegger**, “Crossing the Line: Direct Estimation of Cross-Border Cigarette Sales and the Effect on Tax Revenue,” *The B.E. Journal of Economic Analysis & Policy*, 2008, 8 (1).
- Fabra, Natalia and Mar Reguant**, “Pass-Through of Emissions Costs in Electricity Markets,” *American Economic Review*, September 2014, 104 (9), 2872–2899.
- Fevrier, Philippe and Laurent Linnemer**, “Idiosyncratic shocks in an asymmetric Cournot oligopoly,” *International Journal of Industrial Organization*, June 2004, 22 (6), 835–848.
- 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.
- Harding, Matthew, Ephraim Leibtag, and Michael F. Lovenheim**, “The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data,” *American Economic Journal: Economic Policy*, 2012, 4 (4), 169–198.
- Jr., Joseph J. Doyle and Krislert Samphantharak**, “\$2.00 Gas! Studying the effects of a gas tax moratorium,” *Journal of Public Economics*, April 2008, 92, 869–884.

- Knittel, Christopher R., Ben S. Meiselman, and James H. Stock**, “The Pass-Through of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard,” Working Paper 21343, National Bureau of Economic Research July 2015. DOI: 10.3386/w21343.
- Lade, Gabriel E. and James Bushnell**, “Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard,” 2016.
- Marion, Justin and Erich Muehlegger**, “Fuel tax incidence and supply conditions,” *Journal of Public Economics*, October 2011, 95 (9), 1202–1212.
- Miller, Nathan H., Marc Remer, and Gloria Sheu**, “Using cost pass-through to calibrate demand,” *Economics Letters*, March 2013, 118 (3), 451–454.
- , **Matthew Osborne, and Gloria Sheu**, “Pass-through in a concentrated industry: empirical evidence and regulatory implications,” *The RAND Journal of Economics*, March 2017, 48 (1), 69–93.
- Rozema, Kyle**, “Tax Incidence in a Vertical Supply Chain: Evidence from Cigarette Wholesale Prices,” SSRN Scholarly Paper ID 3022786, Social Science Research Network, Rochester, NY August 2017.
- Stolper, Samuel**, “Who Bears the Burden of Energy Taxes? The Role of Local Pass-Through,” October 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/>.