Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom*

Erich Muehlegger †

Richard L. Sweeney ‡

February 2021

Abstract

In imperfectly competitive settings, a firm's price depends on its own costs as well as those of its competitors. We demonstrate that this has important

*We thank Ryan Kellogg, Katheryn Russ, Joe Shapiro, Jim Stock, Reed Walker and seminar participants at Harvard, Duke, University of Connecticut, University of Maryland, the NBER Hydrocarbon Infrastructure and Transportation workshop, the Empirical Methods in Energy Economics workshop, and Berkeley Energy Camp for helpful comments. Both authors declare they have no interests, financial or otherwise, that relate to the research described 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.

implications for the estimation and interpretation of pass-through. Leveraging a large input cost shock resulting from the fracking boom, we isolate price responses to firm-specific, regional and industry-wide input cost shocks in the US oil refining industry. The pass-through of these components vary from near zero to full pass-through, reconciling seemingly disparate results from the literature. We illustrate the policy implications of rival cost pass-through in the context of a tax on refinery carbon emissions.

JEL Codes: H22, H23, Q40, Q54

1 Introduction

Pass-through of taxes or input costs onto prices is a central policy parameter with wide ranging economic implications. Pass-through is used to estimate the incidence of actual taxes (e.g., Marion and Muehlegger (2011); Fabra and Reguant (2014); Stolper (2016)), the distributional impacts of trade tariffs (e.g., Amiti et al. (2019b); Fajgelbaum et al. (2019)), the beneficiaries of major entitlement program subsidies (e.g., Duggan et al. (2016)), and to forecast the incidence of hypothetical taxes (e.g., Ganapati et al. (2017)). It is used as a sufficient statistic for welfare analysis (e.g., Chetty (2009)) and to recover other objects of economic importance, such as trade costs (Atkin and Donaldson (2015)), international propagation of inflation shocks (Gopinath and Rigobon (2008); Auer et al. (2017) or demand elasticities (Miller et al. (2013)).

In this paper, we demonstrate that the nature of the tax or input cost shock used to estimate pass-through is *directly* related to the pass-through rate recovered. In oligopolistic settings, the price a firm sets depends directly on its own costs, but also indirectly on its rivals' costs. This simple observation has two important implications for pass-through analysis. First, when estimating a firm's response to a change in its own costs, it is important to account for rival responses as well. Beyond simple omitted variable bias from correlated cost shocks, the strategic response of seemingly untreated firms to rival cost changes invalidates them as a control group.[1] Second, when using pass-through for policy prediction, the identifying variation used in estimation must match the policy application. For exam-

[1] The empirical pass-through literature rarely incorporates the indirect effects of competitors' cost changes explicitly. Instead, papers often consider the robustness of estimates by limiting the control group to the most distant "untreated" firms. As examples, Dube et al. (2010) uses interior rather than border counties when estimating the effect of minimum-wage laws. Fowlie et al. (2012) considers geographically distant and unaffiliated facilities when examining the impacts of the RECLAIM program.

ple, a pass-through estimate identified from firm-level cost or tax shocks may substantially underestimate the price impact of an industry-wide policy change.

We demonstrate these points empirically by studying the response of prices in the U.S. oil refining industry to large input cost shocks resulting from the fracking boom. Our setting allows us to overcome three challenges that have limited the consideration of rival cost pass-through in past papers: (1) the need for comprehensive input and output prices for the universe of firms, (2) the ability to observe which firms are direct competitors, and (3) an exogenous input cost shock that differentially affects firms. We examine restricted-use microdata on the universe of oil refiners in the United States from the Energy Information Administration. For each firm, we observe crude oil procurement costs, detailed plant-level production decisions, and monthly prices and quantities sold for all refined products at the state-level. Such rich firm-level data on production, input and output prices are rarely available.[2] The confidential sales data allow us to incorporate the spatial supply patterns of the industry directly, rather than make assumptions about which firms directly compete.[3] Finally, our data span the U.S. fracking boom, which led to a near doubling of U.S. oil production between 2008 and 2015. Despite the scale of the shock, a set of regulatory, logistical and technological constraints limited the extent to which some refineries could benefit. The net result was a large reduction in oil input costs of some refiners, while the costs of other firms, producing identical products and often located nearby, remained largely unchanged.

The heterogenous impact of this cost shock allows us to separately identify the response

[2] Fabra and Reguant (2014) is a notable exception, examining bids and carbon intensities for the universe of Spanish electricity generators.

[3] Miller et al. (2017) also studies rival cost pass-through, documenting more-than-full pass-through of energy prices in the Portland cement industry, but does not observe firm-specific prices nor which firms directly compete.

to refiners' own-costs, their direct competitors' costs, and industry-wide cost shocks.[4] We estimate pass-through rates that vary from near zero, for firm-specific shocks, to full, for industry-wide shocks. Viewed as a continuum, these estimates reconcile seemingly conflicting pass-through estimates from the recent fuel pass-through literature. Using regional variation in energy prices, Ganapati et al. (2017) finds refineries are largely unable to pass-through costs; while using national variation in the cost of renewable fuel credits, Knittel et al. (2017) and Lade and Bushnell (2019) finds that the incidence falls fully on consumers.[5] We interpret the disconnect in these estimates as reflective of the distinction between idiosyncratic and shared cost variation, and show analytically that price changes will naturally vary for these two shocks, even after conditioning on a firm's own costs.

We conclude by demonstrating the relevance of these results for policy analysis. We consider a hypothetical carbon tax on refineries, the second largest industrial point-source of emissions in the U.S. When paired with a border adjustment tax, a carbon tax is equivalent to an industry-wide cost shock. Properly accounting for the indirect effects of competitors, our results suggest consumers bear virtually all of the tax. Moreover, 45 percent of firms, which are relatively less carbon-intensive than their rivals, more than fully pass-on an industry-wide carbon tax. Yet, an estimate of tax pass-through based on within market-time variation in costs would suggest prices rise by a mere five cents per dollar of tax imposed. This exercise highlights the importance of matching the cost variation used to estimate pass-through with the scope of the policy to which the estimate is applied.

---

[4]Sharing the spirit of this exercise, Amiti and Weinstein (2018) illustrates the importance of decomposing firm-borrowing and bank-supply shocks as drivers of firm investment and Amiti et al. (2019a) examines strategic complementarity in price setting amongst manufacturing firms.

[5]Similar results on full pass-through are obtained in the transmission of state fuel taxes (e.g., Doyle Jr. and Samphantharak (2008), Marion and Muehlegger (2011), and Stolper (2016)))

2 Pass-through and Imperfect Competition

A long-standing theoretical literature studies pass-through in imperfectly competitive markets. Bulow and Pfleiderer (1983) and Fevrier and Linnemer (2004) consider pass-through in Bertrand and Cournot markets, illustrating the possibility that pass-through can exceed one in oligopolistic markets; Weyl and Fabinger (2013) provide a more general framework. Similar links between pass-through and competition arise in the pricing-to-market literature (e.g., Bernard et al. (2003); Atkeson and Burstein (2008); De Blas and Russ (2015); Amiti et al. (2019a)) wherein endogenous firm-level markups are an outcome of the competitive environment. We draw upon these papers to illustrate how pass-through relates to the scope of a cost shock (that is whether a cost shock affects a single firm or a firm and its rivals), and to highlight the importance of accounting for competitor costs in empirical estimation.

Here, we consider the pass-through (ρ) of a tax (τ) onto the price of firm i, although an input cost shock, like the oil price shock we study, could be considered analogously. For expositional simplicity, we consider competition in prices, although we present examples for other models competition in Appendix A. Firm i sets profit-maximizing price p_i and faces marginal costs inclusive of taxes of α_i. We allow each firm to be differentially exposed to the tax, with $\frac{\partial \alpha_i}{\partial \tau}$ capturing the marginal per-unit tax rate faced by firm i.[6] Formally, we decompose the pass-through of the tax onto firm i's price as a direct effect and an indirect effect operating through i's competitors:

$$\rho_i(\tau) = \frac{\partial p_i}{\partial \alpha_i}\frac{\partial \alpha_i}{\partial \tau} + \sum_{j \neq i}\left[\frac{\partial p_i}{\partial p_j}\frac{\partial p_j}{\partial \alpha_j}\frac{\partial \bar{\alpha}}{\partial \tau} + \frac{\partial p_i}{\partial p_j}\frac{\partial p_j}{\partial \alpha_j}(\frac{\partial \alpha_j}{\partial \tau} - \frac{\partial \bar{\alpha}}{\partial \tau})\right] \tag{1}$$

where $\frac{\partial \bar{\alpha}}{\partial \tau}$ is the average marginal tax rate, and $\frac{\partial p_i}{\partial p_j}$ reflects firm i's optimal response to a change in firm j's price.

A tax affects a firm's choice of strategy both *directly*, through the firm's own costs, and

[6] In our empirical context, variation in exposure arises from differential access to low-cost crude oil and in our policy exercise, from differential carbon-intensity of production.

indirectly, in response to the strategic choice of the firm's competitors. Thus, a tax borne by the firm *alone* will be passed-through at a different rate than a tax affecting the firm and its competitors. We further decompose the indirect effect into two components - the effect of the mean shock to the firm's competitors and the competitor-specific deviation from the shock, reflected in the second and third terms of equation (1). If firm i's competitors face different tax exposure, pass-through depends on the covariance between relative exposure $(\frac{\partial \alpha_j}{\partial \tau} - \frac{\partial \bar{\alpha}}{\partial \tau})$ and the degree to which firms i and j are close rivals, $\frac{\partial p_i}{\partial p_j}$. If a firm's closest rivals are more heavily impacted by a tax, the firm can pass-through their shock to a greater degree.

These results have three implications relevant to the estimation and interpretation of pass-through. First, omitting shocks to rival firms will bias own-cost pass-through estimates if shocks are correlated across firms.[7] Second, firms which are not directly affected by a shock may still adjust their prices if their rivals are affected. This behavior violates the stable unit treatment value assumption implicit in common estimation strategies, which use the price of rivals not directly affected by a cost shock as contemporaneous counterfactual. Finally, empirical pass-through rates are specific to the identifying variation used for estimation. Even after conditioning on the shock faced by a firm, the resulting price change will vary depending on whether that shock affects only the firm or is shared by its rivals as well. To compute the full price response to a cost shock or policy change, it is necessary to incorporate both the direct and indirect effects. However, empirical models that rely on rich temporal or spatial fixed effects may subsume the latter. This constrains the econometrician to estimating the pass-through rate of a firm-specific shock, which may or may not be of primary policy interest. We illustrate each of these points in our empirical results.

[7] Pennings (2017) notes a similar bias if real exchange rate shocks are correlated across countries.

3 Empirical Setting and Institutional Details

Oil refineries process crude oil into refined end-products like gasoline and diesel. Although the United States imports substantial amounts of crude oil, U.S. refineries produce almost all end-products consumed domestically. Most U.S. refineries are located proximate to traditional oil deposits, such as Texas, Louisiana and California, rather than areas that demand refined fuels.[8] The industry relies heavily on pipelines to move crude oil to refineries and to ship refined products to demand centers. These pipeline networks dictate both the locations from which refineries can cheaply acquire crude and end markets to which refiners can competitively ship. Despite this spatial mismatch of supply and demand, entry is limited - no new refineries have been built since the early 1980's.

3.1 Data

Through a confidential data request, we obtained data on refinery operations from the Energy Information Administration, described in detail in appendix Table B.1. A key feature of the EIA data is that we directly observe output prices and crude input costs (which account for over 90% of variable costs), allowing us to directly observe margins for each firm. Every firm that owns a refinery in the United States reports the total volume and cost of crude oil acquired both domestically and abroad each month, by Petroleum Administration Defense District (PADD).[9] These firms also report detailed monthly production data at the *refinery* level, including the quality of crude oil used. On an annual basis, they also report operable capacity and exhaustive information on the technology installed. Finally, at the *state* level, each firm reports the monthly volume sold and average price of all end products,

[8]Figure B.1 maps refineries by capacity.

[9]PADD's are a commonly used geographic aggregation for the industry dating back to World War II. For a map, see Figure B.1.

broken out by sales to end users (retail) and sales for resale (wholesale).[10] The sample in which all surveys are observed spans 2004-2015, and the resulting 9,215 firm-PADD-month observations are summarized in Table B.3.

3.2 Fracking background

The EIA data spans the U.S. fracking boom, which is the source of our input cost shock. Shortly after the turn of the century, two complementary technologies, hydraulic fracturing and horizontal drilling, rapidly matured, unlocking trillions of dollars of previously uneconomical domestic oil reserves. The ensuing transformation of the U.S. oil industry is documented in the four panels of Figure 1.

U.S. crude oil production nearly doubled between 2008 and 2015, abruptly reversing a decades-long trend towards the use of imported oil (Figure 1a). As a consequence, a previously superfluous 1977 ban on the export of crude oil became binding, and U.S. crude oil spot prices diverged from international crude markets (Figure 1b). At the start of the boom, from 2010 to 2015, the West Texas Intermediate ("WTI") spot price averaged $12 per barrel below the Brent spot price, in comparison to trading at an average $1.41 per barrel premium from 2000 to 2010.

Within the United States, the gains from the fracking boom were not equally shared among refiners. The most productive shale plays were in the middle of the country, while domestic production elsewhere continued to decline (Figure 1a). In the Gulf Coast, shale oil was proximate to major conventional oil fields, the largest concentration of refineries, and the epicenter of the crude distribution system. But in the Midwest, primarily North Dakota, relatively little pipeline capacity existed to move "tight" oil from shale fields to refineries. As

[10] Where a firm owns multiple refineries in a PADD, we assume that the firm supplies states from the refinery with the lowest transportation costs, described further in appendix B.2.1.

Figure 1: Shale Boom and U.S. Refining

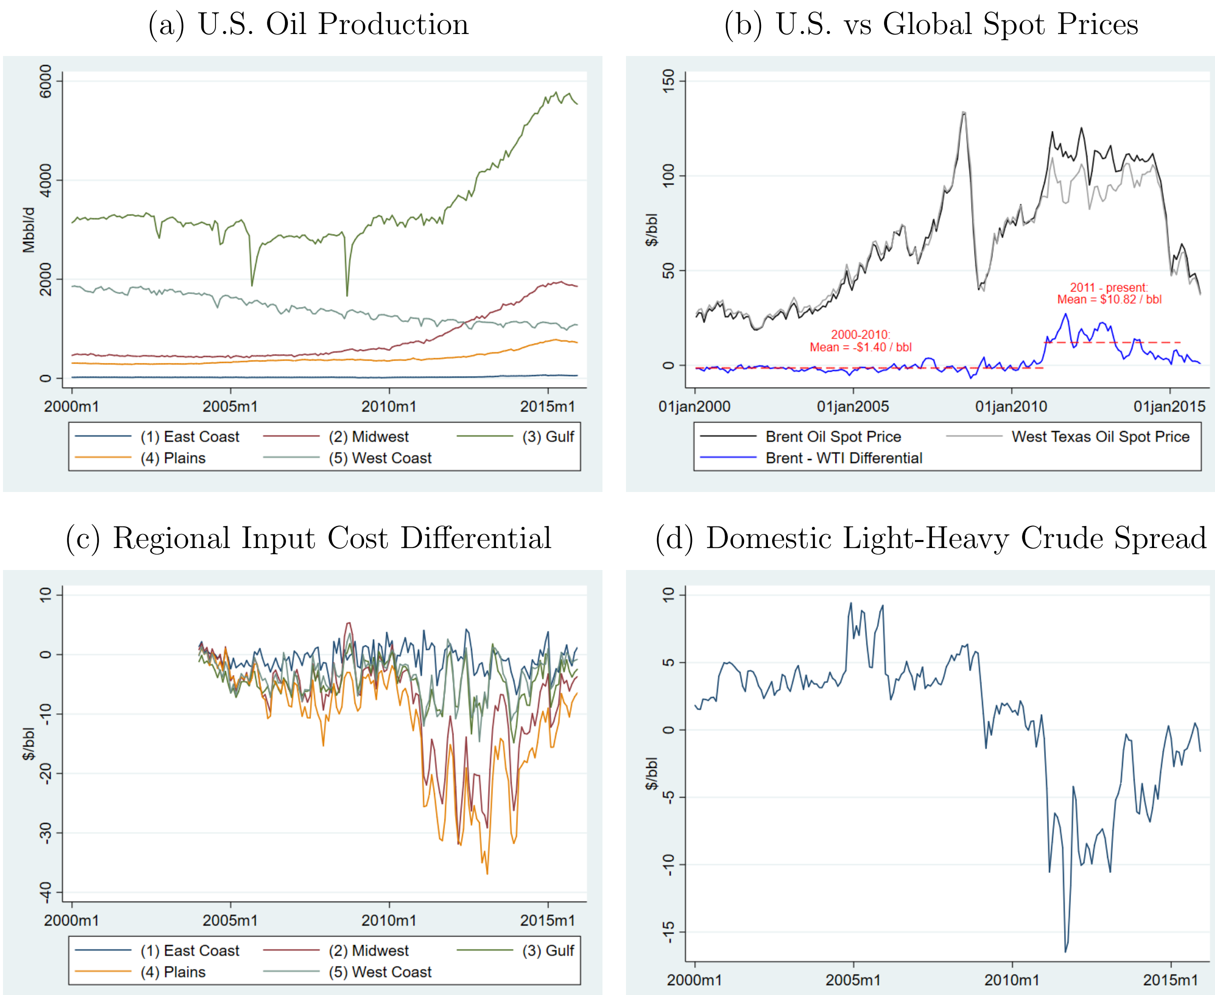

Description: (a) U.S. oil production by PADD; (b) West Texas Intermediate (WTI) and Brent spot prices; (c) Average crude price paid by refiners, minus the contemporaneous Brent spot price, by PADD (PADD-level data only available beginning in 2004); and, (d) U.S. light - heavy crude spread, as approximated by the difference between the average price for crudes with API gravity between 35 and 40, minus the average price for crudes with API gravity between 20 and 25.

a result, refineries in the Rocky Mountains (PADD 4) and Midwest (PADD 2) could acquire crude oil at a significant discounts, while crude costs for coastal refineries tracked foreign markets more closely (Figure 1c).

Within region, some refiners were better positioned to benefit from the boom than others, due to the attributes of shale oil. Crude oil is highly differentiated, and refineries are finely tuned to process specific types of crude oil. The most important dimension along which crudes are differentiated is density. Dense crudes (measured by a low API gravity) require more processing, more sophisticated capital, and typically trade at a discount relative to "lighter" crudes. But, the oil produced from the most active shale plays in the United States has been very light (high API gravity). The boom in light oil production reversed the historical correlation between price and crude density (Figure 1d). This provided a large cost advantage for refineries designed to process light oil, while refineries set up to process heavy crude were unable to benefit without incurring large adjustment costs or capital investment. Gradually, as infrastructure investments relaxed transportation constraints, the advantage enjoyed by certain refiners diminished. In December 2015, the export ban was lifted, and the gap between U.S. spot prices and world prices disappeared.

In sum, domestic refineries located near constrained shale deposits, tailored to process light domestic crude benefited disproportionately from the shale boom. In Appendix C.4, we confirm this with a panel regression of refiner input costs spanning the shale boom. Refineries proximate to shale deposits in the Midwest and Rocky Mountains saw crude costs fall by relative to the rest of the industry (by 5 and 7 dollars per barrel on average, respectively). However, even conditional on region, refineries processing "light", domestic crude in the pre-period, or with the technological capability to more easily switch between crudes, experienced statistically and economically significant relative reductions in crude prices during the shale boom.

4 Estimation and Results

We use the exogenous shock to input prices caused by the fracking boom to demonstrate how the nature of an input cost shock and the inclusion of rivals' costs are central to the estimation and interpretation of pass-through. We begin with the canonical pass-through regression, which projects the price a firm f receives in a given state m at time t onto its own costs.[11]

Consider first,

$$Price_{fmt} = \alpha Cost_{ft} + X'_{fmt}\delta + \nu_{fm} + \mu_t + \epsilon_{fmt} \tag{2}$$

where X contains other demand and supply-side factors which may shift the level of prices.[12] Here we estimate the model at the firm-state-month level, including firm-state and month-of-sample fixed effects. The dependent variable is average wholesale revenue per gallon[13], across all end products sold into the state.[14] $Cost_{ft}$ is the average price of crude oil per

[11]Although in our case, we focus on input costs, a regression of prices on own-input costs, own-tax liability or own-regulatory exposure is a standard approach for estimating pass-through when comparing "treated" and "untreated" firms.

[12]Consistent with Borenstein et al. (1997) that finds oil price changes are incorporated within several weeks into terminal prices, we focus on contemporaneous changes in costs and prices. In contrast, the macro and trade literature (e.g., Boivin et al. (2009); Maćkowiak et al. (2009); Andrade and Zachariadis (2016)) emphasizes menu costs, sticky prices and rational inattention as explanations for gradual incorporation of global and sectoral shocks.

[13]In this sample, roughly 15 percent of refiner sales are retail. Results using a total average price, inclusive of retail sales, instead of the wholesale price as the dependent variable are similar.

[14]The literature to date has typically focused on gasoline prices. However, all refiners are multi-product firms, with less than half of a barrel being converted into gasoline. Importantly, these products are produced *jointly*, in a fundamentally non-separable production

gallon.[15] Column 1 in Table 1 presents the results. The coefficient on own-cost indicates that a \$1 increase in a firm's crude costs leads to an average increase in price of \$0.067.

A natural concern with the estimate in column 1 is the possibility of omitted factors that affect prices and also co-vary with a firm's own costs within a time period. Of particular interest in our setting are spatially-correlated cost or demand shocks. If present, the estimated coefficient conflates the direct response to a firm's own cost with the indirect response to the omitted shock to competitors. In column 2, we replace the month-of-sample fixed effects with state-month-of-sample fixed effects as a way to address spatially correlated cost shocks. Here, deviations in a firm's own costs relative to the average costs firms serving the same state identify the parameter on a firm's own-cost. Our estimate of own-cost pass-through falls slightly, to 5.3 percent. This is consistent with positive correlation in cost shocks and strategic complementarity of prices, although the two point estimates are not statistically distinguishable.

An alternative would be to condition on competitors' costs directly. Consider augmenting equation (2) as follows,

$$Price_{fmt} = \alpha Cost_{ft} + \beta f(RivalCost_{-f,mt}) + X'_{fmt}\delta + \nu_{fm} + \mu_t + \epsilon_{fmt} \tag{3}$$

where $f(RivalCost_{-f,mt})$ denotes a weighted average of f's competitors' costs in m at time t.[16] Model 3 returns to the use of firm-state and month-of-sample fixed effects from column

process. As such, single product markups will be misleading, since they do not reflect substitution between products or the impact of this substitution on other products' prices. With this caveat, product-specific results are available from the authors upon request.

[15]Ideally $Cost_{ft}$ would reflect the marginal cost of supplying market m at time t. We observe average crude costs, which in our setting, closely track the WTI crude spot price, as documented in Appendix D.1.

[16]We consider only reduced-form pass-through approximations, consistent with the vast

Table 1: Fixed Effect Comparison Table

	(1)	(2)	(3)	(4)	(5)
Own	0.0668***	0.0534***	0.0521***	0.0393***	0.0447***
	(0.0131)	(0.0133)	(0.0131)	(0.0141)	(0.0133)
Rival			0.173***	0.938***	0.282***
			(0.0427)	(0.0147)	(0.0164)
Brent Spot					0.625***
					(0.0113)
Time FE	MoS	MoS-St	MoS	Y,M	Y,M
N	71570	71489	71570	71570	71570
r2	0.962	0.973	0.962	0.939	0.945

This table presents the results of estimating Equation (3) using total average wholesale prices as the dependent variable. Regressions are run at the firm-state-month level, and include firm-state fixed effects. Time FEs "Mos" reflect month-of-sample dummies, "Y-M" reflect year-month dummies, while "Y,M" implies separate year and month dummies. Rival costs include the average crude price of all other firms, weighted by the inverse shipping cost of serving a particular state. Standard errors, clustered at the firm-state level, are presented in parentheses. All models include demand shifters (state population, income, heating and cooling degree days) and supply shifters (proportion of retail sales, API gravity, and operating refinery capacity).

1, but includes the average cost of other refiners, weighted by inverse shipping costs.[17] After accounting for rivals' costs, we see a modest reduction in own-cost pass-through, similar to column 2. In this specification, strategic complementarity in prices is now observable – a firm's price increases by 17 cents for every dollar increase in the average costs of competitors. This implies a 22 percent pass-through rate for a common shock experienced by a firm and its competitors.

Just as the inclusion of market-time fixed effects precluded our ability to recover the pass-through of cost shocks faced by local rivals, the inclusion of time fixed effects in model 3 subsumes industry-wide cost shocks common to all firms in a given period. To estimate the effects of an industry-wide shock, we must further coarsen the fixed effects. Column 4 replaces month-of-sample fixed effects with year fixed effects and calendar-month fixed effects, to control for time trends and seasonality. As in column 1, one might worry that omitted shocks (here to foreign refiners) might be correlated with domestic input prices. Thus, in Column 5 we add the Brent spot price, to capture variation in world oil prices and act as a proxy for the costs international refineries. As in column 1, the coefficient on rivals' costs in column 4 are biased upwards (significantly) by the omission of the Brent crude spot price. Once we include the Brent crude spot price in column 5, the magnitude of the coefficients on own-costs and rival-costs are similar to those in column 3. This suggests

majority of the empirical pass-through literature, and the literature on pass-through in the refining industry more specifically (e.g., Hastings and Gilbert (2005), Brown et al. (2008)). This representation, with average competitor costs entering linearly, will only recover the structural pass-through rate under Cournot competition with linear demand. For a discussion of bias in reduced form measures of pass-through, see MacKay et al. (2014). In our setting, our reduced-form estimates reflect the net effect of strategic behavior with respect to both pricing and production quantities.

[17] Additional information on shipping cost construction is provided in Appendix B.2.1. Results using geographic distance are similar. Alternative weightings are discussed below.

that foreign competition disciplines U.S. refiners' ability to pass-along domestic cost shocks.

Looking progressively down the rows in column 5 demonstrates the relationship between the scope of a cost-shock and the degree to which it is passed-onto downstream prices. Focusing solely on firm responses to own cost shocks suggests a low rate of pass-through, conditional on the average cost of all rivals. However, including the second row reveals that a cost shock experienced by all domestic refineries would be passed-through at a higher rate of 33 percent. Broadening scope of the shock further, an industry-wide cost shock which affects both domestic and foreign refineries would be (roughly) fully passed-through to the consumer. [18] While it is natural to question whether eschewing fine time-space fixed effects to simultaneously recover all three margins will bias the estimates idiosyncratic cost shock pass-through, a comparison of the own and rival cost pass-through estimates across columns reveals little tradeoff in this example.

Next, we consider how the competitive proximity of a rival affects the indirect pass-through rate. In many settings, it is difficult to know if two firms are direct competitors. Products may be unobservably differentiated, or logistical constraints may limit the competition of spatially-proximate firms.[19] In our case, the data allow us to observe the frequency with which two rival firms supply the same (undifferentiated) product to the same state. A

[18]Increasing rates of pass-through for progressively more "expansive" shocks echo the literature on exchange rate pass-through. Exploiting variation in real exchange rates (akin to a domestic cost shock), Gopinath and Rigobon (2008); Fitzgerald and Haller (2013); Auer and Schoenle (2016) find endogenous markups manifest in relatively low rates of pass-through onto import prices, on the order of 0.15 - 0.30. In contrast, Nakamura and Zerom (2010) find close to full pass-through from the wholesale coffee price indices to retail prices, consistent with our finding of nearly full pass-through of industry-wide shocks.

[19]For example, due to transportation constraints, refineries in the Gulf Coast supply the vast majority of refined products to New York, rather than "nearer" refineries in the Upper Midwest.

natural decomposition is, therefore, to estimate separate pass-through rates for the costs of "direct" rivals, which typically serve a state, and the costs of the "fringe" of firms who do not.[20] Model 2 in the top panel of Table 2 presents the results (for comparison, model 5 from the previous table is reproduced in column 1). Less than half of the weighted rival cost pass-through estimated in model 1 comes from firms directly serving the state, while the remainder comes from "fringe" competitors.[21]

[20]One justification for including firms in other states is that they are potential entrants, disciplining markups. Another is that seemingly non-competing firms are often bound strategically when some firms serve multiple markets, as discussed in (Bulow et al., 1985).

[21]In some markets, such as electricity, strict capacity constraints and heterogeneous costs impose a natural "dispatch" ordering across firms. In these models, the price would be set by the "marginal firm", rather than the average over inframarginal firms serving a market. While these forces are not present in refined product markets, we present results on pass-through of the highest cost rival in a market in appendix Table D.4.

Table 2: Competition measure results

(a) State Level Results

	(1)	(2)	(3)	(4)
Own	0.0447***	0.0485***	0.0606**	0.0704***
	(0.0133)	(0.0134)	(0.0255)	(0.0252)
Rival	0.282***	0.128***	0.146***	0.0357
	(0.0164)	(0.0214)	(0.0299)	(0.0387)
Fringe		0.159***		0.112***
		(0.0220)		(0.0357)
Brent Spot	0.625***	0.617***	0.732***	0.722***
	(0.0113)	(0.0101)	(0.0124)	(0.0105)
Rival Measure		Avg		Avg
IV			Yes	Yes
fstat			4651	3137
N	71570	71529	71570	71529

(b) Firm Level Results

	(1)	(2)	(3)	(4)
Own	0.0493	0.0552*	-0.00498	0.0102
	(0.0312)	(0.0286)	(0.0517)	(0.0501)
Rival	0.285***	0.204***	0.211***	0.163*
	(0.0368)	(0.0408)	(0.0663)	(0.0858)
Fringe		0.0604		0.00838
		(0.0379)		(0.0654)
Brent Spot	0.622***	0.635***	0.736***	0.757***

	(0.0181)	(0.0168)	(0.0252)	(0.0235)
Rival Measure		Avg		Avg
IV			Yes	Yes
First-stage F			288	190
N	9169	9169	9169	9169

This table presents the results of estimating Equation (3) 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-PADD-month level and includes firm-PADD fixed effects. All specifications include year fixed effects and month fixed effects. In columns 1 and 3, rival costs are the average costs of all firms, inverse shipping cost weighted. In columns 2 and 4, rival costs are the average costs of firms that serve the location in the same calendar year, and fringe costs are the average of all other firms' costs, inverse shipping cost weighted. Standard errors are in parentheses, clustered at the firm-state level in panel (a) and the firm-PADD level in panel (b). All models include demand shifters and supply shifters.

When researchers have multiple price observations per firm in a single period, a decision must also be made about the appropriate level of aggregation. In our setting, we observe refiners serving multiple states each month. In panel (a), "fringe" firms in one state might be direct rivals in another. Moreover, since the products supplied are homogenous and produced in the same facility, a profit maximizing firm will equate expected marginal revenues (net of transportation costs) across states. In panel (b) we repeat these two specifications after aggregating to the firm-PADD level. The dependent variable is the average wholesale price earned by the firm in the PADD. Intuitively, the own-cost results are very similar a higher level of geographic aggregation. At either the state or the regional level, firms have relatively little ability to pass-on cost shocks unique to themselves. However the direct rival costs in column 2 of panel (b) are now larger than they were in panel (a), while the fringe firm's costs are attenuated. This suggests that the importance of "fringe" firms in panel (a) was primarily driven by firms serving other states in the same PADD rather than more distant competitors.

We repeat these specifications using instrumental variables to address two potential problems. First, if realized costs are correlated with unobserved demand shocks, this would bias estimated pass-through rates upwards. Second, while we observe firm-specific crude prices, which comprise over 90% of variable costs, storage might lead to measurement error in costs. If large, this would attenuate our pass-through estimates. Moreover, given that our rival and fringe cost variables are actually averaged over many firms, while own costs reflect a single refinery's reported costs, the latter could be more prone to measurement error, complicating a direct comparison of their relative import for own prices.

For each refinery-month, we construct two sets of instruments. As was discussed in Section 3.2, the shale boom primarily benefitted refiners processing light, domestic crude, located in the Plains or upper Midwest. Following the intuition of Bartik (1991), we construct own-price instruments by interacting time-invariant measures of pre-shock refinery characteristics, namely the density of crude used and an indicator for being located in PADD 2 or 4, with time-varying national benchmark crude input prices.[22] Instruments for rival and fringe

[22] In our setting, we take pre-shock refinery characteristics to be exogenous as they are determined by capital decisions made well before the innovation of hydraulic fracturing, similar to framework highlighted in Goldsmith-Pinkham et al. (2018). We provide further details on the construction of the instruments in Appendix D.2. We also present the first-stage estimates, which have intuitive signs and considerable explanatory power. In our setting refiners sell almost-identical products, and hence, unobserved shocks correlated with

firms are constructed by averaging the own-cost instruments for these firms accordingly.

Columns 3 and 4 of Table 2 present the results. Overall, the IV and OLS estimates are qualitatively consistent. At the state level, the IV estimates of own-cost pass-through are slightly higher, suggesting some degree of measurement error. Conversely, rival costs, which are averaged and thus less prone to measurement error, are passed-through at a considerably lower rate. This is also consistent with the possibility of unobserved demand shocks at the state level being correlated with the costs of firms serving those states. At the PADD level, where own cost measurement error is less severe, the own cost estimates now appear smaller (although the confidence intervals include the OLS estimate). However at this level, rival costs are only sightly smaller than the OLS estimates.

5 Implications for a Tax on Refinery Emissions

We illustrate the implications of our estimates for a hypothetical tax on refinery carbon emissions.[23] Approximately 20 percent of the lifecycle emissions from gasoline occur prior to the pump, of which half come from the refining process (Lattanzio, 2014). With average annual emissions of 1.22 million metric tons of CO2 equivalent (MMT CO2e), refineries are the second highest-emitting domestic point-sources, behind the electric power sector. Collectively, the 145 domestic refineries produce approximately 3% of total U.S. GHG emissions.[24]

---

pre-shock characteristics are less likely to lead us to overstate precision, than in the setting of regional manufacturing sector shares highlighted in Adao et al. (2019).

[23]In this section, we only consider a tax on carbon emissions from processing fuel not from the fuel itself. Taxing embedded emissions would imply a cost shock that is approximately nine times larger than the one considered here, which may in turn induce complicated input and demand substitution response that complicate the simple counterfactual considered here.

[24]Authors' calculations, based on data from EPA Greenhouse Gas Reporting Program, https://www.epa.gov/ghgreporting, accessed January 12, 2018.

Taxes levied on facility emissions at the current social cost of carbon ($51/ton) would raise refiners cost of production by $1.43 per barrel (bbl) processed on average.[25]

We begin by considering the implications of our estimates for the incidence of taxes of different scope, *on average*. Consider first a tax applying to a single refinery. The estimates from column (2), panel (b) of Table 2 predict that prices for the average refinery would only increase by $0.08/bbl. If the scope of the policy were expanded to include a firm and all of its regional competitors, prices would rise by $0.38/bbl. Applying the $51/ton tax to all domestic refiners would increase prices by $0.46/bbl, while pairing it with an equivalent border adjustment tax on imports (effectively taxing the industry) would cause prices to rise by $1.37/bbl on average. In the first case, producers bear close to the full burden of the tax, while in the final case consumers do.

Interpreting pass-through predictions as a function of the scope of a shock, our results reconcile seemingly disparate estimates of refinery pass-through from the recent literature. Ganapati et al. (2017) use regional variation in energy input costs and estimate that 24 - 33% of a carbon tax on refineries would be born by consumers. This lines up closely with our estimates of pass-through for a shock affecting a firm and at its regional rivals, of 26 percent. In contrast, Knittel et al. (2017) studies common input cost shocks stemming from the Renewable Fuel Standard (RFS) that apply to every gallon of surface transportation fuel sold in the U.S., regardless of origin. In this case, incidence of RFS credit prices falls almost fully on consumers, consistent with our estimates of pass-through of a border adjustment tax in both spirit and magnitude.

While average incidence provides a useful guide for the policy impacts a carbon tax, the carbon intensity of the refining process varies substantially across domestic refineries. Differences in capital investment allow some refineries to use denser crudes or to maximize gasoline yields, both of which increase CO2 emissions per barrel. These differences translate

[25]For context, the average reported refining margin from 2004-2009 (the last year available) in the EIA Financial Reporting System was $2.98/bbl.

Figure 2: Within-industry Heterogeneity in Carbon Intensity and Pass-through Rates

(a) Carbon Taxes (b) Pass-through

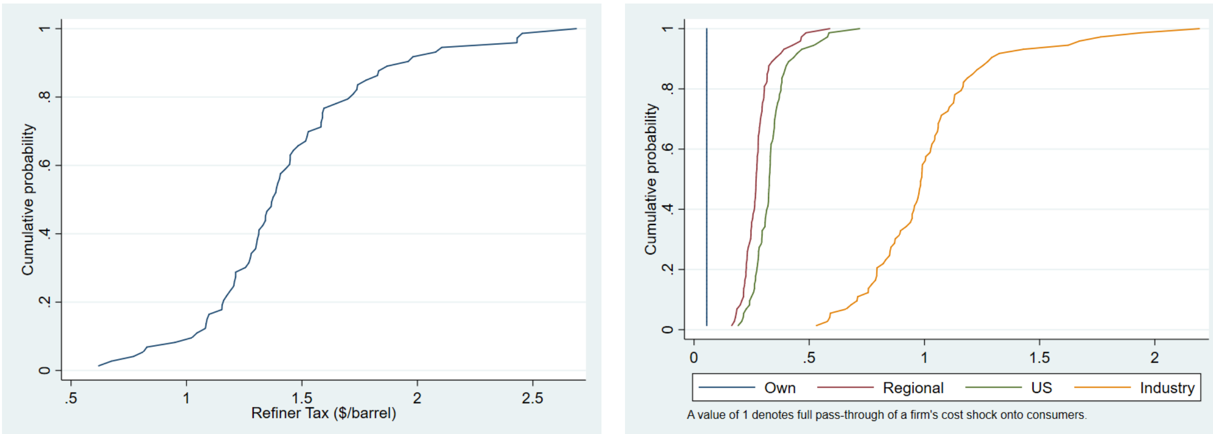

Notes: Pass-through estimates in panel (b) based on model (5) of panel (b) in Table 2.

into substantial heterogeneity in tax exposure under a $51 per ton carbon tax (graphed in Figure 2a). The effective tax in dollars per barrel of crude varies from under $1 for the least carbon-intensive decile, to over $2 for the most carbon-intensive decile.

Heterogeneity in carbon intensity leads to dispersion in firm-level pass-through rates. As was noted earlier when discussing equation (1), if a firm's closest rivals are more carbon-intensive, a tax provides the firm a competitive advantage, allowing the firm to pass-through a greater fraction of the tax. In contrast, competition with less carbon-intensive rivals may discipline the firm's ability to pass-through the tax.

Figure 2b graphs the distribution of firm-level pass-through rates for the four policies described above. Under a tax that applies to a single firm (leftmost line), firm-level pass-through rates are identical, since rivals are unaffected.[26] Extending to a regional tax, firm-level pass-through rates begin to differ. Some firms pass-through as little as 20 percent of their cost change, while others pass-though more than half. Here, a carbon tax disad-

[26]Heterogeneity in pass-through may also arise from differences in regional demand or supply conditions. Tables that include interaction with firm size and HHI are not substantively different than our main results, and are provided in Appendix Table D.5.

vantages carbon-intensive firms relative to less carbon-intensive rivals, limiting the ability of relatively carbon-intensive firms from passing-on the carbon tax. Extending the tax to the national level increases the average pass-through rate, but also further increases the pass-through heterogeneity. Finally, the rightmost line presents the distribution using the full-pass through estimates across an domestic carbon tax, coupled with a commensurate border tax on imported fuels. Although consumers bear virtually all of the tax on average, firm-level pass-through rates vary substantially. Fifty-five percent of refineries pass-through the carbon tax less than fully, bearing some of the costs of the carbon tax. But, the remaining 45 percent of refiners experience an *increase* in markups under the tax. Although the tax affects all firms, less carbon-intensive firms are competitively advantaged relative to more carbon-intensive rivals, allowing them to *more than fully* pass on the carbon tax. Such variation in firm-level pass-through rates is of direct political importance when considering whether particular firms might oppose (or support) such a policy.

6 Conclusion

Pass-through is an important tool with wide ranging economic applications. The extent to which prices change after a policy or event is also of paramount interest to policymakers. Due to its policy import and conceptual simplicity, pass-through is widely estimated. In this paper, we call attention to an underappreciated aspect of pass-through analysis: strategic responses. In imperfectly competitive settings, the price a firm sets is a function of not just its own costs, but those of its rivals as well. Using a simple framework, we demonstrate that the link between the price a firm sets and competitor costs has important implications for estimation, interpretation and application of pass-through. These spillovers not only confound many standard research designs, but also they cannot simply be subsumed or ignored in estimation if the parameter of interest is the full response of prices to a common shock.

Using rich data from the U.S. oil refining industry, we demonstrate the empirical relevance of these points. We study the fracking boom, which upended the U.S. oil market and generated deep input cost discounts for some refineries but not others. Leveraging this shock, we demonstrate that firm prices do respond to competitor costs in practice, and that these own price adjustments vary intuitively with the proximity and scope of the shock. Comparing input cost shocks of different scope, our estimates of pass-through vary from near zero for firm-specific shocks to near one for shocks affecting all firms in the industry, reconciling seemingly disparate estimates of input cost pass-through from the literature.

Finally, we illustrate the benefit of explicitly considering indirect cost pass-through in the context of a hypothetical carbon tax on refinery emissions. Conditional on the direct effect on a firm's costs, the actual incidence borne varies dramatically depending on the extent to which that firm's rivals are taxed as well. Furthermore, we demonstrate that strategic spillovers generate large heterogeneity in pass-through across firms. This wide heterogeneity amongst firms may partially explain some of the intransigence of industry groups to carbon prices, despite the fact that the average incidence will fall primarily on consumers.

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Amiti, Mary and David E Weinstein, "How much do idiosyncratic bank shocks affect investment? Evidence from matched bank-firm loan data," *Journal of Political Economy*, 2018, *126* (2), 525–587.

_ , Oleg Itskhoki, and Jozef Konings, "International shocks, variable markups, and domestic prices," *The Review of Economic Studies*, 2019.

_ , Stephen J Redding, and David Weinstein, "The Impact of the 2018 Trade War on US Prices and Welfare," Technical Report, National Bureau of Economic Research 2019.

Andrade, Philippe and Marios Zachariadis, "Global versus local shocks in micro price dynamics," *Journal of International Economics*, 2016, *98*, 78–92.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 2008, *98* (5), 1998–2031.

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.

Auer, Raphael A and Raphael S Schoenle, "Market structure and exchange rate pass-through," *Journal of International Economics*, 2016, *98*, 60–77.

_ , Andrei A Levchenko, and Philip Sauré, "International inflation spillovers through input linkages," *Review of Economics and Statistics*, 2017, (0).

Bartik, Timothy J., "Who Benefits from State and Local Economic Development Policies?," Books from Upjohn Press, W.E. Upjohn Institute for Employment Research 1991.

Bernard, Andrew B, Jonathan Eaton, J Bradford Jensen, and Samuel Kortum, "Plants and productivity in international trade," *American Economic Review*, 2003, *93* (4), 1268–1290.

Blas, Beatriz De and Katheryn N Russ, "Understanding markups in the open economy," *American Economic Journal: Macroeconomics*, 2015, *7* (2), 157–80.

Boivin, Jean, Marc P Giannoni, and Ilian Mihov, "Sticky prices and monetary policy: Evidence from disaggregated US data," *American Economic Review*, 2009, *99* (1), 350–84.

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.

Brown, Jennifer, Justine Hastings, Erin T. Mansur, and Sofia B. Villas-Boas, "Reformulating competition? Gasoline content regulation and wholesale gasoline prices," *Journal of Environmental Economics and Management*, January 2008, *55* (1), 1–19.

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.

_ **, John D. Geanakoplos, and Paul D. Klemperer**, "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 1985, *93* (3), 488–511.

Chetty, Raj, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 2009, *1* (1), 451–488.

Dube, Arindrajit, T William Lester, and Michael Reich, "Minimum wage effects across state borders: Estimates using contiguous counties," *The review of economics and statistics*, 2010, *92* (4), 945–964.

Duggan, Mark, Amanda Starc, and Boris Vabson, "Who benefits when the government pays more? Pass-through in the Medicare Advantage program," *Journal of Public Economics*, 2016, *141*, 50–67.

Fabra, Natalia and Mar Reguant, "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, September 2014, *104* (9), 2872–2899.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The Return to Protectionism," Technical Report, National Bureau of Economic Research 2019.

Fevrier, Philippe and Laurent Linnemer, "Idiosyncratic shocks in an asymmetric Cournot oligopoly," *International Journal of Industrial Organization*, June 2004, *22* (6), 835–848.

Fitzgerald, Doireann and Stefanie Haller, "Pricing-to-market: evidence from plant-level prices," *Review of Economic Studies*, 2013, *81* (2), 761–786.

Fowlie, Meredith, Stephen P Holland, and Erin T Mansur, "What do emissions markets deliver and to whom? Evidence from Southern California's NOx trading program," *American Economic Review*, 2012, *102* (2), 965–93.

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.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," Technical Report, National Bureau of Economic Research 2018.

Gopinath, Gita and Roberto Rigobon, "Sticky borders," *The Quarterly Journal of Economics*, 2008, *123* (2), 531–575.

Hastings, Justine and Richard J. Gilbert, "Market power, vertical integration and the wholesale price of gasoline," *The Journal of Industrial Economics*, 2005, *53* (4), 469–492.

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," *Journal of the Association of Environmental and Resource Economists*, March 2017, *4* (4), 1081–1119.

Lade, Gabriel E. and James Bushnell, "Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard," *Journal of the Association of Environmental and Resource Economists*, 2019, *6* (3), 563–592.

Lattanzio, Richard K., "Canadian Oil Sands: Life Cycle Assessments of Greenhouse Gas Emissions," 2014.

MacKay, Alexander, Nathan H Miller, Marc Remer, and Gloria Sheu, "Bias in reduced-form estimates of pass-through," *Economics Letters*, May 2014, *123* (2), 200–202.

Maćkowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt, "Sectoral price data and models of price setting," *Journal of Monetary Economics*, 2009, *56*, S78–S99.

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.

Nakamura, Emi and Dawit Zerom, "Accounting for Incomplete Pass-Through," *The Review of Economic Studies*, July 2010, *77* (3), 1192–1230.

Pennings, Steven, "Pass-through of competitors' exchange rates to US import and producer prices," *Journal of International Economics*, 2017, (105), 41–56.

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.

Online Appendix for ”Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom”

A Pass-through and Competition

In this appendix, we illustrate the main points from Section 2 in the context of specific assumptions about the nature of competition.

In each of the cases below, we show that 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.

Following the notation in Weyl and Fabinger (2013), we use $\boldsymbol{\rho_\alpha}$ to denote the pass-through of a shock α onto the vector of firm-specific prices. Firm i chooses a unidimensional strategic variable, σ_i, to maximize profits. To illustrate the distinct effect of idiosyncratic versus common cost shocks, 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 shock ($\bar{\alpha}$) and a firm-specific shock (α_i). We characterize the pass-through of firm i as a direct effect of shock on firm i and an indirect effect operating through firm i's rivals.

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

A.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:

$$\frac{dQ}{d\alpha_i} = \frac{1}{(n+1)P'(Q) + QP''(Q)}; \quad \frac{dQ}{d\bar{\alpha}} = \frac{n}{(n+1)P'(Q) + QP''(Q)}. \tag{4}$$

These equations illustrate that the general points highlighted in Section 2 clearly apply to the static Cournot model. First, a firm-specific shock is passed along at a rate of $1/n$ relative to commensurate market-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).

A.2 Differentiated Nash-in-Prices

Next we consider the differentiated Nash-in-price model, where 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 $\frac{d\boldsymbol{p}}{d\sigma_i}$ is the one corresponding to a firm's own price, simplifying the expression for $\boldsymbol{\rho_\alpha}$. For expositional simplicity, we focus on the two firm case, for which we can express the pass-through of a cost shock $\boldsymbol{\alpha}$ onto firm i's price as:

$$\boldsymbol{\rho_\alpha} = \frac{\frac{\partial\sigma_i}{\partial\alpha} + \frac{\partial\sigma_i}{\partial\sigma_j}\frac{\partial\sigma_j}{\partial\alpha}}{1 - \frac{\partial\sigma_i}{\partial\sigma_j}\frac{\partial\sigma_j}{\partial\sigma_i}}. \tag{5}$$

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:

$$\frac{\boldsymbol{\rho_{\alpha_i}}}{\boldsymbol{\rho_{\bar{\alpha}}}} = \frac{\frac{\partial\sigma_i}{\partial\alpha_i}}{\frac{\partial\sigma_i}{\partial\alpha_i} + \frac{\partial\sigma_i}{\partial\sigma_j}\frac{\partial\sigma_j}{\partial\alpha_j}}. \tag{6}$$

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.

B Refining Industry Background and Data

B.1 Background on the U.S. Refining Industry

Oil refineries process crude oil into refined products. Although the U.S. imports substantial amounts of crude oil, much of the gasoline and diesel fuel consumed in the U.S. is produced U.S. refineries. Figure B.1 maps the locations of refineries scaled by their distillation capacity. The regions are Petroleum Administration Defense Districts (PADDs), which are commonly used geographic aggregation for the industry dating back to World War II. The regions correspond to: (1) East Coast, (2) Midwest, (3) Gulf Coast, (4) Rockies, and (5) West Coast.

Roughly fifty percent of U.S. refining capacity is located in areas with historical petroleum deposits, Texas, Louisiana and California. Other refineries are located proximate to major end markets like Chicago and Philadelphia and are served either by crude oil pipeline from the gulf coast or import terminals that deliver crude from international markets. From these refineries, a network of refined product pipelines transport gasoline and diesel fuel to wholesale terminals near most major metropolitan areas. A key feature of both pipeline systems is that they are incomplete

Figure B.1: Refinery locations and PADD Map

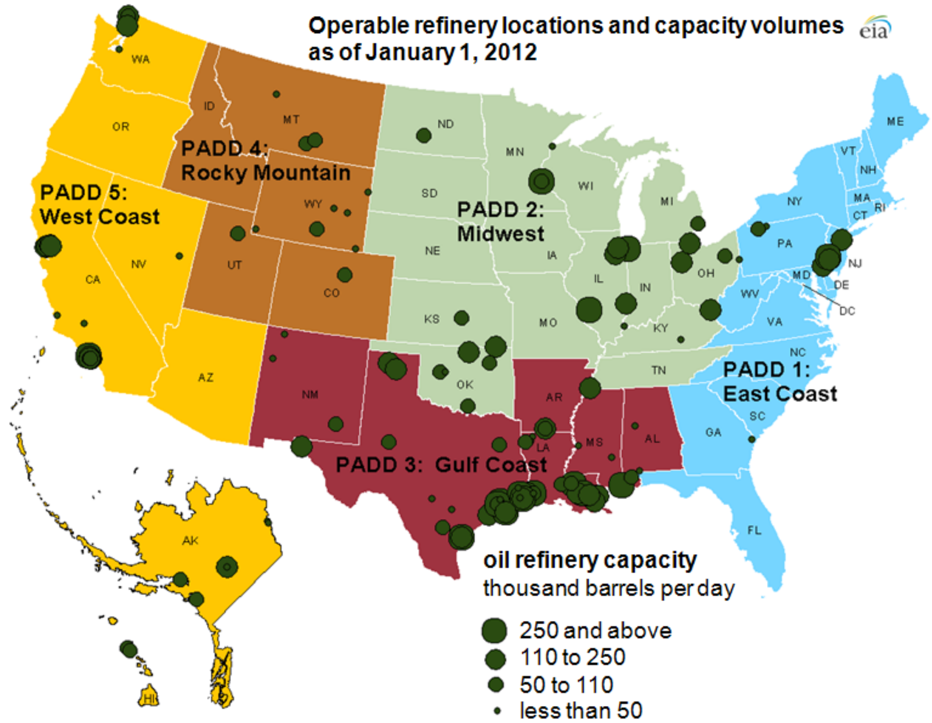

Source: EIA "Today in Energy"

(Figure B.2). As a result, not every refinery can obtain crude from every domestic oil field (the black lines) or send refined product to every end market (the dotted red lines). The pipeline map also leads to spatially complex patterns of competition. For example, refineries in Philadelphia face intense competition from refineries in the Gulf Coast, but operate in markets that are distinct from seemingly closer refineries in the Midwest. Refineries tend to operate at relatively high utilization rates. Refineries undergo planned maintenance for a couple of weeks every three to five years.(see e.g., "Refinery Outages: Description and Potential Impact on Petroleum Product Prices," EIA March 27, 2007. https://www.eia.gov/petroleum/articles/refoutagesindex.php.) These planned outages are scheduled (possibly years) in advance and may result in processing units being offline for several weeks at a time.

Crude oil is a differentiated input, and oil refineries are highly tailored to process specific types of crude. Crudes are differentiated mainly based on how dense it is and how much sulfur it contains. These characteristics define what refining equipment is necessary to convert crude into 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. In response to this input cost differential, some refineries made large capital investments in order to process lower quality crudes. As one example, BP's recent installation of a cracking unit, which processes heavy crude, was described as the largest private investment in Indiana history.

Figure B.2: Crude and Refined Product Pipelines Map

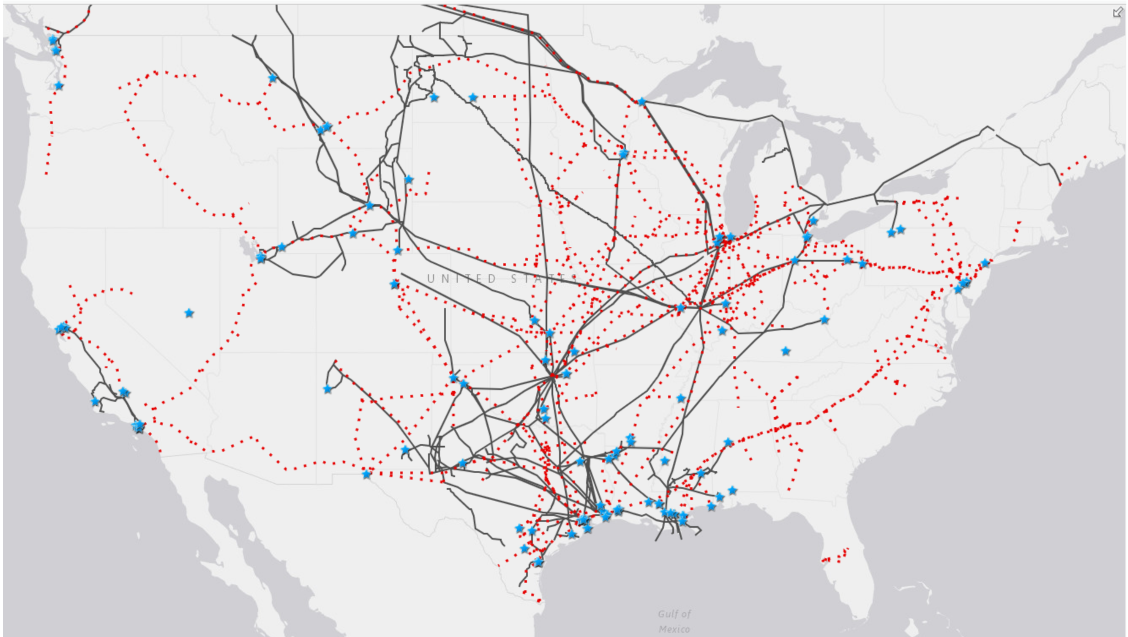

Refined product pipelines are denoted with dashed red lines, and crude oil pipelines are depicted with black lines. Shapefiles for all spatial data come from the EIA state energy mapping system. All location information comes from the EIA state energy map

B.2 Data Appendix

The Energy Information Administration administers detailed surveys at every level of the petroleum industry. The surveys used in this paper are described in Table B.1. The survey forms and additional information,, as well as instructions for requesting access, can be found at http://www.eia.gov/survey/.

Every firm that owns a refinery in the United States reports the total volume and cost of crude oil acquired both domestically and abroad each month, by Petroleum Administration Defense Districts (PADDs), a commonly used geographic aggregation for the industry dating back to World War II mapped in Figure B.1. Every firm reports detailed monthly production data at the refinery level, including the quality of crude oil used, and the exact mix of all end products. This monthly data is supplemented with an annual refinery survey which records exhaustive information about the capacity and technology installed at each refinery at the start of each year. On the sales side, every firm which owns a refinery in the United States reports sales for each state where a transfer of title occurred (regardless of where that product is ultimately consumed). Both the volume sold and the price are reported, by end product, broken out by sales to end users (retail) and sales for resale (wholesale).

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, if a firm owns more than one refinery in a PADD, we only observe a single, aggregated input cost for these facilities. For this reason, we conduct our analysis at the firm-PADD-month level (rather than the refinery level). A breakdown of the number of unique firms by PADD in each year is provided in Table B.2.

A similar aggregation issue exists on the sales side, as sales into a state are reported at the

Table B.1: Description of EIA Surveys

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.

Table B.2: Number of Firms

Year	PADD 1	PADD 2	PADD 3	PADD 4	PADD 5
2004	8	17	21	10	10
2005	7	18	22	11	10
2006	6	17	21	12	11
2007	6	17	21	12	11
2008	6	18	20	12	11
2009	5	20	20	12	10
2010	5	19	20	12	11
2011	5	21	21	13	11
2012	8	15	21	12	11
2013	7	15	21	11	11
2014	7	15	20	12	10
2015	7	16	21	12	11

firm level, not for each refinery. In cases where firms own refineries making the same product in multiple regions in a given month, we proceed by matching each sales state the refinery with the lowest shipping cost to each state. Table B.3 reports the summary statistics for our sample period.

B.2.1 Shipping cost construction

One limitation of the EIA data is that sales in each state are reported by firm, as opposed to the refinery level. Thus, in situations where a firm owns multiple refineries, we do not observe which refinery the firm used to supply a particular state. Here, assume that the firms minimize transportation costs when serving end markets. That is, a firm serves end markets from its refinery with the lowest transportation costs.

To construct transportation costs from each refinery to each end market, we obtained GIS maps of the US refined product pipeline system and waterways suitable for petroleum transportation from EIA, along with GIS coordinates of each refinery. Costs for transporting petroleum products by pipelines, barges and trucks of 2, 4.5 and 30 cents per gallon per thousand miles are taken from estimates presented before the Federal Trade Commission (Jacobs 2002).

C The Impact of the Shale Boom on Refining

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 the late 2000's unlocked billions of barrels of previously uneconomical crude oil reserves. As a result, U.S. oil production nearly doubled during the ensuing decade (Figure 1a).

In this appendix, we illustrate that this surge in domestic oil extraction resulted large, heterogenous reductions in input costs are U.S. refineries. Specifically, we highlight that: (1) U.S. crude oil prices fell relative to the prevailing world price, (2) prices fell more in regions proximate

Table B.3: Summary statistics

	mean	sd
Crude cost (2013 $/gal)	1.860	0.573
Crude - Brent	-0.157	0.218
% Domestic	0.574	0.368
Price Gas	2.362	0.583
Price Distillate	2.492	0.676
Price Total	2.388	0.613
Resale Price Total	2.373	0.612
% Gas	0.531	0.220
% Distillate	0.386	0.200
% Resale	0.851	0.158
N	9215	

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.

C.1 The fracking boom lowered U.S. crude prices relative to world crude prices

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, allowing only a handful of exceptions: (1) export of crude to U.S. territories, (2) export of North Shore crude, and (3) export of California heavy oil. 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 1b 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 $1.40 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 an $11 per barrel reduction on average between 2011 and 2015.

C.2 Initially, fracking primarily benefitted 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, such as North Dakota and Texas, oil production has increased tremendously, while production from conventional resources, such as Alaskan and federal offshore deposits fell approximately 20 percent between 2010 and 2015. Figure 1a, also includes domestic oil production broken out by Petroleum Administration Defense Districts.

PADDs 2 and 3, which contain North Dakota and Texas respectively, both show sharp increases in production around 2010. However, these two areas differ significantly in their pre shale boom conventional production. The Gulf Coast (PADD 3) is home to the most productive on-shore conventional resources. It also contains almost half of U.S. refining capacity, as well as the most active trading hubs. Conversely, PADD 2 had very little conventional production. As a crude oil transportation infrastructure in the upper Midwest was essentially non-existent at the start of the shale boom. This severely limited the ability of oil producers to move product to locations with greater refining capacity.

The result was an unprecedented divergence in crude acquisition costs across refining regions within the United States. In Figure 1c, we plot the average crude acquisition price discount (relative to the Brent spot) at refineries located in each region. 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) fell significantly, relative to the Brent spot price. Over this period, refineries in these regions acquired crude at average costs \$20 to \$30 lower than the Brent Spot price, consistent with production of crude exceeding refinery and transportation capacity in these PADDs.

C.3 Even within region, realized cost reductions varied across refineries

These two divergences, between U.S. and foreign refinery input costs due to the export ban and between Midwest refiners and the rest of the country due to pipeline constraints, are easily observable in the publicly available EIA data. However, using the rich microdata, we are also able to document substantial heterogeneity *within* U.S. refining regions. Figure C.1 presents the distribution of average crude oil acquisition costs by year, 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. The circles in the figure represent the mean difference between the crude acquisition price and the Brent crude oil spot price, by year and PADD and follow a similar pattern to the monthly values in Figure 1c. The average acquisition price in PADDs 2 and 4 fell differentially during the early years of the fracking boom, at a substantial discount relative to the Brent crude spot price. However, the bars in the figure represent the middle 95 percent of the distribution of acquisition costs. Although average prices fell substantially in PADDs 2 and 4, the spread of oil acquisition prices within these regions also increases, with some refiners continuing to pay a premium above global spot prices, despite the glut of oil in their region.

One source of within-region variation comes from the fact the some refineries get their crude from outside of the United States. Figure C.2 plots the average share of crude consumed in each region that is extracted domestically. While refineries near ocean ports were more reliant on imports than others, no region was entirely reliant on or devoid of imports. However, what is remarkable about this graph is how stable these import shares are across time, in light of the massive discounts in domestic crude prices. The one exception is the East Coast, which was initially shut out of domestic crude markets via pipeline, then belatedly obtained access via rail. It was this rail transportation which closed the inter-regional differential in Figure 1c.

The primary explanation for low substitution towards domestic inputs is that crude oils are highly differentiated, and refineries are finely tuned to process a particular type of crude, although some refineries are vertically integrated with international oil companies like Citgo (Venezuela) or Aramco (Saudi Arabia).

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 require either months of

Figure C.1: Average Refinery Crude Price by PADD

Circles represent the mean difference between the crude acquisition price and the Brent crude oil spot price, by year and PADD. The bars represent the range of the middle 95 percent of the distribution, within each year and PADD.

Figure C.2: Domestic Crude Acquisition Shares by PADD

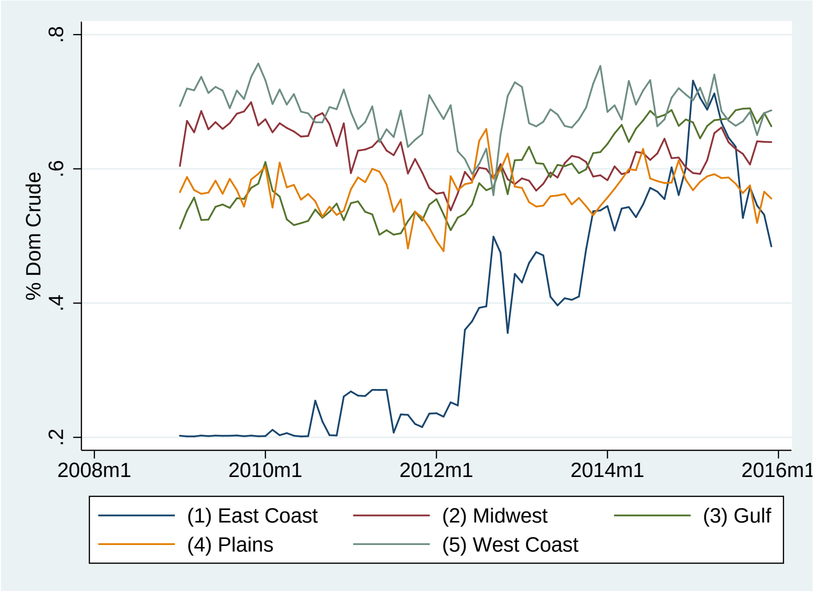

reconfiguration or large capital changes. This is important, because the fracking boom has largely increased domestic supply of "light" (low density) crudes. Figure 1d shows the price domestic oil producers of "light" crude received, less the contemporaneous amount received by producers of "heavy" crude. Historically, lighter crudes traded at a 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.

C.4 Fracking Boom Regressions

We document that location and technological sophistication, which were determined before fracking transformed the U.S. oil market, correlated strongly with the extent to which firms benefited from the fracking boom. To do this, we first split the sample into a "pre" fracking period, prior to 2008, and a "post" fracking period, beginning in 2010. For simplicity we exclude 2008 and 2009 from this "pre"-"post" analysis, although the results are similar with them included.

We calculate "pre" period average crude density (API), domestic crude share, size and technology ("downstream capacity") for each refinery. These pre-period averages are interacted with an indictor for the post-period, as are indicators for which region the refinery is in (excluding PADD 1). Finally, we project monthly crude price paid by each firm onto these interacted variables, along with firm and month-of-sample dummies.

$$CrudePrice_{ft} = 1\{Post_t\}\bar{X}_f^{pre}\beta + \nu_f + \mu_t + \epsilon_{ft} \tag{7}$$

Table C.1 presents the results. Consistent with the exposition above, refineries proximate to shale deposits in the Midwest and Plains states saw large declines in crude costs relative to the omitted group, the East Coast. However, geography alone does not tell the whole story. Refineries processing "light" (high API gravity) crude in the pre-period, or with the technological capability to more easily switch between crudes, experienced larger gains. This regression confirms that pre-period factors like location and the type of crude a refinery is designed for, which were chosen before the shale boom, determined which refineries benefited. This panel variation underpins the identification strategy employed in the paper.

D Supplementary Analysis

D.1 Inputs Costs and the WTI Spot Price

To evaluate whether the average input costs reported by refineries reflect marginal costs, we compare input costs to our closest observable proxy for marginal costs, the West Texas Intermediate Spot Price. If long-run contracts that are not-indexed to spot prices dominate firms' crude acquisition, we would expect a relatively weak link between average input costs and spot prices. In contrast, if firms acquire crude entirely on the spot market or through long-run contracts indexed to the spot price, we would expect high correlation between the two price series - in which case, average crude costs might offer a close approximation to marginal crude costs.

Comparing average input costs and WTI crude costs, we find evidence that the two move in concert - consistent with average crude costs being a relatively good proxy for marginal costs. As a graphical illustration, Figure D.1 plots U.S. average crude costs for all (composite), domestic and international crude streams against WTI crude spot price. More formally, we can empirical test whether changes in average input costs are closely correlated with crude spot prices. Table D.1

Table C.1: Within-PADD Fracking Boom Beneficiaries

API Gravity	-0.00743**
	(0.00294)
% Domestic Crude	-0.101***
	(0.0376)
Log(Capacity)	0.000155
	(0.0123)
Downstream Capacity	-0.0673*
	(0.0359)
(P2) Midwest	-0.116*
	(0.0597)
(P3) Gulf	0.0557
	(0.0632)
(P4) Plains	-0.173***
	(0.0613)
(P5) West Coast	0.0758
	(0.0661)
N	7653
r2	0.958

The dependent variable is the crude price paid ($/gal). All models contain firm-PADD fixed effects and month of sample dummies. The presented coefficients are the average pre-2008 values interacted with an indicator for the post boom period (post 2009). Years 2008 and 2009 are omitted. The sample is restricted to firms with at least 24 months of observations in both periods. Standard errors, clustered at the firm-PADD level, presented in parentheses.

Figure D.1: Average Input Crude Costs and WTI Spot Prices, by Source

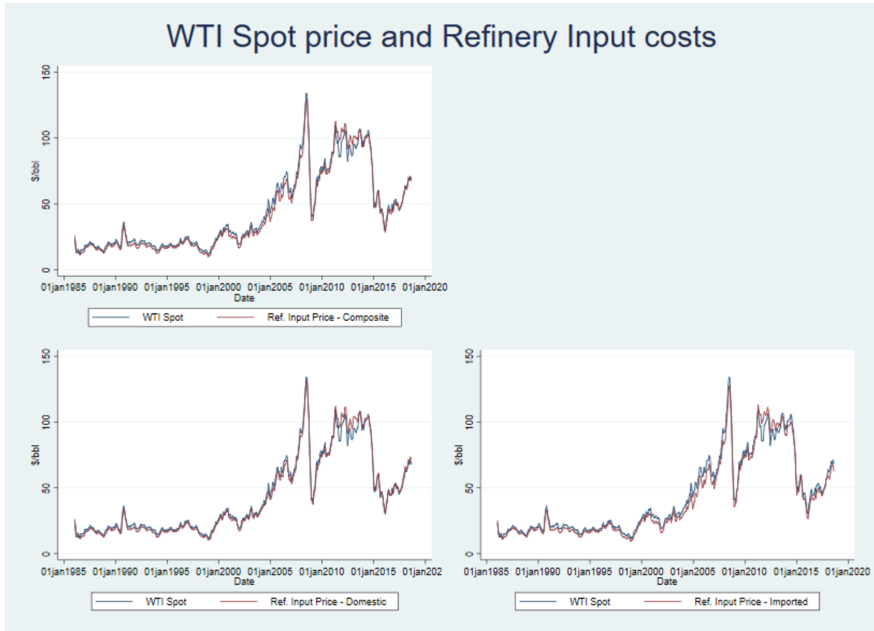

presents the coefficients from regressing first-differenced average input costs against contemporaneous and lagged first-differenced crude spot prices. Consistent with the graphical evidence, we find changes in input costs are correlated with 80 percent of contemporaneous changes in spot prices and effectively 100 percent of changes in spot prices, after including the first lagged month.

D.2 Instrumental Variables

Based on the preceding discussion on refinery heterogeneity (Appendix B) and the fracking boom (Appendix C), we construct four instrumental variables at the firm-month level.

The first two leverage cross-sectional differences in the type of crude oil refineries use. The logic of these instruments is similar to Bartik (1991).

Gravity-based index price (API) - The single biggest driver of price differentials across different crudes is oil density, measured by API gravity. As was discussed in Section B, refineries are configured to process a specific crude density, and deviate little from that in the data.

The EIA reports monthly average delivered crude prices by API gravity, binned into five degree intervals. We calculate the average API for each refinery over the sample, and instrument using the price of the corresponding bin in the EIA data each month.

Index price interacted with downstream technology ($API - Downstream$) - All refineries contain the same basic technology. However, many supplement this with additional "downstream" technology, which provides the refinery a moderate amount of flexibility in transforming a given crude into different end products. For this instrument, we construct the ratio of this downstream technology and distillation capacity (this ratio ranges from 0 to slightly above 2), and then interact this ratio with the API instrument.

In addition to these instruments based on heterogeneity in refinery design, we also construct two instruments based on refinery location. As was discussed in Section 3, due to transportation

Table D.1: Average Input Costs and WTI Spot Prices, by Source

	(1) Composite	(2) Domestic	(3) Imported
$\Delta WTISpot_t$	0.80***	0.76***	0.81***
	(0.012)	(0.013)	(0.019)
$\Delta WTISpot_{t-1}$	0.17***	0.22***	0.15***
	(0.018)	(0.018)	(0.022)
$\Delta WTISpot_{t-2}$	0.0046	0.019	-0.0020
	(0.013)	(0.015)	(0.019)
$\Delta WTISpot_{t-3}$	-0.016	-0.025**	-0.013
	(0.014)	(0.011)	(0.019)
$\Delta WTISpot_{t-4}$	-0.030*	0.010	-0.054***
	(0.015)	(0.016)	(0.020)
Constant	0.012	0.015	0.0033
	(0.038)	(0.039)	(0.052)
Observations	387	387	387
R-Squared	0.96	0.96	0.93

Notes: Dependent variables in Cols 1 - 3 are the average nationwide refinery costs of input from all, domestic and imported sources, respectively.

constraints, the large price declines resulting from the shale boom were concentrated primarily in PADD's 2 and 4. Following a similar logic, we interact time invariant indicators for whether a refinery is in one of these two regions, with time-varying price spread measures.

WTI - Brent spot differential ($PADD2,4-WTI$) - This instrument interacts the US-Global spot price differential with an indicator for being in PADD 2 or 4. As shown in Figure 1b, this difference was small historically, but opened up to unprecedented levels during the shale boom.

US Heavy - Light Crude Spread ($PADD2,4-HL$) - This instrument interacts the US average difference between heavy and light crudes with an indicator for being in PADD 2 or 4. As shown in Figure 1d, this hedonic relationship was reversed during the shale boom, due to the light quality of shale oil.

These own cost instruments are created for every firm-month. We then construct instruments for rival and fringe competitors' costs analogously, by aggregating these own cost measures using the same weighting function employed in each rival and fringe cost calculation.

In sum, this generates eight instruments in the specifications which include own and rival costs as endogenous explanatory variables, and twelve instruments in models which separate out direct rival and fringe firm costs. Tables D.2 and D.3 present the first stage results from columns (3) and (4) in Table 2 of the main text. In general, the coefficients have sensible signs, and magnitudes, particularly comparing across columns. As the API instrument is a price index designed to track each refinery's crude type, a perfect index would have a coefficient close to one. Here the estimated coefficients are all large, but range from 0.65 for own costs, 0.85 for rival costs, and nearly 1 for fringe costs, suggesting measurement error at the individual refinery level. The downstream instrument is actually small in magnitude, and generally not significantly. However, as expected, the two geographic based instruments are large and significant, suggesting refineries in the midwest and plains states capture a large share of deviations in these benchmark differentials.

D.3 Alternative Rival Definitions

In our main specification, rival costs enter as the inverse shipping cost weighted average of competing firms. In this appendix, we consider two alternatives. First, rather than using shipping cost, we also consider simple inverse linear distance between refineries. Those results are presented in columns 1 and 3 of Table D.4.

In some markets, such as electricity, strict capacity constraints and heterogeneous costs impose a natural "dispatch" ordering across firms. In these situations, the market price is set by the cost of the "marginal firm", rather than the average over inframarginal firms serving a market. Such a simple dispatch model does not well describe refinery competition, where capacity constraints are less stark and storage is possible. Nevertheless, we estimate a dispatch-curve model in this setting, replacing the average of rival costs with the contemporaneous cost of the rival which is highest cost on average in the other months during the same year. These results are presented in columns 2 and 4.

D.4 Heterogeneity

The models in the text estimate a single average reduced-form pass-through rate for the entire industry. Structurally, we know that pass-through varies with the relative convexity of residual demand and supply (Weyl and Fabinger, 2013). Mapping these to observed market structure requires assumptions about the nature of competition. Nevertheless, we can check if average pass-through varies with observables likely correlated with these factors.

Table D.5 presents the results from models 1 and 3 from Table 2 with the main pass-through terms interacted with correlates for market structure. In models 1 and 3, pass-through is interacted with an indicator for whether the average HHI for that market during other months in the same year is above the median for the sample. In models 2 and 4, the interaction term is an indicator for whether the firm had more than a 10 percent share of refining capacity in that PADD at the start of the year.

Table D.2: Instrumental Variables First Stage - State Level

	(1)	(2)	(3)	(4)	(5)
Own: API	0.641***	-0.00223	0.629***	-0.00309	-0.00235
	(0.0253)	(0.00550)	(0.0252)	(0.00684)	(0.00589)
Own: API-Downstream	0.0118	0.00340*	0.0124*	0.00667**	-0.00519***
	(0.00731)	(0.00201)	(0.00735)	(0.00264)	(0.00190)
Own: PADD 2,4-HL	0.444***	-0.00583	0.440***	0.00857	0.000231
	(0.0241)	(0.00452)	(0.0239)	(0.00846)	(0.00543)
Own: PADD 2,4-WTI	0.404***	0.00945*	0.405***	-0.0172**	0.00693
	(0.0255)	(0.00508)	(0.0252)	(0.00758)	(0.00473)
Rival: API	0.205***	0.847***	0.175***	0.819***	-0.0348***
	(0.0300)	(0.0108)	(0.0354)	(0.0171)	(0.0102)
Rival: API-Downstream	0.000453	-0.0319***	0.0331***	0.0309***	-0.00624**
	(0.0235)	(0.00810)	(0.0116)	(0.00698)	(0.00314)
Rival: PADD 2,4-HL	-0.187***	0.314***	-0.122**	0.203***	0.0467***
	(0.0527)	(0.00854)	(0.0507)	(0.0206)	(0.00998)
Rival: PADD 2,4-WTI	0.106**	0.489***	0.0430	0.565***	0.0142*
	(0.0464)	(0.0114)	(0.0359)	(0.0167)	(0.00847)
Fringe: API			0.0438	0.0273*	0.958***
			(0.0408)	(0.0151)	(0.0115)
Fringe: API-Downstream			-0.0176	-0.0257***	-0.0770***
			(0.0149)	(0.00473)	(0.00480)
Fringe: PADD 2,4-HL			-0.0692	0.0364	0.215***
			(0.0630)	(0.0223)	(0.0130)
Fringe: PADD 2,4-WTI			0.0293	-0.000599	0.319***
			(0.0537)	(0.0200)	(0.0155)
Rival Measure			Avg	Avg	Avg
Endogenous Var	Own	Rival	Own	Rival	Fringe
N	71570	71570	71529	71529	71529
r2	0.975	0.997	0.975	0.995	0.997

This table presents the first stage results from the state level instrumental variable regressions presented in panel (a) of Table 2. Models (1)-(2) contain the first stage results for the three crude cost variables in the second stage model with month of sample fixed effects (regression 6 from Table 2), and models (3)-(4) contain the first stage results for the second stage model with month and year fixed effects (regression 7 from Table 2). The rows list the excluded variables from each regression, with "Domestic" referring to the domestic crude share instrument, and "API" referring to the API gravity instrument. These are averaged over rival and non-rival firms to match the structure of these variables in the second stage.

Table D.3: Instrumental Variables First Stage - Firm Level

	(1)	(2)	(3)	(4)	(5)
Own: API	0.619***	0.0178	0.597***	0.000388	-0.0306
	(0.0845)	(0.0153)	(0.0876)	(0.0172)	(0.0194)
Own: API-Downstream	0.00224	0.00294	0.00555	0.00658	-0.000777
	(0.0172)	(0.00407)	(0.0164)	(0.00473)	(0.00352)
Own: PADD 2,4-HL	0.406***	0.0257	0.381***	0.0239	0.0250
	(0.0896)	(0.0225)	(0.0974)	(0.0314)	(0.0252)
Own: PADD 2,4-WTI	0.200**	0.00612	0.184**	-0.0361	-0.00335
	(0.0980)	(0.0228)	(0.0755)	(0.0379)	(0.0181)
Rival: API	0.0795	0.841***	0.103	0.862***	-0.112***
	(0.0748)	(0.0300)	(0.163)	(0.0431)	(0.0308)
Rival: API-Downstream	0.0568	-0.0318*	0.132**	0.0692***	-0.0499***
	(0.0694)	(0.0188)	(0.0607)	(0.0208)	(0.0131)
Rival: PADD 2,4-HL	-0.108	0.219***	-0.0817	0.129*	0.0730
	(0.216)	(0.0449)	(0.243)	(0.0662)	(0.0452)
Rival: PADD 2,4-WTI	0.540***	0.500***	0.441***	0.604***	-0.0652*
	(0.196)	(0.0528)	(0.143)	(0.0773)	(0.0380)
Fringe: API			-0.0237	-0.0162	1.111***
			(0.163)	(0.0282)	(0.0322)
Fringe: API-Downstream			-0.0114	-0.00820	-0.0666***
			(0.0362)	(0.00641)	(0.00877)
Fringe: PADD 2,4-HL			0.153	0.156***	0.154***
			(0.195)	(0.0403)	(0.0566)
Fringe: PADD 2,4-WTI			-0.0528	-0.0851***	0.366***
			(0.153)	(0.0320)	(0.0361)
Rival Measure			Avg	Avg	Avg
Endogenous Var	Own	Rival	Own	Rival	Fringe
N	9169	9169	9169	9169	9169
r2	0.969	0.997	0.969	0.996	0.995

This table presents the first stage results from the firm-PADD level instrumental variable regressions presented in panel (b) of Table 2. Models (1)-(2) contain the first stage results for the three crude cost variables in the second stage model with month of sample fixed effects (regression 6 from Table 2), and models (3)-(4) contain the first stage results for the second stage model with month and year fixed effects (regression 7 from Table 2). The rows list the excluded variables from each regression, with “Domestic” referring to the domestic crude share instrument, and “API” referring to the API gravity instrument. These are averaged over rival and non-rival firms to match the structure of these variables in the second stage.

Table D.4: Competition measure results

(a) State Level Results

	(1)	(2)	(3)	(4)
Own	0.0513***	0.0562***	0.0705***	0.0476**
	(0.0131)	(0.0132)	(0.0248)	(0.0237)
Rival	0.0854***	0.0435***	0.00624	0.118***
	(0.0171)	(0.0109)	(0.0313)	(0.0272)
Fringe	0.196***	0.253***	0.146***	0.0820***
	(0.0195)	(0.0159)	(0.0311)	(0.0286)
Brent Spot	0.620***	0.601***	0.718***	0.691***
	(0.0102)	(0.0102)	(0.0107)	(0.0116)
Rival Measure	AvDist	Max	AvDist	Max
IV			Yes	Yes
fstat			3041	1339
N	71529	71517	71529	71517

(b) Firm Level Results

	(1)	(2)	(3)	(4)
Own	0.0508*	0.0782***	-0.0252	0.00687
	(0.0298)	(0.0269)	(0.0707)	(0.0435)
Rival	0.171***	0.105***	0.165	0.247***
	(0.0369)	(0.0238)	(0.106)	(0.0710)
Fringe	0.0907**	0.150***	0.0479	-0.0256
	(0.0351)	(0.0321)	(0.0599)	(0.0559)
Brent Spot	0.642***	0.621***	0.752***	0.707***
	(0.0159)	(0.0185)	(0.0231)	(0.0273)
Rival Measure	AvDist	Max	AvDist	Max
IV			Yes	Yes
First-stage F			123	167
N	9169	9169	9169	9169

This table presents the results of estimating Equation (3) 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-PADD-month level and includes firm-PADD fixed effects. All specifications include year fixed effect and month fixed effects. 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).

Table D.5: Competition measure results

(a) State Level Results

	(1)	(2)	(3)	(4)
Own	0.0445**	0.0373**	0.0478	0.0348
	(0.0180)	(0.0171)	(0.0320)	(0.0369)
Own X Int	0.000282	0.0286	0.0273	0.0549
	(0.0211)	(0.0289)	(0.0382)	(0.0600)
Rival	0.282***	0.295***	0.157***	0.181***
	(0.0200)	(0.0214)	(0.0354)	(0.0452)
Rival X Int	-0.00213	-0.0418	-0.0299	-0.0685
	(0.0213)	(0.0307)	(0.0384)	(0.0618)
Brent Spot	0.626***	0.625***	0.734***	0.729***
	(0.0112)	(0.0115)	(0.0124)	(0.0134)
Interaction	HHI	Cap.PADD	HHI	Cap.PADD
Time FE	Y,M	Y,M	Y,M	Y,M
IV			Yes	Yes
fstat			2329	1551
N	71570	71570	71570	71570

(b) Firm Level Results

	(1)	(2)	(3)	(4)
Own	0.0486	0.0368	-0.0433	-0.0822
	(0.0400)	(0.0369)	(0.0624)	(0.0677)
Own X Int	0.00100	0.0795	0.136	0.243***
	(0.0406)	(0.0572)	(0.0860)	(0.0825)
Rival	0.285***	0.306***	0.238***	0.309***
	(0.0444)	(0.0402)	(0.0741)	(0.0810)
Rival X Int	-0.00189	-0.105*	-0.137	-0.268***
	(0.0386)	(0.0559)	(0.0845)	(0.0821)
Brent Spot	0.622***	0.619***	0.747***	0.722***
	(0.0176)	(0.0180)	(0.0242)	(0.0246)
Interaction	HHI	Cap.PADD	HHI	Cap.PADD
Time FE	Y,M	Y,M	Y,M	Y,M
IV			Yes	Yes
fstat			120	114
N	9169	9169	9169	9169

This table presents the results of estimating Equation (3) 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-PADD-month level and includes firm-PADD fixed effects. All specifications include year fixed effect and month fixed effects. "X Int" variables reflect the cost measure listed, interacted with a binary indictor based on the measure listed in the Interaction row. Standard errors are presented in parentheses, clustered at the firm-state level. All models include demand shifters and supply shifters.

D.5 Carbon Tax

Figure D.2: GHG heterogeneity

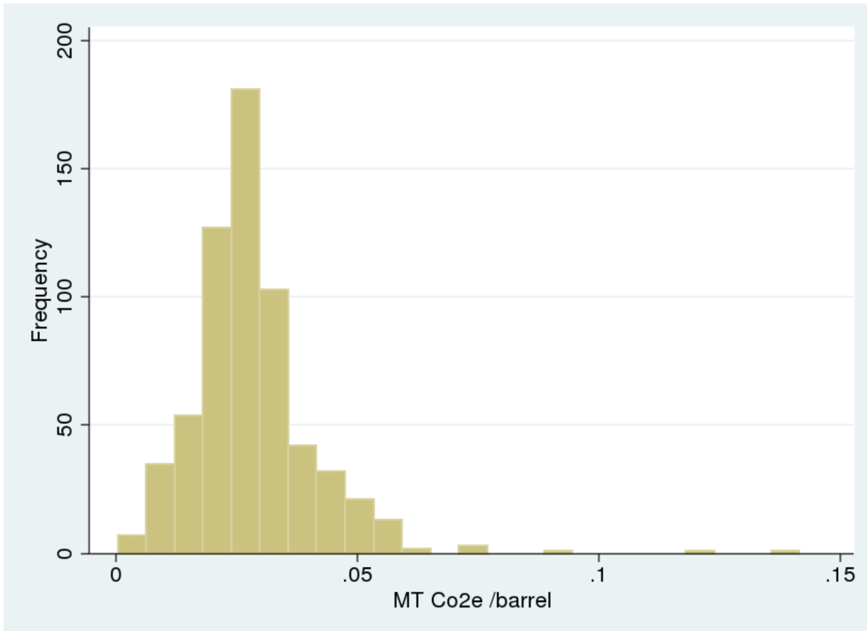

Based on annual data (2011-2015) in EPA GHGRP.

Table D.6: Determinants of CO2 Heterogeneity

	(1)
API Gravity	0.000174***
	(0.0000613)
log(Capacity)	0.0000794
	(0.000473)
% Coking	0.0238***
	(0.00373)
% Cracking	0.0201***
	(0.00203)
Constant	0.0220***
	(0.00574)
mean(Y)	.027
TimeFes	Y
N	615
r2	0.347

Dependent variable: Metric tons of CO2 equivalent per barrel of inputs processed. Data from EPA GHGRP (2011-2015). Model includes PADD and year dummies.