

NBER WORKING PAPER SERIES

ENERGY COST PASS-THROUGH IN U.S. MANUFACTURING:
ESTIMATES AND IMPLICATIONS FOR CARBON TAXES

Sharat Ganapati
Joseph S. Shapiro
Reed Walker

Working Paper 22281
<http://www.nber.org/papers/w22281>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
May 2016, Revised April 2019

Previously circulated as “The Incidence of Carbon Taxes in U.S. Manufacturing: Lessons from Energy Cost Pass-Through.” We thank Joe Altonji, Antonio Bento, Ernesto dal Bó, Severin Borenstein, Lucas Davis, David Donaldson, Penny Goldberg, Kostas Metaxoglou, Nate Miller, Jim Poterba, Jim Sallee, Glen Weyl, Danny Yagan and various seminar participants for useful comments and discussions. Funding from the Department of Energy, the National Science Foundation (SES-1850790), the Sloan Foundation, and the NBER is gratefully acknowledged. We would also like to thank Randy Becker, Cheryl Grim, and Kirk White for sharing code and data. Jonathan Kadish and Carla Johnston provided excellent research assistance. The research in this paper was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau at the Berkeley and Yale Census Research Data Center. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau or the National Bureau of Economic Research. This paper has been screened to insure that no confidential data are revealed.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Sharat Ganapati, Joseph S. Shapiro, and Reed Walker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Energy Cost Pass-Through in U.S. Manufacturing: Estimates and Implications for Carbon Taxes

Sharat Ganapati, Joseph S. Shapiro, and Reed Walker

NBER Working Paper No. 22281

May 2016, Revised April 2019

JEL No. H22,H23,L11,Q40,Q54

ABSTRACT

This paper studies how changes in energy input costs for U.S. manufacturers affect the relative welfare of manufacturing producers and consumers (i.e., incidence). In doing so, we develop a novel partial equilibrium methodology designed to estimate the incidence of input taxes. This method simultaneously accounts for three determinants of incidence that are typically studied in isolation: incomplete pass-through of input costs, differences in industry competitiveness, and substitution amongst inputs used for production. We apply this methodology to a set of U.S. manufacturing industries for which we observe plant-level unit prices and input choices. We find that about 70 percent of energy price-driven changes in input costs are passed through to consumers in the short- to medium-run. We combine industry-specific pass-through rates with estimates of industry competitiveness to show that the share of welfare cost borne by consumers is 25-75 percent smaller (and the share borne by producers is correspondingly larger) than models featuring complete pass-through and perfect competition would suggest.

Sharat Ganapati
Georgetown University
1301 33rd Street NW
UNIT 3
Washington, DC 20057
United States
sganapati@gmail.com

Reed Walker
Haas School of Business
University of California, Berkeley
2220 Piedmont Ave
Berkeley, CA 94720
and NBER
rwalker@haas.berkeley.edu

Joseph S. Shapiro
University of California at Berkeley
714 University Hall #3310
Berkeley, CA 94720
and NBER
joseph.shapiro@berkeley.edu

1 Introduction

Greenhouse gas emissions are classic environmental externalities that have not faced stringent U.S. or global policy, even in the face of mounting evidence that the economic costs from climate change could be severe (Carleton and Hsiang, 2016). Existing and proposed policies designed to address future climate change implicitly or explicitly raise the price on carbon dioxide emissions. Pricing carbon emissions, by design, will make fossil-fuel based energy consumption more expensive. This has led to the dual concern by policymakers that increased energy costs will not only make industries that rely on these energy inputs less competitive but also make consumers of industry products worse off due to higher prices. Political debates about the extent to which greenhouse gas mitigation policies would hurt consumers or firms have made these policies difficult to enact, despite widespread consensus among economists that a carbon tax would be efficient.¹ Despite these concerns, relatively little is known about how changes in energy input costs would impact the prices that consumers and producers face, or, more generally, consumer and producer welfare.²

This paper uses administrative data from the Census Bureau’s Census of Manufactures to estimate the degree to which energy price-driven changes in production costs impact consumers relative to producers (i.e., incidence) for a select set of U.S. industries. In so doing, this paper offers a novel partial equilibrium approach to analyzing the incidence of changes in input costs that accounts for three important issues that the existing public finance literature considers in isolation or not at all: substitution among productive inputs, incomplete pass-through of input costs, and various forms of imperfect competition. Formally, we generalize recent theories of incidence to derive a sufficient statistic representation for the incidence of producer input taxes that depends on two parameters that we estimate: a pass-through rate and a measure of industry competitiveness. The goal of this paper is to use these statistics to characterize industry-level changes in consumer relative to producer surplus associated with changes in energy input costs.

The framework we develop has several important features. While the study of tax incidence has played a central role in public finance since at least Jenkin (1872), existing theory focuses on the incidence of taxes on a firm’s outputs (Weyl and Fabinger, 2013). As a result, most partial equilibrium empirical studies on tax incidence to date have had to make tradeoffs in order to study changes in input costs, such as assuming that policymakers only tax outputs but not inputs, or that

¹For example, the National Association of Manufacturers lamented that a U.S. cap-and-trade program for CO₂ would have a “devastating impact to manufacturing” due in part to inability “to adjust the price of their goods and services quickly enough to match potentially steep energy cost increases” (Streeter, 2009). Similarly, the Heritage Foundation review of climate regulation argues, “Americans will be hit repeatedly with higher prices as businesses pass higher costs onto consumers.” They go on to say, “If a company had to absorb the costs, high energy costs would shrink profit margins” (Loris and Jolevski, 2014).

²The lack of empirical evidence has often been noted in the literature. Bento (2013) comments, “More research that carefully quantifies the effects of environmental policies on the prices of final goods is also needed. Existing research typically assumes that firms will have the ability to fully pass along the costs of environmental policies in the form of higher consumer prices. Under many circumstances, this assumption may not be applicable.” Similarly, Parry, Sigman, Walls, and Williams (2006) highlight. “Empirical studies on the extent to which the costs of environmental policies are passed forward into higher prices of consumer products would be extremely valuable; at present, empirical analyses typically assume 100% pass-through. . .”

firms use one input to production, or that firms cannot substitute between inputs. By generalizing standard incidence formulas to explicitly examine input taxes, our approach removes the need for such assumptions. The methodology explicitly nests output taxes as a special case where all inputs are taxed, and we present several analytically equivalent versions of our incidence formula, each of which have different data requirements. While we focus on the incidence of energy prices, one of our goals is to illustrate how this framework could be used to analyze the product market incidence of a wide array of policies that change input prices, such as studying how minimum wage laws or employee health insurance mandates affect the surplus of firms versus their consumers.

The paper has two main empirical findings. First, for the several manufacturing industries we study, we estimate a pass-through rate of around 0.70; that is, on average a 1 dollar increase in marginal costs due to higher energy prices translates into a 70 cent increase in unit output prices. Second, many papers assume that consumers bear 100 percent of the welfare burden of shocks to industrial energy prices. Our estimates for these six industries imply instead that consumers bear only 25-75 percent of this burden; firms bear the rest. Most of our estimates reject both the common assumptions of complete pass through and perfect competition.³

To help frame ideas, consider a simple example of a manufacturing plant that faces increases in carbon taxes. Economists have long recognized that the incidence of a tax, defined in this paper as its impact on the welfare of producers relative to consumers, is independent of who physically writes the check to pay it. The same principle applies for shocks to production costs (e.g., oil supply disruptions arising from political events in oil-producing countries, or fracking). The reason is that changes in production costs due to taxes or other market forces lead to changes in prices and quantities of both outputs and inputs. In this example, a carbon tax levied at the level of a manufacturing plant only physically applies to the plant since the government directly collects tax revenue from the plant. If the tax causes plants to increase prices, then the tax burden will shift forward to consumers. If the tax causes plants to invest in energy efficient production technologies, then producers will have to pay less in energy taxes, minimizing their burden. Thus, pass-through and input substitution describe means by which the party which physically pays a tax can either transmit the effects of that tax to others or avoid paying the tax altogether.

The paper proceeds in three steps. First, we formalize an expression for the incidence of changes in the costs of a specific input like energy. The partial equilibrium focus delivers empirical tractability while abstracting away from important general equilibrium phenomena associated with the production supply chain or vertical integration more generally. The key parameter determining incidence is the degree to which marginal costs are passed through to output prices.

The paper's second part focuses on six U.S. manufacturing industries for which we observe both

³Most existing work on the incidence of energy costs assumes that the industry supply curve is infinitely elastic and that producers are perfectly competitive, or equivalently, they assume that pass through is complete and that consumers bear the entire welfare burden. Ultimately all welfare changes affect households, either through ownership of firms (producer incidence) or through expenditure patterns (consumer incidence). The question of who owns firms is beyond the scope of this paper. However, it is well established that capital ownership is not evenly distributed across the U.S. population, so welfare losses to manufacturing producers and consumers have different incidence.

unit prices and input costs: boxes, bread, cement, concrete, gasoline refining, and plywood.⁴ We compute plant-level marginal costs using methods originally proposed by Hall (1986) and further developed by De Loecker and Warzynski (2012). In practice, this amounts to estimating time-varying, plant-level markups using production function estimates and backing out marginal costs as the difference between the price and the markup.

The paper’s third part estimates how energy price-driven changes in marginal costs are passed through to output prices.⁵ Marginal cost pass-through is an important ingredient for any analysis of input cost incidence, yet has received relatively little attention in the empirical literature. We focus on energy-price induced variation in marginal costs both to address possible endogeneity issues stemming from measurement error and to deliver a policy-relevant, local average treatment effect for energy-price induced, marginal cost pass-through. We generate proxies for changes in plant-level energy costs that we use to instrument our measure of plant-level marginal costs. Both of these instruments interact the shares of different fuels used for energy with time-series variation in the prices of these fuels (a “shift-share” approach related to Bartik (1991)). The first instrument relies on the fact that electricity prices vary over time and space depending on the local fuels used for generation in a region. We calculate the share of a state’s electricity generated by coal, petroleum, and natural gas, and we interact these (lagged) generation shares with national time-series variation in the prices of these fuels. For example, when coal prices rise, industrial electricity prices disproportionately rise in areas where coal-fired power plants generate a large share of electricity.⁶ The second instrument for plant-level marginal costs uses the fact that manufacturing industries rely on a range of different energy inputs into the production process. Thus, industries whose production process mainly uses coal will see energy costs increase more when coal prices rise, and industries that rely more on natural gas will see energy costs increase more when the price of natural gas rises. We interact (lagged) industry fuel input shares with national time-series variation in the prices of these fuels. Both instruments predict changes in marginal costs, and in a majority of estimates (though not all) these instruments provide a strong first-stage. Using shocks to energy costs as instruments for marginal cost delivers a local average treatment effect for how energy price-induced

⁴The Census Bureau primarily collects quantity data for industries in which quantity data is readily interpretable, such as those industries that produce homogeneous products. A few additional industries produce homogeneous products and have Census-collected price/quantity data. However, much of the quantity data in these industries has been imputed, and the remaining, non-imputed sample sizes preclude the estimation of costs. Outside of industries that produce homogeneous products, plant-level output price data is less readily available at the plant-level, and in cases where the data is available (e.g., BLS producer price data), it is often not possible to observe or estimate changes in input costs. Plant level linkages between the BLS producer price survey and plant-level production surveys such as the Census of Manufacturers should be a priority for future research.

⁵The theory and the empirical setting are designed to explore the incidence of changes in energy input costs as a lens through which to understand the incidence of a carbon tax. Recent research suggests that incidence of taxes can differ from incidences of market based prices changes due to differences in consumer salience (Li, Linn, and Muehlegger, 2014) or due to various forms of tax evasion (Marion and Muehlegger, 2008; Kopczuk, Marion, Muehlegger, and Slemrod, 2016a).

⁶Coal prices likely affected electricity prices more in the 1972-1997 years we study than they do today. In the years we study, essentially all states had regulated electricity markets in which electricity prices depended on the average and not the marginal cost of generating electricity. By contrast, after the period we study (i.e., since the late 1990s), many states have operated deregulated electricity markets in which marginal units (which more often run on gas than coal) rather than average costs determine prices.

changes in marginal cost are passed through to product prices. The paper concludes by combining information on cost pass-through with information on industry competitiveness to characterize the economic costs of energy price increases for consumers relative to producers, separately by industry.

Some questions of external validity are worth clarifying. We study six homogeneous industries, over the period 1972-1997, using variation in fuel prices rather than policy. This warrants caution in taking exact numbers from our regressions to specific policies today. Nonetheless, our general finding that no industries or estimates exhibit complete pass-through or perfect competition suggests it is worth carefully considering these common assumptions. Second, one of our research designs exploits variation in fuel prices across U.S. regions; how relevant is this analysis to national carbon taxes? This research design may be most relevant to regional carbon pricing schemes, such as the Northeast's Regional Greenhouse Gas Initiative and California's AB32 cap-and-trade market. We also obtain similar results from a second research design that exploits differences in energy prices across industries, which may have more external validity to national changes in energy prices. Third, our use of panel data and regressions implies that we identify short-to-medium run variation in pass-through rates; cross-sectional regressions would identify a longer-run pass-through rate, but face stronger identifying assumptions. Short-and medium-run estimates may be disproportionately relevant for the political economy of policies affecting energy prices.

It is also worth clarifying the applicability of our approach. We study homogeneous products; is this approach relevant to differentiated products? We study these products primarily because the census collects output quantity data for them; the methods we describe are relevant to differentiated products (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016). How applicable is our approach given its high data demands? The required data are available in other important settings.⁷ We also outline variants of the approach which are applicable in different data environments. Finally, what is the relationship of our industry-by-industry analysis to economy-wide effects of a national carbon tax? There is often a gap between economy-wide general equilibrium analyses, which make assumptions like perfect competition in part for tractability, and industry-specific partial equilibrium analyses like . We see this paper as one step towards bridging that gap.

A last clarification involves welfare interpretation. Several markets we study mainly provide intermediate goods; can we measure consumer surplus from such markets? A literature shows that under strong assumptions, the change in surplus due to taxing intermediate inputs is identical whether surplus is estimated in the market for intermediate or final goods (Schmalensee, 1976).

This paper builds on several literatures. Researchers in public finance have used pass-through to understand partial equilibrium tax incidence for decades (see e.g., Poterba (1996)).⁸ Typically, this

⁷For example, many countries collect firm-level price microdata for constructing Producer Price Indices. Some surveys such as India's Prowess data and Columbia's Annual Manufacturing Survey report on both quantity and price of outputs.

⁸Some general equilibrium work studies taxes on capital but not labor (e.g., Harberger (1962)). An important literature in public finance studies the economy-wide incidence of carbon taxes and energy prices, typically by using input-output matrices and detailed expenditure data (e.g., Grainger and Kolstad (2010)). However, this literature typically assumes that industries are perfectly competitive, pass-through is complete, and firms cannot adjust input demands in response to changes in input costs. In this paper, we try to relax these assumptions in a partial equilibrium, empirical setting. Some computable general equilibrium models allow for dynamic adjustment costs, which provide for

work focuses on estimating pass-through of sales taxes or cost pass-through of one input in specific industries.⁹ Researchers typically combine these pass-through estimates with an assumption about market structure (e.g., perfect competition) to characterize incidence. When industries use a single input and when the market is perfectly competitive, input cost pass-through is a sufficient statistic for the incidence of input costs. However, when firms can substitute across different inputs used for production or when the market is not perfectly competitive, pass-through alone is no longer sufficient. For example, Miller, Osborne, and Sheu (2015) analyze the pass-through of energy costs in a single manufacturing industry – cement – though assume a single factor of production, which precludes factor substitution. One motivation for our paper is that an essential feature of energy taxes, tariff reforms for intermediate inputs, minimum wage laws, or other regulation of inputs is that firms can change their use of the regulated input. This changes the welfare burden, formulas for measuring incidence, and methods required to estimate the relevant parameters. Another motivation for this paper is that, conditional on a given pass-through rate, incidence varies with industry competitiveness (Katz and Rosen, 1985; Weyl and Fabinger, 2013); many of the most concentrated industries in the United States are also industries that produce large amounts of CO₂ emissions (e.g., steel, petroleum refining, electricity, and cement).

A related literature highlights that welfare and incidence of environmental and energy policy may differ dramatically for imperfectly competitive industries (Buchanan, 1969; Barnett, 1980). The nascent empirical work here typically studies a single industry with detailed data or information on industry costs, consumer preferences, and market structure (e.g., cement or electricity). With detailed data and industry-specific insight, researchers can estimate both supply and demand and then explore welfare and incidence through a variety of policy counterfactuals (Ryan, 2012; Fowlie, Reguant, and Ryan, 2016). However, the data requirements and modeling choices in these studies are typically industry specific, and it is not clear how the methodologies or results generalize into other sectors. Our finding that cement is quite different from the other industries we study (in both cost pass-through and competitiveness) underscores that these findings may be hard to generalize to other industries. In contrast to this existing literature, we develop an empirical methodology to estimate the incidence of input costs for any industry in which we observe price data and input choices. The approach is flexible enough to allow for changes in energy prices that lead to factor substitution in production, and the approach is general enough to calculate incidence in the presence of incomplete pass-through or deviations from perfect competition.

This paper is also related to a set of work analyzing how energy prices differentially affect firms and industries. The literature on industry impacts of carbon policies is based largely on simulation modeling, although a number of statistical analyses also exist (e.g., Kahn and Mansur

incomplete pass-through in the short run (Goulder and Hafstead, 2017), though many computable general equilibrium models assume complete pass-through.

⁹See e.g., the pass-through of gasoline taxes or ethanol subsidies into retail gasoline or diesel prices (Doyle and Samphantharak, 2008; Marion and Muehlegger, 2011; Kopczuk, Marion, Muehlegger, and Slemrod, 2016b; Lade and Bushnell, 2016), European Union Emissions Trading System (ETS) allowance prices to electricity wholesale prices (Fabra and Reguant, 2014), natural gas prices to nitrogenous fertilizer (Bushnell and Humber, 2017), or raw coffee bean prices to retail coffee (Nakamura and Zerom, 2010; Bonnet, Dubois, Boas, and Klapper, 2013).

(2012), Aldy and Pizer (2015), and Martin, de Preux, and Wagner (2014)). The simulation analyses include both short-term partial equilibrium assessments as well as long-term computable general equilibrium (CGE) models.¹⁰ Davis, Grim, Haltiwanger, and Streitwieser (2013) document patterns in plant-level electricity input prices for U.S. manufacturing using some of the same data we use, though their focus is on explaining what determines patterns of electricity input prices, rather than assessing their consequences. Relative to existing empirical literature in this area, we examine unexplored outcomes such as how energy prices affect plant-level marginal costs and output prices for a select group of industries.

Lastly, this project relates to a substantial empirical literature on pass-through of costs other than energy prices. This literature spans many fields and explores the pass-through of exchange rates (Goldberg and Hellerstein, 2008; Gopinath, Gourinchas, Hsieh, and Li, 2011; Campa and Goldberg, 2005), taxes, subsidies, and permit fees (Poterba, 1996; Marion and Muehlegger, 2011; Lade and Bushnell, 2016; Stolper, 2017), healthcare capitation payments (Cabral, Geruso, and Mahoney, 2015; Duggan, Starc, and Vabson, 2014), major commodity inputs (Nakamura and Zerom, 2010; Bonnet, Dubois, Boas, and Klapper, 2013) and minimum wage laws (Harasztosi and Lindner, 2015). Few papers observe or estimate marginal cost pass-through, with recent exceptions including a paper focused on Indian manufacturing (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016), another constructing marginal costs from fuel data in the Spanish electricity market (Fabra and Reguant, 2014), and a recent paper exploring pass-through of cost shocks for U.S. oil refining (Muehlegger and Sweeney, 2017). Our main innovation here is estimating incidence and deriving a formula to do so for changes in input prices. Many of these studies analyze the pass-through of one important component of a firm’s costs, but to the best of our knowledge, no literature estimates the pass-through of a firm’s full marginal cost in U.S. manufacturing. This parameter is a critical ingredient for characterizing the incidence of input taxes levied on U.S. manufacturers.

The rest of the paper proceeds as follows. Section 2 describes the role of energy in the output of U.S. manufacturing products. Section 3 describes a general theory of incidence that motivates our empirical analysis. Section 4 describes the data, and Section 5 describes the econometric setting. Section 6 presents results, and Section 7 concludes.

2 Energy and U.S. Manufacturing

A brief background on energy use in U.S. manufacturing may clarify this paper’s analysis.¹¹ Manufacturing accounts for a large share of energy demand. Industrial energy consumption (which includes manufacturing along with agriculture, mining, and construction) accounts for about 30 percent of U.S. end-use energy consumption and also about 30 percent of end-use greenhouse gas emissions (EIA, 2015; USEPA, 2015). Energy is a limited direct cost for manufacturing on average, at about two percent of revenues for the entire manufacturing sector, though energy costs are

¹⁰Ho, Morgenstern, and Shih (2008) review more than a dozen prior U.S. and European analyses using CGE models.

¹¹Except where otherwise noted, this section describes data for the year 2010.

much greater in some industries (Becker, Gray, and Marvakov, 2013). In alkali and chlorine manufacturing, cement, gasoline refining, lime manufacturing, and primary aluminum production, for example, energy costs (including energy that is physically formed into the manufactured product, or “feedstock”) exceed 20 percent of revenues.

Manufacturing generally uses two categories of energy—electricity and primary fuels. Electricity’s price per British thermal unit (BTU) is two to five times the mean price of other energy sources (EIA, 2010), partly because much of the fuel used to produce electricity is lost as heat. The main primary fuels used in manufacturing are oil, natural gas, and coal. About 75 percent of BTUs used for fuel in manufacturing come from natural gas, 20 percent from coal, and the rest from oil and assorted sources (EIA, 2010). Natural gas is increasingly common, partly because hydraulic fracturing decreased the domestic price of natural gas beginning around 2008 (see e.g., Hausman and Kellogg (2015)), though even in 1990, natural gas provided 75 percent of BTUs used for fuel.

Manufacturing plants use energy for four general tasks: boiler fuels (about 25 percent of BTUs), process production (40 percent), other on-site purposes (10 percent), and feedstock (25 percent). Boiler fuels mainly come from natural gas and coal and are used for combined heat-and power, cogeneration, or related purposes. Process production includes heating or cooling parts of the manufacturing product itself, driving manufacturing machines, or electro-chemical processes. Driving machines almost exclusively use electricity, but other process production tasks use a mix of natural gas, oil, and coal. Other on-site uses of energy include plant lighting, heating, cooling, ventilation, and on-site transportation. Industries differ in their use of fuels based on the prevailing production processes and regional availability of fuel inputs.

Energy is costly to store. Batteries are expensive enough that mass storage of electricity is economically infeasible, and most electricity is consumed at the instant it is generated. Most manufacturing plants obtain natural gas from distribution pipelines and do not store it on-site. Oil and coal can be stored, though their weight and bulk mean they are stored in limited quantities. In addition, industries differ in their use of different fuels based on the prevailing production processes and regional availability of fuel inputs.

Different energy sources also have different spatial market structures. Electricity prices vary over time and space depending on fuels, efficiency, and scheduling of electricity generating units. In the years we study (pre-1998), which precede deregulation, an electric utility had a monopoly over customers in its service territory, so the electricity prices an industrial customer faced depended heavily on the fuels used by the electric utility serving it. Utilities supply most electricity used in manufacturing, though additional electricity comes from non-utility generators (e.g., merchant plants) and from on-site generation.¹² Crude oil is traded on a global market. Manufacturing plants generally buy distillate or residual fuel oil, which is processed by refineries. Lack of spatial integration in refinery markets introduces additional spatial variation in prices of these petroleum products that is mostly driven by idiosyncratic changes in local supply and demand. Natural gas

¹²Our empirical analysis counts on-site electricity generation as primary fuel consumption since in this case the data record the plant buying fuels rather than the plant buying electricity.

is transported by pipeline within the U.S. from producing to consuming regions. In the years we study, much natural gas was extracted in Texas and Louisiana, and natural gas prices increase with distance from those areas due to pipeline transportation costs. The price at the location where an interstate pipeline reaches a population center differs from the price that industrial plants pay for natural gas due to local distribution costs and to distributor markups. Coal is more costly than other fuels to transport, so coal prices vary more over space due to local market conditions.

In sum, energy prices vary substantially and are an important input into many manufactured goods. The goal of this paper is to better understand how this temporal and spatial variation in energy prices affect both manufacturing producers and their consumers.

3 Theory of Incidence

There exists a body of research at the intersection of public finance and industrial organization which describes how incidence might vary with industry competitiveness (Katz and Rosen, 1985; Stern, 1987; Weyl and Fabinger, 2013). We believe this literature is useful because some of the most energy intensive manufacturing industries, are also industries characterized by very large fixed costs or transportation costs which may lead to varying degrees of market power. (e.g., cement and oil refining). The existing literature is primarily concerned with the incidence of output taxes, and we extend this literature to analyze changes in input taxes (or input costs).¹³ The difference is that firms must pay an output tax on every dollar of revenue but must only pay an input tax on each unit of the input purchased. Firms can substitute away from the taxed input.

We first describe incidence assuming firms are either perfectly competitive or have a monopoly. These polar cases illustrate the basic intuition for how incidence differs with market power. We then describe a general setting where firms are characterized by arbitrary forms of competition.

A few assumptions guide the analysis. The following discussion and subsequent analysis is partial equilibrium. We consider taxes on variable input costs, not fixed costs. We make the assumption, consistent with most of the literature, that all goods outside the focal industry (including markets for the taxed input) are supplied perfectly competitively, and thus that the welfare of producers arising from consumer substitution to these goods may be ignored.¹⁴ Additionally, we assume the taxed input has perfectly elastic supply.

We begin with key definitions. Let I denote the incidence of a marginal increase in the tax rate τ , defined as the ratio of its effects on consumer and producer surplus (CS and PS):

$$I \equiv \frac{dCS/d\tau}{dPS/d\tau}$$

Incidence above one implies that consumers bear a majority of the welfare loss, while incidence

¹³We abstract from the use of potential tax revenue.

¹⁴Relaxing this assumption may be possible but would require estimates of cross-price elasticities across industries, which are difficult to estimate. Future work could combine the insights from Goulder and Williams III (2003) to derive empirically implementable formulas for incidence in the presence of pre-existing distortions in other markets.

below one implies that producers do. Let ρ denote the pass-through rate of a tax, defined as the marginal change in the level of output prices P due to a change in input tax rates:

$$\rho \equiv \frac{dP}{d\tau}$$

Let γ denote the cost-shift rate, defined as the marginal effect of the input tax rate τ on marginal costs: $\gamma \equiv dMC/d\tau$. The cost-shift rate γ can be less than or greater than one. Finally, let $dAVC/d\tau$ denote the change in average variable costs (equal to variable costs divided by total output) due to a marginal increase in the tax rate.

We now turn to describe the key incidence results. Perfect competition provides a useful baseline since its results are simple and intuitive. In perfectly competitive markets with input taxation, the pass-through rate and the cost shift rate fully characterize tax incidence. The incidence simply equals the pass-through rate divided by the cost shift rate minus the pass-through rate:

$$\begin{aligned} I^{Competitive} &= \frac{\rho}{dAVC/d\tau - \rho} \\ &= \frac{\rho}{\gamma - \rho} \end{aligned}$$

The second equality follows under the assumption $AVC = MC$.¹⁵

This result has an intuitive basis in the envelope theorem. A marginal increase in a tax decreases consumer surplus by the equilibrium quantity, Q^* , times the change in consumer prices, ρ . A marginal increase in a tax decreases producer surplus by Q^* times the change in producer prices relative to average variable costs, $dAVC/d\tau - \rho$. The incidence of a tax in a competitive market equals the ratio of these two terms, $\rho/(dAVC/d\tau - \rho)$.¹⁶ Describing tax incidence in a perfectly competitive market only requires knowing the pass-through rate and the effect of tax on average variable costs (if $MC = AVC$, the latter is replaced with the cost-shift rate), making pass-through and cost-shifting parameters “sufficient statistics” for incidence of changes in input costs.

Figure 1 illustrates incidence under perfect competition. Panel A shows a shift in average variable costs, $dAVC/d\tau$, due to an input tax, and Panel B shows the levels of consumer and producer surpluses in the new equilibrium. As visualized, a small change in average variable costs causes consumer surplus to decrease by the change in prices, ρ , times the output quantity Q^* . Producers now receive an additional ρ per unit sold, however this is offset by the change in

¹⁵Our empirical framework is designed to estimate marginal cost pass-through. In order to map theory to our empirical results, we assume that average variable costs equal marginal costs: $AVC = MC$. Some production technologies are sufficient to ensure this assumption; examples include constant elasticity of substitution (CES) production functions or technologies with constant marginal costs (e.g., Melitz (2003)). For others like Cobb-Douglas or the translog function that we use, constant returns to scale is a sufficient condition for $AVC = MC$. In general, the theory allows a specific application to work with either $dAVC/d\tau$ or the combination of $dMC/d\tau = \gamma$ and the assumption $AVC = MC$, depending on data availability in a specific setting.

¹⁶We describe the incidence of infinitesimal changes in tax rates. Characterizing the incidence of discrete changes in tax rates requires integrating over changes in consumer and producer surplus from the initial tax rate to the new tax rate. The incidence of a discrete change in tax rates then depends on the average pass-through rate between the baseline and new tax rate, where the average is weighted by the quantities consumed at each tax rate (Weyl and Fabinger, 2013).

additional production costs, ΔAVC . An output tax simplifies this analysis, as $dAVC/d\tau = \gamma = 1$, reflecting how an output tax cannot be avoided through input factor substitution.

Similar results are available for an input tax faced by a monopolist—the incidence of a tax on an input for a monopolist is $I^{Monopoly} = \rho/(dAVC/d\tau) = \rho/\gamma$.¹⁷ Again the second equality follows by the assumption $AVC = MC$. For any standard pass-through rate (greater than zero), consumers bear a greater share of the burden under perfect competition than under monopoly. Figure 2 illustrates incidence under monopoly with increasing marginal costs. The change in consumer surplus is identical to the perfect competition case, and equals the change in price ρ times quantity Q^* . However, the change in producer surplus is entirely determined by the change in average variable costs ΔAVC , as the change in prices is offset by the additional change in quantity.

We now turn to a more general form of competition that nests both perfect competition and monopoly. We assume all firms in the market are identical. Let $\epsilon_D \equiv -[dQ/dP][P/Q]$ denote the elasticity of demand. This is a market-level elasticity, i.e., it describes the change in total market-wise sales in response to a change in the prevailing market price.¹⁸ Let $L \equiv (P - MC)/P$ denote the Lerner (1934) index, a measure of markups, which equals the gap between price and marginal cost, divided by price. In the presence of arbitrary forms of competition, the incidence of an input tax depends on four statistics: the pass-through rate ρ , the cost-shift rate γ , the Lerner index L , and the demand elasticity ϵ_D . Note that perfect competition is a special case where $L\epsilon_D = 0$, and monopoly is a special case where $L\epsilon_D = 1$.

Proposition 1 *Under generalized oligopoly, with N symmetrical firms and $AVC=MC$, incidence takes the form:*

$$I = \frac{\rho}{\gamma - (1 - L\epsilon_D)\rho} \quad (1)$$

Proof. Each of the N symmetric producers maximizes profits selling Q_i units at price P :

$$\pi = (P - MC) Q_i$$

Differentiating profits with respect to an input tax τ and substituting in definitions of L , ϵ_D , ρ , and γ produces the following relationship:

$$\frac{d\pi}{d\tau} = Q_i [(1 - L\epsilon_D)\rho - \gamma].$$

Consumer surplus is simply given by $dCS/d\tau = -Q\rho$. Aggregating across all producer, incidence can then be written as:

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{-Q\rho}{N \cdot Q_i [(1 - L\epsilon_D)\rho - \gamma]},$$

¹⁷The consumer side of the market is calculated in the same way as with perfect competition. The producer side of the market is derived by differentiating producer surplus with respect to the tax rate, then aggregating across firms (Weyl and Fabinger, 2013).

¹⁸See, e.g., Genesove and Mullin (1998). When all firms are homogeneous and so have the same price and quantity, a firm-level demand elasticity is the same as a market-level demand elasticity.

where $Q = N \cdot Q_i$ is the total quantity produced by all N symmetrical producers. This term simplifies to equation (1). ■

Equation (1) has an intuitive explanation. The loss to consumers equals the change in product price, ρ . The loss to producers equals the change in marginal costs, γ , minus the change in product price, ρ . The firm’s change in product prices depends on the term $1 - L\epsilon_D$. Graphically, for most market structures, generalized oligopoly lies somewhere between Figures 1 and 2; the extent to which it resembles one graph versus another depends on the structure of a particular market.

These equations help contrast input and output taxes. For all three cases — perfect competition, monopoly, and general oligopoly — the incidence of input taxes differs from the incidence of output taxes. For an output tax, $\gamma = 1$ since firms cannot substitute away from the tax on a given product. For an input tax, the most likely scenario is $\gamma < 1$, i.e., marginal costs increase less than 1 for 1 with a tax on a single input. This is because firms can substitute away from a taxed input. Thus, conditional on the pass-through rate, input taxes are likely to put a greater share of the burden on consumers (and a smaller share on firms) than output taxes do. Relative to output taxes, input taxes allow firms to substitute away some of the potential increase in marginal cost, thereby lessening the cost to profits.

When taking equation (1) to the data, we use a slightly simpler though analytically equivalent version. Recall that ρ is the pass-through of the tax rate to product prices. Let ρ_{MC} denote the pass-through of marginal costs to product prices, so $\rho_{MC} \equiv dP/dMC$. Dividing the numerator and denominator of (1) through by γ and using $\rho = \gamma\rho_{MC}$ gives

$$I = \frac{\rho_{MC}}{1 - (1 - L\epsilon_D)\rho_{MC}} \quad (2)$$

This second version requires estimating only three parameters: marginal cost pass-through ρ_{MC} , the Lerner index L , and the demand elasticity ϵ_D .¹⁹ To limit the influence of potential outliers in plant-level data, we estimate the pass-through rate $\rho_{MC} \equiv dP/d\tau$ by calculating the proportional pass-through rate $d\log(P)/d\log(\tau)$ and translating it to levels using the relationship $d\log(P)/d\log(\tau) = \rho_{MC}\overline{MC}/\bar{p}$, where \bar{x} denotes the sample mean; we also show that directly estimating the pass-through rate in levels produces generally similar results.

4 Data

The primary data for our analysis comes from administrative survey records collected by the U.S. Census Bureau. We supplement this data with information from the Energy Information Agency (EIA) on energy prices, consumption, and generation.

¹⁹We estimate marginal cost pass-through for two reasons. First, for this homogeneous firms case, it avoids the need to estimate the cost-shift parameter γ . Second, we lack data on plant-level energy input prices.

Census of Manufacturers (CM)

We use administrative data on annual plant-level inputs and outputs from the Census Bureau’s Census of Manufacturers (CM). We use this data to measure plant level inputs, such as capital, labor, and materials.²⁰ The CM is conducted quinquennially in years ending with a 2 or 7, and we draw upon CM years from 1972 through 1997. These sample years are chosen based upon the availability and quality of physical output data in the CM.²¹ We measure labor inputs in hours, capital as plants’ reported book values of equipment and structures, and materials and energy inputs as the reported expenditures on each. We deflate capital, material, and energy expenditures using the corresponding industry-specific input price indices from the NBER Productivity Database.

CM unit prices, which we calculate as product-level revenue divided by quantity, involve several challenges. Since output prices can reflect unobserved product quality, we follow Foster, Haltiwanger, and Syverson (2008, hereafter FHS) in limiting analysis to single-product plants in six industries that produce homogenous products: Boxes, Bread, Cement, Concrete, Gasoline, and Plywood.^{22,23} A minority of firms within these industries have multiple products. For the plants that satisfy these criteria but still produce other products, we follow FHS and scale the focal product output by the inverse of the revenue share. This input-adjustment method assumes inputs are used proportionately to each product’s revenue share. Another challenge is widespread imputation. We exclude any observation identified as an “administrative record” since many of their values are imputed, and we also exclude records where any input or output is imputed.²⁴

Since a few observations still appear to be errors, we make additional sample restrictions similar to those of Roberts and Supina (1996, 2000) and FHS. We exclude a small number of plants reporting physical quantities that imply prices greater than ten times or less than one-tenth the median price in a given industry-year. We also exclude observations missing any one of the main production function variables (i.e., labor, capital, materials, or output quantity). Additionally, we exclude observations where the plant’s labor or materials cost share is less than one-tenth of the

²⁰Labor inputs are measured as plants’ reported production-worker hours adjusted using the method of Baily, Hulten, and Campbell (1992) (i.e., multiplying production-worker hours by the ratio of total payroll to payroll for production workers).

²¹In more recent years, the amount of quantity data collected in the CM has declined considerably, making analyses of more recent time periods infeasible for most, if not all, industries.

²²Following FHS, we define a plant as single product if it receives over half its revenue from the homogeneous product of interest. This definition uses revenue and not quantity shares since different products are measured in different units. When selecting single-product plants, we ignore revenues from product codes for contract work, miscellaneous receipts, product resales, and balancing codes. The Census Bureau creates balancing codes when the summed value of shipments for reported individual products does not equal the plant’s reported total value of shipments.

²³This adds cement to the FHS industries but excludes several of FHS’ industries (sugar, carbon black, coffee, flooring, block and processed ice), for two reasons. Unlike FHS, we exclude observations with imputed quantity, which substantially reduces sample sizes. Also unlike FHS, we estimate industry-specific coefficients in translog production functions. To have sufficient sample size, we exclude industries which, after imposing our sample restrictions, have less than 100 plant-year observations.

²⁴We thank Kirk White for providing the product-level imputation data for the first half of our sample. Many, if not all, researchers using the census microdata for empirical research in economics drop imputed values such as administrative records (see e.g., Doms, Dunne, and Troske (1997); Foster, Haltiwanger, and Syverson (2008); Linn (2008); Atalay, Hortaçsu, and Syverson (2014)).

corresponding industry’s average cost share for that year or when the cost share is more than one. Finally, we trim the one-percent tails of a productivity index measure.²⁵ All unit output prices are adjusted to a common 1987 basis using the revenue-weighted geometric mean of the product price in a given year across all of the plants producing the product in our sample.

It is worth commenting a bit further on the six industries that remain the focus of this paper. Some of these industries are particularly important consumers of energy in U.S. manufacturing. A fourth of U.S. greenhouse gas emissions come from transportation, and most fuels for the transportation sector pass through oil refineries. Cement is one of the largest sources of greenhouse gas emissions in the world. We have chosen these industries because they have price/quantity data and are relatively homogenous, though the relative homogeneity of their products limits product differentiation as a source of market power.

Manufacturing Energy Consumption Survey (MECS) and Annual Survey of Manufacturers Fuels Trailer

We supplement the CM with plant-level data on fuels from the Manufacturing Energy Consumption Survey (MECS), which was collected about every 3 years beginning in 1981, and from the Annual Survey of Manufacturers (ASM) Fuels Trailer, which accompanied the ASM annually between 1973 and 1979, excluding 1977. These surveys report physical fuel consumption separately for each fuel. Energy expenditures in these data exclude feedstocks used for production. Since gasoline refining spends a significant fraction of material costs on energy feedstocks, we augment the energy cost shares for gasoline refining to include the costs of energy feedstocks, which we get from the ASM and CM materials input trailers.

Energy Information Association - State Energy Data System (SEDS)

We use data from the EIA’s State Energy Data System (SEDS) to measure the annual national and state fuel prices for coal, oil, and natural gas. We also use these data to measure the share of electricity generation in a state generated with each of these fuels. The EIA compiles SEDS primarily from surveys of energy suppliers. We convert all fuel prices to real 1987 dollars using the average of the industry-specific energy price deflators for the industries in our sample from the NBER Productivity Database.

5 Econometrics of Pass-Through

We now turn to describe our methodology in six steps. The goal of the first five steps is to estimate the pass-through rate of marginal costs into output prices. First, we describe how we recover marginal costs from production data. Second, we describe how we use production functions to

²⁵The productivity index is constructed using a gross-output, Cobb-Douglas production function with labor, capital, and materials as inputs. The output elasticities are computed using industry-level cost shares under the assumption of constant returns to scale, and output is measured using physical quantities.

recover output elasticities, which are needed to calculate marginal costs. Third, we describe the two research designs for energy prices. Fourth, we describe the analysis of how energy prices affect marginal costs and unit prices. Fifth, we describe how we estimate pass-through. Finally, we describe how we estimate demand elasticities.

Before proceeding, we note two general points about our methodology. One alternative way to estimate pass-through would be to estimate how a plant’s output price changes with respect to a plant’s input price for energy. Our application does not use this approach for two reasons. First, the incidence of input cost shocks typically requires estimates of both the pass-through rate and the cost-shift parameter (γ), but marginal cost pass-through is a sufficient statistic for both input cost pass through and the cost-shift rate. Second, we do not observe plant-level energy prices.²⁶ A separate point is that the method for estimating output elasticities involves strong assumptions that may lead to estimation error. In part for this reason we instrument the resulting estimates of marginal costs with energy price shocks; these instruments may help address any attenuation bias due to measurement error.

5.1 Recovering Marginal Costs

We recover marginal costs by combining plant-level production data with assumptions on firm cost minimization. As originally shown by Hall (1986) and further developed by De Loecker and Warzynski (2012), a firm’s first-order condition implies that the plant’s multiplicative markup (i.e., its price divided by its marginal cost) equals the output elasticity of a variable input like energy or materials divided by the revenue share of that input.²⁷ Thus, if we can identify an output elasticity of a variable input, we can recover a markup. Moreover, price data then let us back out marginal costs as price divided by the markup.

Appendix A.1 describes this methodology formally. Essentially, we compute a time-varying, plant-level markup by using the estimated output elasticity of a variable input and the revenue share of that input. We then recover marginal costs from the accounting identity that price equals markups time marginal costs.

5.2 Output Elasticities and Production Functions

As described in Appendix A.2, to estimate output elasticities we rely on proxy methods (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Akerberg, Caves, and Frazer, 2015). We focus

²⁶The closest information is the “cost of fuels,” which represents total expenditure on all fuels, but does not distinguish each fuel or identify price or quantity. The Manufacturing Energy Consumption Survey does report these variables for a small sample of plants but is not longitudinal. The Census of Manufacturers reports plant-level unit prices for a single energy input—electricity. Using variation in these prices is subject to concerns about bulk discounts (Davis, Grim, Haltiwanger, and Streitwieser, 2013) and endogeneity.

²⁷The output elasticity is defined as the change in a plant’s physical output due to a change in a variable input like materials; the revenue share of a variable input like materials is defined as the plant’s expenditure on that variable input divided by the plant’s total revenue. An intuition for why this approach identifies markups is that in an imperfectly competitive market, input growth must be associated with disproportionate revenue growth.

on production functions with a scalar, Hicks-neutral productivity term and estimate elasticities separately by industry, assuming common technology across firms and over time within an industry.

Essentially, because OLS estimates of a production function may suffer from omitted variables bias due to unobserved plant-year productivity, we use a control-function. We assume a general demand function for materials and invert it to express productivity as a flexible function of a plant’s inputs. We implement this using translog, gross output production functions, separately for each industry. These production functions produce estimates of output elasticities that vary by plant and by year. These estimates include plant fixed effects and exploit variation in energy prices over time, and hence may be thought of as short-to-medium run output elasticities.

5.3 Two Instrumental Variables for Marginal Costs

The previous two subsections show how we estimate output elasticities and marginal costs. These two objects are theoretically sufficient to identify pass-through from a regression of unit prices on marginal costs. However, marginal costs may be measured with error. To address the possible endogeneity of marginal costs, we construct two sets of instrumental variables using variation in energy input prices. Existing research uses other types of instrumental variables for marginal costs, such as tariffs for imported intermediate inputs (e.g., De Loecker, Goldberg, Khandelwal, and Pavcnik (2016)), though not in the context of U.S. manufacturing.

Our two measures of variation in local energy prices leverage the fact that national changes in the price of a fuel disproportionately affect regions and industries heavily dependent on that fuel. For example, when the national price of natural gas rises more than the national price of oil or coal, energy prices in regions and industries heavily dependent on natural gas will disproportionately increase. These “shift-share” designs are related to the Bartik (1991) instrument which is widely used in labor and public economics to study exogenous movements in labor demand (see e.g., Blanchard and Katz (1992); Moretti (2011); Notowidigdo (2011)).

We first focus on variation in industrial *electricity* prices that are driven by regional differences in electricity generation by fuel type. Figure 3 depicts the share of total state-level electricity generation that comes from coal, natural gas, and petroleum/oil, respectively. The maps make clear that the primary fuels used for electricity generation vary considerably over space. Coal accounts for more than three-fourths of fuel for electricity generation in the Upper Midwest but practically no electricity generation in the Western coastal states. Natural gas accounts for 15 percent of generation in the South but over 40 percent of generation in California. We interact this cross-sectional variation in energy input shares (i.e., the shares of fuel costs devoted to electricity generation in a state) with national trends in fuel prices to generate predicted changes in regional electricity prices. Figure 4 shows time-series patterns in the real price of the three primary fuel inputs for electricity — coal, oil, and natural gas. All three fuels had low prices around 1970, a spike during the OPEC crisis in 1975, a decline in the mid-1980s as the crisis subsided, and lastly an increase in the 2000s due in part to rapid economic growth in Asia. While the secular trends are similar among all three fuels, each fuel has independent variation. Coal, for example, was the

most costly fuel in the 1970s but the cheapest in the 2000s. While the 1980s crisis produced abrupt changes in oil and natural gas prices, it led to only gradual and smooth changes in coal prices. As fuel prices vary nationally, differences in marginal electricity generation units' fuel input shares cause that national fuel price variation to differentially affect regional electricity prices or industries dependent on those fuels.

We formalize the relationship between regional electricity prices and regional heterogeneity in fuel inputs used for electricity as follows. We interact (lagged) cross-sectional differences in the share of fuels used to generate electricity in a state-year with national time-series variation in the prices of these fuels to generate three instruments summarized in the vector $z_{s,t}^A$:

$$z_{s,t}^A = [e_{-s,t,f}^A \cdot \sigma_{s,t-k,f}^A], f \in \{\text{coal, gas, oil}\} \quad (3)$$

The variable $e_{-s,t,f}^A$ represents the unweighted national mean over state-level log mean fuel prices f in year t , excluding the state s mean. This research design considers three fuels: coal, oil, and natural gas. The variable $\sigma_{s,t-k,f}^A$ represents the cost of fuel f in year $t-k$, expressed as a share of total fuel expenditure of these three fuels for electricity generation in state s . We present results using lags k of zero, two, and five years. We use the leave-out mean and lagged shares to ensure that energy price variation is independent of local demand for fuels and electricity.²⁸

The second source of variation in marginal costs stems from the fact that different industries use different fuel inputs. Table 1 shows the allocation of energy expenditure across fuels as a fraction of total input expenditures, by industry. Total input expenditures are defined as annual expenditures on salary and wages, capital rental rates, materials, electricity, and fuels. For example, 0.7% of total input costs for box manufacturing come from natural gas, but 13 percent of total input costs in cement come from coal. We formalize the predicted variation in industrial energy prices by interacting national, leave-out mean energy input prices for industrial consumers with the (lagged) share of energy input costs in an industry-year devoted to a particular fuel.

$$z_{n,t}^B = [e_{-n,t,f}^B \cdot \sigma_{n,t-k,f}^B], f \in \{\text{coal, gas, oil, electricity}\} \quad (4)$$

Here, the vector $z_{n,t}^B$ contains four instruments, one for each fuel. The variable $\sigma_{n,t-k,f}^B$ represents the share of total expenditures in industry n and year $t-k$ devoted to fuel f . The variable $e_{-n,t,f}^B$ denotes the national, leave-out mean input price of fuel f for industrial consumers.²⁹ This research design considers four energy inputs: coal, oil, natural gas, and electricity.

There are three main differences between the equations used in the two research designs (3)

²⁸While state-based leave-out mean shares may be susceptible to local demand due to regional spillover effects, we include region-year fixed effects as a robustness exercise and note that these instruments improve on the existing literature Aldy and Pizer (2015).

²⁹We calculate energy input expenditure shares at the industry level using the data from MECS and the ASM fuel trailers. In principle, we could compute energy input expenditure shares by plant or industry×region. Small sample sizes in both the MECS and the ASM preclude the use of more granular industry energy input expenditure share definitions, such as $e_{i,-s,t,f}^B$. Industry×region energy input expenditure share definitions deliver similar but less precise results.

and (4). First, the electricity price research design uses national energy input prices for electricity generation $e_{-s,t,f}^A$, whereas the energy price research design uses national energy input prices for industrial consumers $e_{-n,t,f}^B$. Second, the energy price research design includes electricity along with the three primary fuels as energy inputs, while the electricity price research design only uses the three primary fuels used for electricity generation. Lastly, energy input shares σ are calculated differently in both equations; equation (3) calculates shares as the fraction of electricity generation in a state that comes from fuel f , and equation (4) calculates energy expenditures as a fraction of total expenditures for a given industry, fuel, and year.

5.4 Effects of Energy Prices on Unit Prices, Marginal Costs, and Markups

The goal of these “shift-share” research designs is to use the predicted sources of variation in energy prices to analyze the pass-through of marginal costs to product prices. However, it is also informative to understand how these measures of energy prices are related to unit prices and marginal costs. In the context of instrumental variables, one can interpret the relationship between unit prices and these fuel price \times fuel share interactions as the reduced form, whereas the relationship between marginal costs and the fuel price \times fuel share interactions as the first stage. We also examine the relationship between this energy price variation and plant level markups.

We investigate this question with the following fixed effects regression model:

$$y_{ist} = z'_{nst}\beta + X'_{nst}\gamma + \eta_i + \pi_t + \nu_{ist} \quad (5)$$

Equation (5) describes a regression of outcome y in logs (unit prices, marginal costs, or markups) for plant i in state s and year t . The vector z_{nst} represents either the interaction between national fuel prices and state-level electricity generation shares from z_{st}^A (i.e., equation (3)) or the interaction between national energy prices and industry-level energy input shares from z_{nst}^B (i.e., equation (4)). In either case, we include each of the fuels as separate variables in the vector z . The vector X_{st} includes the leave-out mean energy/fuel prices $e_{-s,t,f}$ or $e_{-n,t,f}$ separately for each fuel. It also includes either the generation share $\sigma_{s,t-k,f}^A$, measured k years ago in state s , or the industry energy input share $\sigma_{n,t-k,f}^B$, measured k years ago in industry i . Some specifications also control separately for differential trends by state, region by year fixed effects, and industry by year fixed effects. The regression also includes plant fixed effects η_i and year fixed effects π_t . While equation (5) describes a plant-level regression from our analysis sample, we estimate it using census weights to obtain a pass-through rate representing the entire market.

The interpretation of the coefficient vector of interest β differs across the two research designs. When estimating equation (5) using the electricity price shift-share z_{st}^A , the vector β describes the elasticity of outcome Y_{ist} ($= \exp(y_{ist})$) with respect to the national (leave-out mean) fuel price for fuel f for a state that produces 100 percent of its electricity using fuel f . In practice, no state generates all of its electricity from a single fuel source, so these coefficients should be evaluated at the mean value of a fuel’s generation share. During our sample, this is about 20 percent for natural

gas. When estimating equation (5) using the energy price shift-share z_{nt}^B , the vector β describes the elasticity of outcome Y_{ist} with respect to the national (leave-out mean) fuel price for fuel f in an industry for which 100 percent of the industry’s total annual expenditures are devoted to energy input f , where f is electricity, fuel oil, natural gas, or coal. Again, no industry uses a single fuel input for production, and so these coefficients should be evaluated at the mean value of respective energy expenditure share (see Table 1).

5.5 Pass-Through

We use our data on unit prices, recovered estimates of marginal costs, and constructed energy price variation to estimate marginal cost pass-through. We estimate the marginal cost pass-through elasticity from the following plant-level regression of (log) output price on (log) marginal costs:

$$p_{ist} = \rho_{MC,\epsilon} mc_{it} + X'_{nst} \gamma + \eta_i + \pi_t + \epsilon_{ist} \quad (6)$$

The main coefficient of interest, $\rho_{MC,\epsilon}$, measures the elasticity of unit prices with respect to marginal costs. Note that $\rho_{MC,\epsilon}$ differs from the marginal cost pass-through rate ρ_{MC} outlined earlier in that $\rho_{MC,\epsilon}$ represents an elasticity whereas ρ_{MC} is pass-through in levels. The marginal cost pass-through rate can be calculated by multiplying the pass-through elasticity by the markup: $\rho_{MC} = \rho_{MC,\epsilon} \times P/MC$.

Equation (6) includes the same vector of controls X_{nst} as above, along with plant fixed effects η_i , year fixed effects π_t , and an idiosyncratic error ϵ_{ist} .³⁰ As above, we report additional regression estimates that control separately for differential trends by state, region by year fixed effects, and industry by year fixed effects. We instrument for mc_{it} using either z_{st}^A or z_{nt}^B . As mentioned above, equation (5) describes both the first stage and reduced form from an instrumental variables regression of price on marginal cost in equation (6).

Since price equals the product of markups and marginal costs ($P_{it} = MC_{it} * \mu_{it}$), one might wonder whether the elasticity $\rho_{MC,\epsilon}$ in equation (6) should equal one. The answer is no. Since this pass-through regression does not explicitly control for markups, the markup term is in the regression error. Even abstracting from potential econometric issues, pass-through estimates capture the extent to which marginal costs do not perfectly predict product prices due to variability of markups (see De Loecker, Goldberg, Khandelwal, and Pavcnik (2016) for related discussion). The other controls in equation (6) we include for two econometric reasons—they adjust for measurement error in estimated marginal costs (e.g., a firm-specific component of measurement error is absorbed by η_i), and since the shift-share instrument is valid conditional on these controls (e.g., the time fixed effects help adjust for effects of energy prices on consumer incomes).

³⁰In both research designs, the excluded instrument is the interaction of the leave-out mean fuel price with the share of that fuel used for energy. We also control for both of these variables in levels as exogenous controls X. The most informative way to interpret the interaction terms is to evaluate them at some share (e.g., what is the effect of gas price increases in an area that uses gas for 100% of energy). Because these energy prices are approximately the national value (subject to leaving out the own-state price), the energy price controls in this regression are similar to flexible year trends.

5.6 Demand Estimation

Finally, in order to consider incidence in markets that are less than perfectly competitive, we need to estimate industry-specific demand elasticities. As discussed earlier, we need a market-level elasticity; some conventional methods, like those of FHS, estimate the elasticity of residual demand for an individual plant. We use data from the NBER-CES manufacturing industry database to estimate demand elasticities from the following equation:

$$\Delta \log Q_t = -\epsilon_D \Delta \log P_t + \alpha_y t + \alpha_c + \eta_t$$

We estimate a demand elasticity separately for each industry. Here Q_t is the industry’s real output in year t , equal to the value of shipments deflated by the industry-specific price deflator. On the right-hand side, P_t is the industry’s output price deflator, t is a time trend, α_c is a constant, and η_t is the regression error. The coefficient of interest, ϵ_D , is the elasticity of demand. We estimate the regression in first-differences since the data may be non-stationary. Because price and quantity suffer from a classic simultaneity problem, we instrument P_t with a total factor productivity index, calculated from a five-factor Cobb-Douglas model.

This methodology estimates the demand elasticity described in the incidence theory section. The formulas in that section depend on the effect of a change in a market-level price on the industry’s quantity, and this elasticity is common across firms. Because this methodology uses the real output level for the industry and year, it corresponds well with the theory.

6 Results

Now that we have described the theory, data, and methodology, we turn to describe four sets of results: mean levels of prices, markups, and marginal costs; the effects of energy input cost shocks on these variables; marginal cost pass-through; and incidence.

6.1 Average Levels of Marginal Costs, Unit Prices, and Markups

Table 2 presents some key summary statistics by industry. Column (1) shows that total annual expenditures on electricity and fuels comprise a limited fraction of input costs.³¹ The mean energy cost share for most industries in our sample is 2 percent of total annual input expenditures, though cement and gasoline refining are quite different from the others, with 32 and 84 percent of total expenditures devoted to electricity and fuels.³² While energy is not a large cost share for several of these industries, it is also a modest cost share of the entire U.S. economy, at 3 to 4 percent of

³¹Energy cost shares are defined as annual energy and electricity expenditures at a plant divided by salary and wage payments, rental payments on capital stocks, and expenditures on materials, electricity, and fuels.

³²As mentioned earlier, the energy cost share for gasoline is high in large part because much of the crude oil used at refineries is physically transformed into gasoline (i.e., it is “feedstock”) rather than being combusted at the refinery for heat or fuel.

gross output. Firms and consumers are not concerned that energy is their largest single cost, but that it is a cost affected substantially by regulation in ways that are politically sensitive.

Columns (2) through (4) show median estimates of output elasticities separately for labor, materials, and capital. As described in the previous section, we estimate production functions, and hence output elasticities, separately for each industry using a translog, gross-output production function, so output elasticities vary by plant. A few points here are worth noting. First, median output elasticities for all industries and factors of production are positive. Second, elasticities vary considerably across industries. The estimated output elasticity for labor, for example, is 0.04 for boxes but 0.91 for cement. For each of the three factors of production, cement has the largest estimated output elasticity among all industries. Third, the output elasticities differ considerably across factors of production. For all industries, materials has the largest output elasticity, at between 0.6 and 1.1. Capital and labor output elasticities are much smaller.

Column (5) of Table 2 presents the median of the sum of the three output elasticities, which is a measure of returns to scale in the industry. The results suggest that, except for cement, all the industries in our study cement have approximately constant returns to scale, with estimated returns to scale of 0.92 to 1.13. Cement is quite obviously different from the others, with strongly increasing returns to scale at 2.46. Foster, Haltiwanger, and Syverson (2008) find constant returns to scale for all of these industries except cement, which they do not analyze. Since cement uses some of the largest industrial equipment in the world, it may not be surprising that we estimate cement to have increasing returns to scale. Column (6) shows average markups of ten to fifty percent across industries except for cement, which is quite different from the others, with markups of 130 percent. The product with the lowest estimated markup is gasoline, with markups of only 11 percent. These levels of industry markups are largely consistent with the existing literature that estimates markups using production function methods (Hall, 1986; De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016; Collard-Wexler and De Loecker, 2015).³³

Columns (7) and (8) present mean unit prices and marginal costs by industry. Unit prices equal revenue divided by output quantities. Marginal costs equal the log output price minus the log markup. Lastly, column (9) presents the share of materials expenditures as a fraction of total revenue, which is used to construct markups in column (6).

6.2 Effects of Energy Prices on Marginal Costs, Unit Prices, and Markups

Table 3 presents our baseline set of results describing how variation in fuel input costs for electricity generation differentially affects markups, marginal costs, and unit prices for U.S. manufacturing plants. Table 4 presents similar estimates using industry heterogeneity in energy input cost shares.

³³ De Loecker, Goldberg, Khandelwal, and Pavcnik (2016) find industry median markups in Indian manufacturing ranging from 1.15 to 2.27. Collard-Wexler and De Loecker (2015) explore markups in the U.S. steel industry ranging from 1.2 to 1.5 depending on the time period studied. Hall (1986) explores markups using more aggregated industry definitions (2-digit SIC) but finds results that are largely consistent with our own. For example, Hall (1986) finds an estimated markup for SIC 34, which consists of both Cement and Ready Mix Concrete, of 1.81. Similarly, Plywood manufacturing, SIC 24, has an estimated markup above marginal cost of 1.0.

These two tables reflect the two research designs in equations (3) and (4), and they stem from various versions of equation (5) which are estimated using OLS. Standard errors in this table and all subsequent regression output are clustered by state, unless otherwise mentioned. Because the first research design focuses on cross-region differences in electricity prices, it has stronger external validity to sub-national greenhouse gas mitigation policies like the Regional Greenhouse Gas Initiative, a cap-and-trade for electricity in the Northeast. By contrast, the second research design has national scope. Ultimately, the two distinct research designs lead to similar pass-through and incidence estimates.

Panel A of Table 3 analyzes how variation in fuel prices affects plant-level marginal costs. We report coefficients on the interactions between fuel prices and (lagged) state-level electricity generation shares. The regression models also control separately for the levels of fuel prices and for generation shares. Since both the dependent variable and the fuel variables are measured in logs, the reported interaction terms represent an elasticity for states in which 100 percent of generation comes from a given fuel. Each column in the table shows a slightly different specification. Columns (1) and (4) use the contemporaneous state-level generation mix for each fuel; columns (2) and (5) use two year lags, and columns (3) and (6) use five-year lags of the generation shares. Columns (1) through (3) control for plant fixed effects, year fixed effects, and state trends, whereas columns (4) through (6) control for product-year fixed effects, region-year fixed effects, and state trends. The identifying variation across all columns comes from within-plant variation in fuel shares and prices, while adjusting for various forms of time-varying observed and unobserved determinants of marginal costs that may be correlated with changes in predicted regional fuel prices.

These regressions give expected signs and plausible magnitudes. The coefficients in Panel A of Table 3 are mostly positive, which is expected since shocks to fuel input prices are likely to increase electricity prices that firms face and hence increase their marginal costs of production. The smaller coefficients in columns (4)-(6), those with region-year and product-year fixed effects, seem to suggest that there may be other correlated unobservables driving costs and prices in the absence of these controls. Note, however, that the ratio of the coefficients in Panels A and B remain largely similar across columns. The magnitudes of these coefficients are also reasonable. Recall that energy is on average 2 percent of production costs for most of these industries, and electricity expenditures represent an even smaller share. Column (6), which includes a 5-year lag in generation shares, implies that if natural gas provides 100 percent of in-state electricity generation, then a 1 percent increase in the price of natural gas used for electricity results in a 0.29 percent increase in the marginal costs of manufacturing production. In practice, natural gas constitutes about 20 percent of electricity generation over this time period.³⁴ Thus, these estimates imply that a 1 percent increase in natural gas prices would cause an 0.06 percent increase in marginal costs.³⁵ The

³⁴20 percent natural gas generation corresponds to the fraction of production that is generated by coal, natural gas, and oil (i.e., excluding other methods of generation such as nuclear or hydroelectric).

³⁵Note that the results in Table 3 do not control for or model industry specific variation in fuel input usage. Thus, to the extent that national fuel prices affect the marginal cost of production through channels other than electricity prices (e.g., through increased fuel costs used directly for production), these estimates implicitly incorporate this variation. The second research design used in Table 4 does account for raw coal, oil, and gas used as direct inputs to

magnitude of this coefficient varies across fuels. Since this research design only exploits variation in electricity markets, it is not surprising that the gas coefficient is more precise than the oil or coal coefficients since oil is used for little electricity generation, coal has more regional markets (so the leave-out mean is less predictive), and gas is more often the marginal generation technology.³⁶ The response of power plants and regulators to energy price shocks through fuel substitution, efficiency improvements, or strategic production decisions, however, may also vary across fuels.

Panel B of Table 3 presents a version of equation (5) which uses log unit prices as the dependent variable. The signs and precision are similar to the marginal cost estimates of Panel A. The positive and precise estimates in Panel B provide a first piece of evidence that energy price shocks are passed through to plant-level unit prices. The fact that the effects of energy price shocks on unit prices (in Panel B) are smaller than the estimated effects on marginal costs (in Panel A), however, suggests that these cost shocks are less than fully passed through to prices.

Finally, Panel C of Table 3 presents a version of equation (5) which uses log markups as the dependent variable. Since markups are defined as the ratio of prices to marginal costs, the effects of changes in energy prices on markups approximately equal the difference between their effects on marginal costs and unit prices. Thus, the numbers in Panel C can largely be inferred from the numbers in Panels A and B. The point estimates suggest that increasing the price of fuels used for electricity causes modest decreases in markups. The more stringent specifications of columns (4)-(6) imply that increasing the natural gas price by 1 percent for in-state electricity generation decreases markups by around 0.01 percent.³⁷

Table 4 presents a similar set of results, exploring how markups, marginal costs, and unit prices differentially respond to changes in energy input costs used for production. As before, we report the coefficients from the interaction between energy input shares, which is defined as the share of total input costs devoted to a specific energy input, with national energy prices, which are defined as the leave out mean national industrial energy price, omitting the focal state's average industrial energy price. We focus on the four primary energy inputs into manufacturing production: coal, natural gas, oil, and electricity. Since we use information from MECS and the ASM fuel trailers to construct energy input shares, and the ASM trailers only begin in 1975, we are not able to lag these shares more than 2 years without losing an extra year in our analysis sample.

Panel A of Table 4 suggests that increases in the costs of energy used for manufacturing increase marginal costs. The interaction term represents the elasticity of marginal costs with respect to energy input prices in a plant for which 100 percent of total expenditures comes from one of the four energy sources. Thus, if we were able to perfectly measure fuel prices faced by a plant, we might expect that the coefficient should be close to 100. In practice, there are many reasons that the coefficients may deviate from 100, not the least of which is that fuel expenditure shares in the

manufacturing.

³⁶According to the State Energy Database System, oil accounted for about 15 percent of BTUs used in the electric power sector in the 1970s, 5 percent in the 1980s and 3 percent in the 1990s

³⁷This comes from calculating the means across columns (4) through (6), weighting by the inverse standard error, and then multiplying by the average generation share of natural gas in our sample of 0.2.

neighborhood of 100 percent are far out of sample. Moreover, we observe average industrial fuel prices which might be quite different than marginal fuel prices faced by the firms in our sample. The broad takeaway is that relative fuel input usage interacted with national fuel prices strongly predicts changes in marginal costs, unit prices, and markups. Subsequent sections use these relationships to quantify by how much fuel and electricity price induced increases in marginal costs are passed through to consumers in the form of higher unit prices.

6.3 Marginal Cost Pass-Through

We now take the estimated relationships between prices, marginal costs, markups, and energy prices and embed them into a pass-through regression of unit prices regressed on marginal costs. Table 5 presents OLS regression estimates corresponding to various versions of equation (6). Column (1) presents estimates with plant fixed effects, year fixed effects, and state trends. Columns (2) and (3) add product-year and region-year fixed effects, respectively. Column (4) includes the full set of product-year and region-year controls to the model from Column (1). Since both the dependent and independent variables are in logs, the coefficient estimates measure an elasticity. The estimates across the four columns are broadly similar, suggesting that a 1 percent increase in marginal costs is associated with a 0.6 percent increase in the unit price. In order to convert this pass-through elasticity into a pass-through rate, we multiply the coefficient by the average markup in the sample of 1.15, which gives a marginal cost pass-through rate of 0.7.

To address the possible endogeneity of marginal costs, Table 6 presents instrumental variable estimates of equation (6), where marginal costs are instrumented in two separate ways. Panel A presents pass-through estimates where marginal costs are instrumented using the interactions between the lagged electricity generation mix in a state and the log of the leave out mean electricity fuel input prices. Panel B instruments marginal costs with the interaction between an industry’s share of annual expenditures devoted to one of four energy inputs multiplied by the national, leave-out mean of that energy input price. As before, columns (1) through (3) include plant fixed effects, year fixed effects, and state-specific trends. These columns vary the lag in the fuel share component of the marginal cost instrument between zero, two, and five years. Columns (4) through (6) of Panel A add region-year fixed effects and product-year fixed effects, whereas Columns (4) through (6) of Panel B add only region-year fixed effects.³⁸ We present estimates showing zero, two, and five lags of the fuel share. All columns include the uninteracted log of the fuel prices and the uninteracted, lagged generation mix (not reported).

Panel A suggests that pass-through elasticities range between 0.62 and 0.72, which translates into pass-through rates of 0.71 to 0.83. The strength of the instrument varies slightly across the specifications, with first stage partial F statistics ranging from 4 to 14. These first stage F-statistics suggest there may be bias stemming from a weak first stage, where the bias is towards the OLS counterpart. The results in Panel B are broadly consistent with those in Panel A, though they are estimated using a different instrument for marginal costs. The energy price instrument in Panel B

³⁸Industry expenditure shares for fuel inputs are collinear with product-year fixed effects in Panel B.

is also a stronger predictor of marginal costs than the electricity price instrument in Panel A, as reflected by the first stage F-statistics of between 66 and 137. The average pass-through elasticity in Panel B is around 0.51, which translates into a pass-through rate of 0.59.

This overall pass-through elasticity may hide important cross-industry heterogeneity. Table 7 reports estimates of pass-through rates separately for each of the six industries in our sample. Panel A presents specifications that instrument marginal costs with our electricity price instrument, controlling for year fixed effects and state specific linear time trends. Panel B presents estimates that also include region-year fixed effects. All regressions control for plant fixed effects. Note that year fixed effects in an industry-level regression correspond to industry-year fixed effects from a pooled cross-industry regression. These tables reveal substantial cross-industry heterogeneity in the pass-through of marginal costs into unit prices. Pass-through elasticities vary from a high of 0.96 for boxes to a low of 0.33 for gasoline refining.³⁹ Table 7 also reports the corresponding pass-through rates, which are calculated by multiplying the elasticity estimate by the industry-specific markup. These estimates suggest that boxes, cement, and plywood “overshift” changes in marginal costs into changes in unit prices. The pass-through rates for these industries vary between 1.02 and 1.78. Conversely, bread, concrete, and gasoline refining have comparatively low pass-through rates ranging from 0.36 to 0.82.

Appendix Table B1 presents industry-specific pass-through estimates. Marginal costs here are instrumented using the interaction between an industry’s share of annual expenditures devoted to one of four energy inputs and the national, leave-out mean of that energy input price. The results remain quite similar to those in Table 7. One challenge with estimating industry-specific pass-through rates using industry-year level instruments is the inability to control for other time-varying shocks that may affect both marginal costs and output prices in a given year (i.e., we cannot include year fixed effects). Since our energy prices are calculated as leave out mean of the national energy price (excluding the focal state), the model is still identified when one includes year fixed effects. However, the identifying variation in a model with year fixed effects stems from idiosyncratic differences in state-level fuel prices driven by the “leave-out variation that may be more subject to endogeneity concerns. With these limitations in mind, Appendix Table B1 presents estimates that are qualitatively similar to industry-specific estimates using the electricity price instrument. As discussed earlier, in order to limit the potential influence of outliers, all these estimates come from estimating the proportional pass-through rate then translating it into levels. Appendix Table B2 estimates the pass-through rate in levels directly, and obtains generally similar results.

How reasonable are these numbers? Some industries exhibit pass-through rates that exceed one, suggesting that producer surplus may actually increase due to a change in tax. Various researchers have shown the conditions under which overshifting can occur under oligopoly (see e.g., Bulow and Pfleiderer (1983); Seade (1985); Delipalla and Keen (1992)). The general intuition from this literature suggests that in the presence of oligopoly, a single firm can raise industry prices by

³⁹In both research designs, the first-stage F-statistics for gasoline and bread are not large, so we interpret pass-through estimates for these two industries cautiously given that they suffer from weak instruments.

reducing its own output, but it fails to do so because it is costly for the single firm in terms of foregone profits. A tax-induced cost increase will necessarily induce output-reductions for all firms, imposing upon the producers some of the collusion they themselves had been unable to achieve.⁴⁰ Studies have found pass-through rates ranging from below unity (Goldberg and Hellerstein, 2008; Gopinath, Gourinchas, Hsieh, and Li, 2011; Campa and Goldberg, 2005), to above unity (Besley and Rosen, 1999), to equal to unity (Poterba, 1996; Fabra and Reguant, 2014), depending on the methodology, market, time-period, or data used. Empirically, we observe higher pass-through rates in industries where we see relatively inelastic demand and higher markups. Conversely, we observe lower pass-through rates in industries with elastic demand and lower markups. As emphasized by Seade (1985) and more recently by Weyl and Fabinger (2013), pass-through in imperfectly competitive product markets is closely related to the curvature of demand — very convex demand will typically have pass-through rates exceeding one, whereas concave demand will have pass-through rates below one.

Whether firms actually benefit from the increase in input costs is an empirical question that turns on whether the pass-through of the tax increases revenues enough to overcome both the lower demand due to higher prices as well as the direct cost increase due to the tax. We now turn to this question by computing incidence for each of the six industries in our study.

6.4 Incidence

Table 8 describes the incidence of an energy input tax separately for each of the six industries. Panel A presents the necessary components for calculating incidence using equation (2).

The full estimates of demand elasticities are presented in Appendix Table B4. We estimate that most industries have fairly elastic demand, with demand elasticities between -2 (plywood) and -5 (bread). Gasoline refining has relatively inelastic demand, with an elasticity of -0.13.⁴¹ These demand elasticity estimates use total factor productivity as an instrumental variable for price. For two industries – bread and concrete – this instrument is not strong. For others, it borders on being a strong estimates (first-stage F of 8.2 to 11.7); the estimates with instruments that are less strong may be biased towards OLS.

We use the average industry pass-through rate from Table 7.⁴² The table shows standard errors for incidence, which we calculate from a bootstrap that is clustered by state.⁴³

⁴⁰Seade (1985) describes this as a “public-goods problem” – restraint by any one firm in the industry in question raises the prices they all face for their outputs, constituting a common benefit. But the cost of this restraint, in the form of profitable revenue foregone, is borne by that one firm alone. Hence, too little of that good is produced (i.e., there is too little restraint). A cost increase will necessarily reduce output, thus raising the supply of the public good “restraint.” In other words, in the absence of explicit collusion, the tax acts as a coordinating device allowing oligopolists to restrict output and thereby increase profit.

⁴¹The most widely cited estimate of the elasticity for short-run retail gasoline demand is -0.03 to -0.08 for the period 2001 to 2006, and -0.21 to -0.34 for the period 1975 to 1980 Hughes, Knittel, and Sperling (2008); estimates for long-run demand are somewhat larger Brons, Nijkamp, Pels, and Rietveld (2008). Ryan (2012) estimates a demand elasticity for cement of -2.96, which is close to our estimate of -2.90.

⁴²We average the two pass-through estimates for each industry weighting by the inverse of the standard error of each estimate.

⁴³We implement the bootstrap by drawing 50 states with replacement from the plant-level data, then using this

To simplify exposition, we report the change in consumer surplus as a share of the total change in surplus between producers and consumers, or $I/(I+1) = \frac{(dCS/d\tau)}{dCS/d\tau + dPS/d\tau}$. We find this exposition simpler since it describes incidence on a scale along 0.0 to 1.0, where 0 means producers bear the entire burden while 1.0 means producers do.

Panel B presents incidence estimates. In all industries, consumers bear substantially less than 100 percent of the burden. For gasoline refining, we estimate that consumers bear 24 percent of the welfare burden of increases in energy input prices. For plywood, we estimate that consumers bear 66 percent of the welfare burden. For the other products we study (concrete, bread, boxes, cement), consumers bear 50 to 60 percent of the welfare burden. In these data, we conclude that both producers and consumers typically bear part of the welfare burden—for each industry, the confidence intervals reject the hypothesis that consumers bear the entire burden (i.e., that the incidence statistic in Panel B equals one), and for every industry except gasoline, we reject the hypothesis that producers bear the entire burden.

In general, the industries with lower pass-through rates force consumers to bear a smaller share of the burden of the increase in input costs. The cement industry has the greatest estimated pass-through rate, and also nearly the greatest share of the burden borne by consumers (0.63). By contrast, the gasoline refining industry has a lower estimated pass-through rate, and correspondingly a lower share of the burden borne by consumers (0.24).

Rows 5 and 6 of Table 8 present our estimates of incidence under the assumption that firms engage in monopoly or perfect competition, respectively. Incidence under monopoly assumes the Lerner index is the inverse of the demand elasticity, and incidence under perfect competition assumes that the Lerner index equals zero, since perfectly competitive firms do not charge markups. Therefore, computing incidence for the polar cases of perfect competition or monopoly only require estimates of marginal cost pass-through and does not require either demand elasticities or firm markups. Moreover, if products in an industry are substitutes and there is no collusion, these polar cases will bound the “true” incidence in the presence of market power. We also note that the pass-through rate itself may depend on the market structure. The conceptual exercise here is that we have estimated the pass-through rate given the existing (unknown) market structure. Conditional on the estimated pass-through rate, the bounding exercise calculates the possible values of incidence that could be consistent with the estimated pass-through rate.

The conventional approach to accounting for market power in estimating partial equilibrium tax incidence assumes the market is a monopoly. For the six industries we study, which appear to have substantial market power, we conclude that monopoly is not a terrible assumption, but may moderately understate the share of the welfare burden that consumers bear.

Specifically, Panel D shows that for three of our industries, the assumption of monopoly is

bootstrap sample to calculate the pass-through rate and Lerner index. For this bootstrap sample, we also draw a value of the demand elasticity from a normal distribution with mean and standard deviation given by Appendix Table B4 . We calculate the incidence statistic for this bootstrap sample. We then repeat this entire procedure 1000 times, generating an empirical distribution of estimates for the incidence statistic. We exclude the few of the 1,000 bootstrap estimates for which the General Method of Moments estimate of output elasticities fails to converge, which mainly occur in the industries with small samples. We report the standard deviation of this incidence statistic.

reasonably accurate; for the other three, the monopoly assumption overstates the share of the welfare burden that firms bear and understates the share that consumers bear. In boxes, cement, and gasoline refining, consumers bear almost exactly the same share of the burden if we estimate markups and a demand elasticity and allow for arbitrary oligopoly as if we assume each of these industries is a monopoly. For example, the consumer share of the welfare burden in the cement industry is 0.63 if we allow for arbitrary oligopoly and 0.64 if we assume the industry is a monopoly. In the other three industries (bread, concrete, and plywood), consumers bear a larger share of the welfare burden if we allow the industry to have arbitrary forms of oligopoly than if we assume the industry is a monopoly.

Panel E presents incidence of input costs under the assumption of perfect competition. Perfect competition is the standard assumption in most partial equilibrium models used to study incidence, and in most general equilibrium models used to study incidence. For most industries we study, the incidence estimates suggest that perfect competition poorly approximates our preferred oligopoly estimates of incidence. For most industries, however, the perfect competition estimates are imprecise, and we cannot statistically reject the hypothesis that they equal the oligopoly estimates.

Specifically, for all six industries, assuming the industry is perfectly competitive obtains a larger estimate of the share of the welfare burden that consumers bear than assuming the industry is an arbitrary form of oligopoly would imply. In the concrete industry, the assumption of arbitrary oligopoly implies that consumers bear half the welfare burden (standard error 0.07), while the assumption of perfect competition implies that consumers bear three-fourths of the welfare burden (standard error 0.16). In the box, cement, and plywood industries, consumers face more than 100% of the burden as producers pass through over 100% of marginal cost increases onto the consumer.

This section reports overall incidence for each industry. Because we estimate markups and therefore the Lerner index (L) for each plant in an industry, and because we can calculate plant-specific pass-through rates, we can also estimate the distribution of incidence across plants in each industry.⁴⁴ We calculate this in Appendix Table B3 by combining plant-level markup estimates and pass-through rates with industry-level demand elasticities. This calculation finds that the mean plant has an incidence close to the overall incidence values from Table 8, and the distribution of incidence across plants is similar to the standard errors for incidence shown in Panel B of Table 8.

7 Conclusion

In this paper we develop a methodology for estimating the partial equilibrium incidence of input cost shocks that is able to relax assumptions pertaining to complete pass-through, factor substitution, and industry competitiveness. The approach leverages the fact that equilibrium prices, and the extent to which they respond to cost shocks, are sufficient statistics for more primitive demand

⁴⁴The plant-specific pass-through rates come from dividing the proportional pass-through rate $d\log(P)/d\log(MC)$ by the ratio MC_i/P_i . The distribution of incidence across plants could differ from results in Table 8 because incidence is a nonlinear function of the Lerner index L . If \bar{x} denotes the sample mean of x and $I(\cdot)$ describes incidence, then Table 8 reports $I(\bar{L}, \bar{\rho}_{MC})$, while this paragraph discusses $\bar{I}(L_i, \rho_i)$.

and supply parameters describing a market. With further information on markups and demand elasticities, we are able to extend this intuition to general forms of imperfect competition.

We consider the specific application of energy cost shocks for U.S. manufacturing, and we assess the extent to which the welfare consequences of these shocks are borne by manufacturing producers versus consumers. Many analyses of the incidence of various types of energy cost shocks, such as climate change mitigation regulations, assume perfect competition and complete pass-through. Our results suggest that the short- to medium-run incidence of a change in input costs, defined as the ratio of the change in consumer to producer surplus, differs dramatically from these assumptions. In one industry consumers bear almost 90 percent of the total change in input costs. This result stems from the ability of producers to more than fully pass through the change in input prices into unit output prices. In other industries, especially those with elastic demand, producers bear a greater incidence of the change in input costs.

We emphasize several caveats. First, we focus on intermediate goods, while traditionally the measurement of consumer surplus focuses on final goods markets. In some cases the change in surplus due to taxing intermediate inputs is identical whether the analysis focuses on intermediate or final goods, though these cases can involve strong assumptions. Second, we focus on six homogeneous products since they have plant-level quantity output data, and since homogeneity diminishes concerns about unobserved product quality biasing estimates. The small number of industries and their relative homogeneity limits external validity. The use of plant-level quantity output data in our framework also limits its potential applicability, though an increasing number of datasets and settings have data on output quantities. Third, our research design relying on regional variation in fuel prices is most relevant to regional (sub-national) greenhouse gas policies, although we obtain similar results using a separate research design that exploits national variation in fuel prices, and the second research design is more relevant to national policies. Fourth, this partial equilibrium analysis for several industries has narrower scope than a general equilibrium analysis of all industries in the economy, though it also can accommodate arbitrary forms of imperfect competition and other features that are more difficult to analyze in general equilibrium analyses. Finally, it uses short- to medium-run variation; an important question we leave for future work is the importance of allowing for non-competitive market structures in long-run analyses.

This paper also arrives at a few additional conclusions that may be of broader interest. First, the common assumptions of complete pass-through and perfect competition in research on the incidence of commodity price shocks may be overly strong; these assumptions appear to be incorrect for all the industries we study. In the short- to medium-run, firm owners may shoulder more of the burden than existing research would suggest. On the other hand, some industries may actually benefit from increasing energy prices through a carbon tax. For example, in industries that “overshift” input taxes, the firm-level gains from reducing output in an economy with market power may exceed the direct detrimental effect of the tax for the firm. While we do not observe energy price-induced increases in producer surplus in this setting, more work is needed to understand how these results generalize into other industries and time periods. The second more general takeaway is

the considerable heterogeneity across industries; the one industry which has been the subject of scrutiny in research on market power and the environment – cement – seems to differ from the five other industries we study in much of the analysis.

One productive avenue for future work is to overcome the existing limitations associated with the paucity of producer price data linked to information on firm input and output decisions in the United States, which might facilitate estimating firm-specific pass-through rates, as opposed to industry-wide averages. Another is to incorporate our conclusions about incomplete pass-through and industry competitiveness into an economy-wide, general equilibrium framework that can both account for pre-existing distortions in other markets while analyzing incidence by different consumer demographics. Lastly, the technology developed here is well-suited to think about incidence of other changes in input costs across industries, for example, stemming from changes in minimum wage laws or increased capital costs. We leave these avenues for future work.

References

- ACKERBERG, D. A., K. CAVES, AND G. FRAZER (2015): "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83(6), 2411–2451.
- ALDY, J. E., AND W. A. PIZER (2015): "The Competitiveness Impacts of Climate Change Mitigation Policies," *Journal of the Association of Environmental and Resource Economists*, 2(4), 565–595.
- ATALAY, E., A. HORTAÇSU, AND C. SYVERSON (2014): "Vertical integration and input flows," *The American Economic Review*, 104(4), 1120–1148.
- BAILY, M. N., C. HULTEN, AND D. CAMPBELL (1992): "The distribution of productivity in manufacturing plants," *Brookings Papers: Macroeconomics*, 187.
- BARNETT, A. H. (1980): "The Pigouvian tax rule under monopoly," *The American Economic Review*, pp. 1037–1041.
- BARTIK, T. J. (1991): "Who Benefits from State and Local Economic Development Policies?," W.E. Upjohn Institute.
- BECKER, R., W. GRAY, AND J. MARVAKOV (2013): "NBER-CES Manufacturing Industry Database: Technical Notes," Discussion paper, NBER-CES.
- BENTO, A. M. (2013): "Equity Impacts of Environmental Policy," *Annual Reviews of Resource Economics*, 5, 181–96.
- BESLEY, T., AND H. S. ROSEN (1999): "Sales taxes and prices: an empirical analysis," *National Tax Journal*, 52, 157–178.
- BLANCHARD, O. J., AND L. KATZ (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, pp. 1–75.
- BLUNDELL, R., AND S. BOND (2000): "GMM estimation with persistent panel data: an application to production functions," *Econometric reviews*, 19(3), 321–340.
- BONNET, C., P. DUBOIS, S. B. V. BOAS, AND D. KLAPPER (2013): "Empirical Evidence on the Role of Nonlinear Wholesale Pricing and Vertical Restraints on Cost Pass-Through," *Review of Economics and Statistics*, 95(2), 500–515.
- BRONS, M., P. NIJKAMP, E. PELS, AND P. RIETVELD (2008): "A meta-analysis of the price elasticity of gasoline demand. A SUR approach," *Energy Economics*, 30(5), 2105–2122.
- BUCHANAN, J. M. (1969): "External Diseconomies, Corrective Taxes, and Market Structure," *American Economic Review*, 59(1), 174–177.
- BULOW, J. I., AND P. PFLEIDERER (1983): "A note on the effect of cost changes on prices," *Journal of political Economy*, 91(1), 182–185.
- BUSHNELL, J., AND J. HUMBER (2017): "Rethinking Trade Exposure: The Incidence of Environmental Charges in the Nitrogenous Fertilizer Industry," *Journal of the Association of Environmental and Resource Economists*, 4(3), 857–894.

- CABRAL, M., M. GERUSO, AND N. MAHONEY (2015): “Does Privatized Health Insurance Benefit Patients or Producers? Evidence from Medicare Advantage,” Mimeo, UT Austin.
- CAMPA, J. M., AND L. S. GOLDBERG (2005): “Exchange rate pass-through into import prices,” *Review of Economics and Statistics*, 87(4), 679–690.
- CARLETON, T. A., AND S. M. HSIANG (2016): “Social and economic impacts of climate,” *Science*, 353(6304), aad9837.
- COLLARD-WEXLER, A., AND J. DE LOECKER (2015): “Reallocation and technology,” *American Economic Review*, 105(1), 131–171.
- DAVIS, S. J., C. GRIM, J. HALTIWANGER, AND M. STREITWIESER (2013): “Electricity Unit Value Prices and Purchase Quantities: U.S. Manufacturing Plants, 1963–2000,” *Review of Economics and Statistics*, 95(4), 1150–1165.
- DE LOECKER, J. (2011): “Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity,” *Econometrica*, 79(5), 1407–1451.
- DE LOECKER, J., AND P. K. GOLDBERG (2014): “Firm Performance in a Global Market,” *Annu. Rev. Econ.*, 6(1), 201–227.
- DE LOECKER, J., P. K. GOLDBERG, A. KHANDELWAL, AND N. PAVCNİK (2016): “Prices, Markups and Trade Reform,” *Econometrica*, 84(2), 445–510.
- DE LOECKER, J., AND F. WARZYŃSKI (2012): “Markups and Firm-Level Export Status,” *The American Economic Review*, 102(6), 2437–2471.
- DELIPALLA, S., AND M. KEEN (1992): “The comparison between ad valorem and specific taxation under imperfect competition,” *Journal of Public Economics*, 49(3)(351–367).
- DOMS, M., T. DUNNE, AND K. R. TROSKE (1997): “Workers, wages, and technology,” *The Quarterly Journal of Economics*, pp. 253–290.
- DOYLE, JR., J. J., AND K. SAMPHANTHARAK (2008): “200 Gas! Studying the Effects of a Gas Tax Moratorium Across State Borders,” *Journal of Public Economics*, 92, 869–884.
- DUGGAN, M., A. STARC, AND B. VABSON (2014): “Who Benefits when the Government Pays More? Pass-Through in the Medicare Advantage Program,” Working Paper 19989, National Bureau of Economic Research.
- EIA (2010): “2010 Manufacturing Energy Consumption Survey, Form EIA-846,” Discussion paper, EIA, Office of Energy Consumption and Efficiency Statistics.
- (2015): “Monthly Energy Review,” Discussion paper, EIA.
- FABRA, N., AND M. REGUANT (2014): “Pass-Through of Emissions Costs in Electricity Markets,” *American Economic Review*, 104(9), 2872–2899.
- FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2008): “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?,” *American Economic Review*, 98(1), 394–425.

- FOWLIE, M., M. REGUANT, AND S. P. RYAN (2016): “Market-Based Emissions Regulation and Industry Dynamics,” *Journal of Political Economy*, 124(1), 249–302.
- GENESOVE, D., AND W. P. MULLIN (1998): “Testing static oligopoly models: conduct and cost in the sugar industry, 1890-1914,” *RAND Journal of Economics*, 29(1), 355–377.
- GOLDBERG, P. K., AND R. HELLERSTEIN (2008): “A structural approach to explaining incomplete exchange-rate pass-through and pricing-to-market,” *The American Economic Review*, 98(2), 423–429.
- GOPINATH, G., P.-O. GOURINCHAS, C.-T. HSIEH, AND N. LI (2011): “International prices, costs, and markup differences,” *The American Economic Review*, 101(6), 2450–2486.
- GOULDER, L., AND M. HAFSTEAD (2017): *Confronting the Climate Challenge: U.S. Policy Options*. Columbia University Press.
- GOULDER, L. H., AND R. C. WILLIAMS III (2003): “The substantial bias from ignoring general equilibrium effects in estimating excess burden, and a practical solution,” *Journal of political Economy*, 111(4), 898–927.
- GRAINGER, C. A., AND C. D. KOLSTAD (2010): “Who Pays a Price on Carbon?,” *Environmental and Resource Economics*, 46, 359–376.
- HALL, R. E. (1986): “Market structure and macroeconomic fluctuations,” *Brookings papers on economic activity*, pp. 285–338.
- HARASZTOSI, P., AND A. LINDNER (2015): “Who Pays for the Minimum Wage?,” .
- HARBERGER, A. C. (1962): “The Incidence of the Corporation Income Tax,” *Journal of Political Economy*, 70(3), 215–240.
- HAUSMAN, C., AND R. KELLOGG (2015): “Welfare and Distributional Implications of Shale Gas,” *NBER Working Paper*, (w21115).
- HO, M. S., R. D. MORGENSTERN, AND J.-S. SHIH (2008): “Impact of carbon price policies on US industry,” .
- HUGHES, J. E., C. R. KNITTEL, AND D. SPERLING (2008): “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand,” *The Energy Journal*, 29(1), 113–134.
- JENKIN, F. (1872): “On the Principles which regulate the Incidence of Taxes,” *Proceedings of the Royal Society of Edinburgh*, 7, 618–631.
- KAHN, M. E., AND E. T. MANSUR (2012): “Do Local Energy Prices and Regulation Affect the Geographic Concentration of Employment? A Border Pairs Approach,” Mimeograph, UCLA.
- KATZ, M. L., AND H. S. ROSEN (1985): “Tax Analysis in an Oligopoly Model,” *Public Finance Quarterly*, 13, 3–19.
- KOPCZUK, W., J. MARION, E. MUEHLEGGGER, AND J. SLEMROD (2016a): “Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes,” *American Economic Journal: Economic Policy*, 8(2), 251–286.

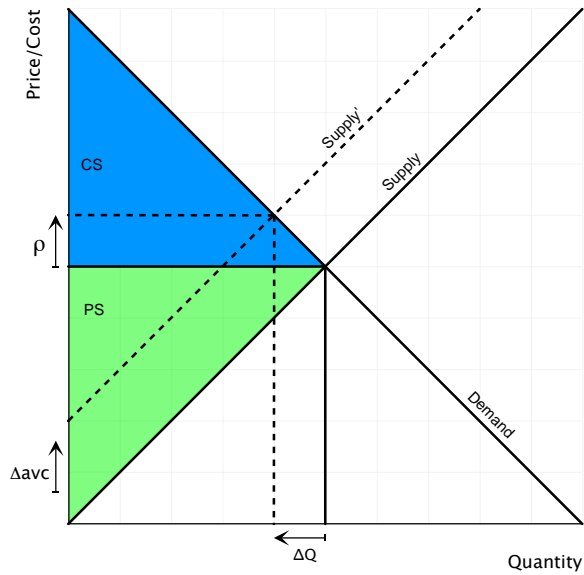
- (2016b): “Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes,” *American Economic Journal: Economic Policy*, 8, 2.
- LADE, G. E., AND J. BUSHNELL (2016): “Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard,” .
- LERNER, A. P. (1934): “The Concept of Monopoly and the Measurement of Monopoly,” *Review of Economic Studies*, 1, 157–175.
- LEVINSOHN, J., AND A. PETRIN (2003): “Estimating Production Functions Using Inputs to Control for Unobservables,” *Review of Economic Studies*, 70(2), 317–341.
- LI, S., J. LINN, AND E. MUEHLEGGGER (2014): “Gasoline Taxes and Consumer Behavior,” *American Economic Journal: Economic Policy*, 6(4), 302–342.
- LINN, J. (2008): “Energy Prices and the Adoption of Energy-Saving Technology,” *The Economic Journal*, 118(533), 1986–2012.
- LORIS, N., AND F. JOLEVSKI (2014): “EPA’s Climate Regulations Will Harm American Manufacturing,” Discussion paper, Heritage Foundation.
- MARION, J., AND E. MUEHLEGGGER (2008): “Measuring Illegal Activity and the Effects of Regulatory Innovation: Tax Evasion and the Dyeing of Untaxed Diesel,” *Journal of Political Economy*, 116(4), 633–666.
- (2011): “Fuel tax incidence and supply conditions,” *Journal of Public Economics*, 95(9), 1202–1212.
- MARSCHAK, J., AND W. H. ANDREWS (1944): “Random simultaneous equations and the theory of production,” *Econometrica, Journal of the Econometric Society*, pp. 143–205.
- MARTIN, R., L. DE PREUX, AND U. WAGNER (2014): “The Impact of a Carbon Tax on Manufacturing: Evidence from Microdata,” *Journal of Public Economics*, 117(1-14).
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71(6), 1695–1725.
- MILLER, N. H., M. OSBORNE, AND G. SHEU (2015): “Pass-Through in a Concentrated Industry: Empirical Evidence and Regulatory Implications,” Mimeo, Georgetown University.
- MORETTI, E. (2011): “Local labor markets,” *Handbook of labor economics*, 4, 1237–1313.
- MUEHLEGGGER, E., AND R. SWEENEY (2017): “Pass-Through of Input Cost Shocks Under Imperfect Competition: Evidence from the U.S. Fracking Boom,” NBER Working Paper No. 24025.
- NAKAMURA, E., AND D. ZEROM (2010): “Accounting for Incomplete Pass-Through,” *Review of Economic Studies*, 77, 1192–1230.
- NEWKEY, W. K., AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55(3), 703–708.
- NOTOWIDIGDO, M. J. (2011): “The incidence of local labor demand shocks,” Discussion paper, National Bureau of Economic Research.

- OLLEY, G. S., AND A. PAKES (1996): “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64(6), 1263–1297.
- PARRY, I. W., H. SIGMAN, M. WALLS, AND R. C. WILLIAMS (2006): *International Yearbook of Environmental and Resource Economics*chap. The Incidence of Pollution Control Policies, pp. 1–42. Edward Elgar.
- POTERBA, J. M. (1996): “Retail price reactions to changes in state and local sales taxes,” *National Tax Journal*, pp. 165–176.
- ROBERTS, M. J., AND D. SUPINA (1996): “Output price, markups, and producer size,” *European Economic Review*, 40(3-5), 909–921.
- (2000): *Advances in Applied Microeconomics, Vol 9, Industrial Organization*chap. Output Price and Markup Dispersion in Producer Micro Data: The Roles of Producer Heterogeneity and Noise. JAI Press.
- RYAN, S. (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80(3), 1019–1061.
- SCHMALENSEE, R. (1976): “Another Look at the Social Valuation of Input Price Changes,” *American Economic Review*, 66(1), 239–243.
- SEADE, J. (1985): “Profitable cost increases and the shifting of taxation: equilibrium response of markets in Oligopoly,” Discussion paper, University of Warwick, Department of Economics.
- STERN, N. (1987): “The Effects of Taxation, Price Control, and Government Contracts in Oligopoly and Monopolistic Competition,” *Journal of Public Economics*, 32, 133–158.
- STOLPER, S. (2017): “Local Pass-Through and the Regressivity of Taxes: Evidence from Automotive Fuel Taxes,” Mimeo, University of Michigan.
- STREETTER, E. (2009): “NAM, NFIB Launch Media Campaign Opposing Waxman-Markey Climate Change Bill,” Press Release.
- SYVERSON, C. (2004): “Market Structure and Productivity: A Concrete Example,” *Journal of Political Economy*, 112(6), 1181–1222.
- USEPA (2015): “Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2013,” Discussion paper, USEPA.
- VAN BIESEBROECK, J. (2007): “Robustness of Productivity Estimates,” *The Journal of Industrial Economics*, 55(3), 529–569.
- WEYL, E. G., AND M. FABINGER (2013): “Pass-through as an economic tool: Principles of incidence under imperfect competition,” *Journal of Political Economy*, 121(3), 528–583.


Figures

Figure 1: Pass-Through and Incidence under Perfect Competition

(a) Shift in Marginal Cost




(b) New Market Equilibrium




NOTES: See text for full description. The top panel shows consumer and producer surplus (denoted CS and PS, respectively) under general supply and demand curves in a market exhibiting perfect competition. The bottom panel shows consumer and producer surplus following a tax rate τ that changes average variables costs by Δavc , prices by ρ , and quantity sold by ΔQ .

Figure 2: Pass-Through and Incidence under Monopoly

(a) Shift in Marginal Cost




(b) New Market Equilibrium




NOTES: See text for full description. The top panel shows consumer surplus, producer surplus, and the deadweight loss (denoted CS, PS, and DWL respectively) in a market with a single monopolist. The MR line denotes marginal revenue and the MC line reflects an increasing marginal cost production function. The bottom panel shows consumer and producer surplus, as well as deadweight loss following a tax t that changes average variables costs by Δavc , prices by ρ and quantity sold by ΔQ .

Figure 3: Electricity Fuel Mix by Region


(a) Coal



(b) Natural Gas




(c) Petroleum



NOTES: These maps show the spatial distribution of electricity generation by fuel type by state averaged over our sample period, 1972-1997. Panel A shows the fraction of electricity generation in a state that comes from coal-fired generation; Panel B shows the fraction of electricity generation in a state that comes from natural gas; Panel C shows the fraction of electricity generation in a state that comes from petroleum.

Source: Energy Information Association, State Energy Data System.

Figure 4: National Fuel Prices, 1967-2012



NOTES: This figure plots a time series of national fuel prices from 1967 to 2012. Prices have been converted to real 2011 dollars using the consumer price index. Crude oil prices reflect the price of U.S. crude oil net imports in dollars per barrel. Natural gas prices reflect wellhead prices in dollars per thousand cubic feet. Coal prices reflect dollars per short ton. Prices for coal and crude oil have been divided by 10 to facilitate a common axis across the 3 fuels. Source: Energy Information Association, Annual Energy Review.

Tables

Table 1: Allocation of Energy Input Expenditures Across Fuels, by Industry

	Coal	Natural Gas	Fuel Oil	Electricity
Boxes	0.000	0.007	0.002	0.009
Bread	0.000	0.007	0.002	0.011
Cement	0.130	0.049	0.010	0.134
Concrete	0.000	0.002	0.011	0.007
Gasoline	0.000	0.014	0.817	0.007
Plywood	0.000	0.005	0.002	0.012

Notes: This table shows the average energy input percentage by industry for each of the four primary energy inputs into manufacturing production. These statistics are calculated by dividing the expenditure on each of the four energy inputs by the total annual expenditures (salary and wages; capital rental rates; materials; electricity; fuels) in the industry. Expenditures on energy inputs for gasoline refining include energy feedstocks. Source: Manufacturing Energy Consumption Survey, Census and Annual Survey of Manufacturers.

Table 2: Summary Statistics

	Energy Cost Share	Output Elasticities			Returns to Scale	Markup	Output Price	Marginal Costs	Materials Share of Revenue	Observations
		Labor	Materials	Capital						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Boxes	0.02	0.04	0.95	0.04	1.00	1.47	0.63	0.43	0.65	1414
Bread	0.02	0.28	0.63	0.09	1.13	1.20	0.47	0.39	0.53	248
Cement	0.32	0.91	1.08	0.19	2.46	2.30	0.05	0.02	0.46	229
Concrete	0.02	0.11	0.68	0.16	1.09	1.12	43.82	39.82	0.61	3369
Gasoline	0.84	0.01	0.99	0.03	1.02	1.11	23.08	20.84	0.89	284
Plywood	0.02	0.02	0.95	0.11	0.92	1.48	0.68	0.57	0.62	139
Sample Mean	0.02	0.10	0.70	0.14	1.09	1.15	—	—	0.62	5683

Notes: This table shows mean values of energy cost shares, output elasticities, and markups. An observation is a plant-year. Energy cost shares are the sum of fuel and electricity expenditures divided by total annual expenditures (salary and wages; capital rental rates; materials; electricity; fuels). Plant-level markups come from estimating production functions by industry using GMM-Proxy methods (Akerberg, Caves, and Frazer, 2015). Output elasticity estimates come from a three factor, gross-output, translog production function, where the inputs consist of labor, capital, and materials. Price and costs are measured in 1000s of 1987 dollars. Boxes are measured in short tons; bread is measured in thousands of pounds; cement is measured in cubic yards; concrete is measured in thousands of cubic yards; gasoline is measured in thousands of barrels; plywood is measured in thousands of square feet surface measure. See text for details. Source: Census of Manufacturers.

Table 3: Relationship Between Marginal Costs, Unit Prices, Markups and Electricity Input Prices

	(1)	(2)	(3)	(4)	(5)	(6)
	Lag (t-0)	Lag (t-2)	Lag (t-5)	Lag (t-0)	Lag (t-2)	Lag (t-5)
Panel A: Marginal Costs						
Coal Price \times Coal Share	0.092 (0.387)	0.156 (0.363)	-0.110 (0.311)	0.357 (0.244)	0.374 (0.293)	0.123 (0.255)
Gas Price \times Gas Share	0.779*** (0.175)	0.788*** (0.140)	0.866*** (0.191)	0.235*** (0.086)	0.225*** (0.084)	0.291*** (0.061)
Oil Price \times Oil Share	0.136 (0.341)	0.229 (0.290)	0.013 (0.201)	-0.070 (0.121)	0.047 (0.118)	-0.029 (0.139)
Panel B: Unit Prices						
Coal Price \times Coal Share	0.081 (0.259)	0.063 (0.254)	-0.061 (0.197)	0.259 (0.263)	0.159 (0.257)	0.065 (0.223)
Gas Price \times Gas Share	0.491*** (0.109)	0.502*** (0.088)	0.532*** (0.101)	0.186** (0.074)	0.204*** (0.061)	0.222*** (0.054)
Oil Price \times Oil Share	0.101 (0.181)	0.172 (0.168)	0.078 (0.131)	-0.008 (0.102)	0.057 (0.094)	0.079 (0.108)
Panel C: Markups						
Coal Price \times Coal Share	-0.012 (0.220)	-0.092 (0.183)	0.049 (0.171)	-0.098 (0.156)	-0.215 (0.172)	-0.058 (0.172)
Gas Price \times Gas Share	-0.288*** (0.080)	-0.286*** (0.066)	-0.334*** (0.096)	-0.049 (0.045)	-0.021 (0.053)	-0.069* (0.038)
Oil Price \times Oil Share	-0.035 (0.181)	-0.056 (0.134)	0.065 (0.097)	0.062 (0.090)	0.010 (0.069)	0.107 (0.077)
N	5892	5892	5892	5892	5892	5892
Plant FE	X	X	X	X	X	X
Year FE	X	X	X			
State Trends	X	X	X	X	X	X
Product-Year FE				X	X	X
Region-Year FE				X	X	X

Notes: This table presents regression coefficients from 18 separate regressions, 6 per panel and 1 per column. An observation is a plant-year. Panel A presents a set of regressions, regressing plant-level marginal costs on national fuel prices interacted with lagged shares of state-level electricity generation mix. Panels B and C present a similar set of regressions, using plant-level unit prices and plant-level markups as the dependent variable, respectively. The regression includes the uninteracted fuel prices and generation shares as controls (not reported). Columns (1) and (4) report results using contemporaneous electricity generation shares, columns (2) and (5) present results lagging generation shares by 2 years, and columns (3) and (6) present results lagging generation shares by 5 years. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census of Manufacturers, EIA-SEDS.

Table 4: Relationship Between Marginal Costs, Unit Prices, Markups and Energy Prices

	(1) Lag (t-0)	(2) Lag (t-2)	(3) Lag (t-0)	(4) Lag (t-2)
Panel A: Marginal Costs				
Coal Price \times Coal Share	38.54*** (7.91)	67.60*** (8.90)	38.69*** (7.95)	66.79*** (8.97)
Gas Price \times Gas Share	139.60*** (29.29)	100.65** (41.61)	136.44*** (29.32)	97.41** (40.44)
Oil Price \times Oil Share	358.58*** (23.31)	427.31*** (84.93)	354.57*** (22.63)	419.99*** (83.01)
Electricity Price \times Electricity Share	-69.48** (32.15)	118.43 (113.81)	-70.59** (32.25)	112.69 (116.56)
Panel B: Unit Prices				
Coal Price \times Coal Share	13.95*** (4.70)	31.68*** (5.75)	13.84*** (4.67)	31.06*** (5.76)
Gas Price \times Gas Share	76.65*** (22.60)	54.92** (24.18)	75.80*** (22.71)	53.09** (23.65)
Oil Price \times Oil Share	183.76*** (20.47)	251.67*** (53.38)	181.91*** (20.05)	247.12*** (52.57)
Electricity Price \times Electricity Share	-12.66 (18.75)	81.70 (69.18)	-12.18 (19.10)	78.83 (68.15)
Panel C: Markups				
Coal Price \times Coal Share	-24.59*** (4.79)	-35.92*** (5.34)	-24.85*** (4.83)	-35.73*** (5.33)
Gas Price \times Gas Share	-62.95** (25.89)	-45.73** (20.41)	-60.64** (25.59)	-44.31** (19.87)
Oil Price \times Oil Share	-174.82*** (17.73)	-175.64*** (33.45)	-172.66*** (17.15)	-172.87*** (32.35)
Electricity Price \times Electricity Share	56.88** (23.70)	-36.72 (47.93)	58.41** (23.64)	-33.85 (46.82)
N	5683	5683	5683	5683
Plant FE	X	X	X	X
Year FE	X	X		
State Trends	X	X	X	X
Region-Year FE			X	X

Notes: This table presents regression coefficients from 12 separate regressions, 4 per panel and 1 per column. An observation is a plant-year. Panel A presents a set of regressions, regressing plant-level marginal costs on national energy input prices interacted with lagged industry energy expenditure shares. Panels B and C present a similar set of regressions, using plant-level unit prices and plant-level markups as the dependent variable, respectively. The regression includes the uninteracted fuel prices and expenditure shares as controls (not reported). Columns (1) and (3) report results using contemporaneous electricity generation shares, and columns (2) and (4) present results lagging expenditure shares by 2 years. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census and Annual Survey of Manufacturers, MECS, and EIA-SEDS.

Table 5: Pass-Through Rate of Marginal Costs into Unit Prices: OLS

	(1)	(2)	(3)	(4)
Marginal Costs	0.598*** (0.018)	0.680*** (0.033)	0.598*** (0.018)	0.681*** (0.032)
N	5892	5892	5892	5892
Plant FE	X	X	X	X
State Trends	X	X	X	X
Year FE	X			
Product-Year FE		X		X
Region-Year FE			X	X

Notes: This table presents regression coefficients from 4 separate regressions, one per column. An observation is a plant-year. The dependent variable is the plant-level unit-price, and the independent variable is plant-level marginal cost. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census of Manufacturers, EIA-SEDS.

Table 6: Pass-Through Rate of Marginal Costs into Unit Prices: Instrumental Variables

	(1) Lag (t-0)	(2) Lag (t-2)	(3) Lag (t-5)	(4) Lag (t-0)	(5) Lag (t-2)	(6) Lag (t-5)
Panel A: Electricity Shift-Share Instrument						
Marginal Costs	0.628*** (0.031)	0.623*** (0.031)	0.625*** (0.029)	0.660*** (0.099)	0.654*** (0.088)	0.715*** (0.086)
N	5892	5892	5892	5892	5892	5892
First Stage F-Statistic	9.53	14.33	6.99	8.89	3.95	12.09
Plant FE	X	X	X	X	X	X
Year FE	X	X	X			
State Trends	X	X	X	X	X	X
Product-Year FE				X	X	X
Region-Year FE				X	X	X
Panel B: Fuel Shift-Share Instrument						
Marginal Costs	0.511*** (0.013)	0.514*** (0.014)	— —	0.505*** (0.012)	0.508*** (0.014)	— —
N	5683	5683	—	5683	5683	—
First Stage F-Statistic	137.48	78.68	—	120.74	66.18	—
Plant FE	X	X	X	X	X	X
Year FE	X	X	X			
State Trends	X	X	X	X	X	X
Region-Year FE				X	X	X

Notes: This table presents regression coefficients from 10 separate regressions, 6 regressions in Panel A and 4 regressions in Panel B. An observation is a plant-year. The dependent variable is the plant-level unit-price, and the independent variable is plant-level marginal cost. In Panel A, marginal cost is instrumented by the interactions between national fuel prices for electricity generation and lagged electricity generation shares. In Panel B, marginal cost is instrumented by the interactions between national fuel prices for industrial production and lagged industry energy expenditure shares. Columns (1) and (4) report results using contemporaneous electricity generation shares, columns (2) and (5) present results lagging generation shares by 2 years, and columns (3) and (6) present results lagging generation shares by 5 years. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census and Annual Survey of Manufacturers, MECS, EIA-SEDS.

Table 7: Pass-Through Rate of Marginal Costs into Unit Prices, by Product: Instrumental Variables

	(1) Boxes	(2) Bread	(3) Cement	(4) Concrete	(5) Gasoline	(6) Plywood
Panel A: Baseline - Electricity Price Instrument						
Marginal Costs	0.963*** (0.038)	0.681*** (0.150)	0.775*** (0.087)	0.711*** (0.082)	0.327** (0.143)	0.692*** (0.082)
N	1414	308	293	3369	345	163
Pass-Through Rate	1.42	0.82	1.78	0.80	0.36	1.02
First Stage F-Statistic	23.41	1.67	49.29	23.36	2.43	38.55
Plant FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
State-Trends FE	X	X	X	X	X	X
Panel B: Region-Year FE - Electricity Price Instrument						
Marginal Costs	0.992*** (0.057)	0.458*** (0.172)	0.801*** (0.100)	0.624*** (0.066)	0.242** (0.111)	0.758*** (0.068)
N	1414	308	293	3369	345	163
Pass-Through Rate	1.46	0.55	1.84	0.70	0.27	1.12
First Stage F-Statistic	2.04	3.39	22.90	13.5	4.22	30.10
Plant FE	X	X	X	X	X	X
Region×Year FE	X	X	X	X	X	X
State-Trends FE	X	X	X	X	X	X

Notes: This table presents regression coefficients from 14 separate regressions; one per column in each of the two panels. Each column represents a separate sample, where the sample is indicated in the column headings. An observation is a plant-year. The dependent variable is the plant-level unit-price, and the independent variable is plant-level marginal cost. Marginal cost is instrumented by the interactions between national fuel prices for electricity generation and 5-year lagged electricity generation shares. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census and Annual Survey of Manufacturers, MECS, EIA-SEDS.

Table 8: Incidence: Change in Consumer Surplus as Share of Change in Total Surplus

	(1)	(2)	(3)	(4)	(5)	(6)
	Boxes	Bread	Cement	Concrete	Gasoline	Plywood
Panel A: Incidence Components						
MC Pass-Through (ρ_{MC})	1.43	1.03	2.62	0.73	0.42	1.02
Demand Elasticity (ϵ_D)	2.76	5.23	2.90	4.28	0.13	1.93
Mean Lerner Index (L)	0.32	0.26	0.61	0.11	0.11	0.35
Panel B: Consumer Share of Burden (by Market Structure)						
Oligopoly	0.63 (0.03)	0.38 (0.17)	0.42 (0.09)	0.54 (0.07)	0.42 (0.22)	0.55 (0.18)
Monopoly	0.59 (0.02)	0.45 (0.17)	0.66 (0.17)	0.42 (0.05)	0.28 (0.12)	0.47 (0.16)
Perfect Competition	1.43 (0.12)	1.03 (0.71)	2.62 (2.04)	0.73 (0.16)	0.42 (0.23)	1.02 (0.53)

Notes: This table presents results for welfare incidence by industry, where incidence is defined as the change in consumer surplus as a share of the change in consumer and producer surplus. Columns (1) - (6) reflect separate calculations for industries defined in the column heading. Panel A displays the necessary components for measuring incidence using equation (2). Marginal cost pass-through (ρ_{MC}) is calculated as the average of the industry specific pass-through rates in Table 7. Panel B displays incidence under different market structures. The first two lines display incidence under arbitrary forms of oligopoly that next both perfect competition and monopoly. The third and fourth rows of Panel B display incidence under monopoly and perfect competition using equation (2) where $L = 1/\epsilon_D$ and $L = 0$ respectively. Numbers in parentheses reflect clustered-bootstrap standard errors, clustering by state. Source: Census and Annual Survey of Manufacturers, MECS, EIA-SEDS.

A Recovering Marginal Costs, Output Elasticities, and Production Functions - For Online Publication

A.1 Recovering Marginal Costs

This section describes the methodology for calculating marginal costs, following Hall (1986); De Loecker and Warzynski (2012). Let Q_{it} denote the physical output Q of plant i in year t . Output is a function of variable inputs V_{it} (i.e., those not subject to adjustment costs) like materials and energy; dynamic inputs D_{it} like capital or sticky labor, which are subject to adjustment costs; and plant-specific productivity Ω_{it} : $Q_{it} = Q_{it}(V_{it}, D_{it}, \Omega_{it})$. We assume a firm minimizes the cost of the variable input(s), conditioning on the dynamic inputs. The firm solves the following Lagrangian:

$$\mathcal{L}(V_{it}, D_{it}, \lambda_{it}) = P_{it}^V V_{it} + R_{it} K_{it} + \lambda_{it} [Q_{it} - Q_{it}(V_{it}, D_{it}, \Omega_{it})]$$

Here P_{it}^V is the price of variable inputs, R_{it} is the price of dynamic inputs, and λ_{it} is the Lagrange multiplier.

The firm's first-order condition for a variable input like materials is

$$\frac{\partial \mathcal{L}}{\partial V_{it}} = P_{it}^V - \lambda_{it} \frac{\partial Q_{it}(\cdot)}{\partial V_{it}}$$

Rearranging terms for an optimum where $\partial \mathcal{L} / \partial V_{it} = 0$ and multiplying the right-hand side by $V_{it} Q_{it} / V_{it} Q_{it}$ and both sides by P_{it} shows how we recover markups:

$$\frac{P_{it}}{\lambda_{it}} = \left[\frac{\partial Q_{it}(\cdot)}{\partial V_{it}} \frac{V_{it}}{Q_{it}} \right] \left[\frac{P_{it}^V V_{it}}{P_{it} Q_{it}} \right]^{-1} \quad (7)$$

The left-hand side of this equation is the multiplicative markup μ_{it} , which equals prices divided the Lagrange multiplier. The Lagrange multiplier represents marginal costs, since it reflects the costs of relaxing the output constraint. The right-hand side is the product of two bracketed terms; we construct empirical analogues to both. The first is the elasticity of output with respect to a variable input, or the ‘‘output elasticity.’’ We estimate the output elasticity from production functions, described in the next subsection. The second bracketed term is the cost of the variable input divided by the firm's revenue, or the ‘‘revenue share.’’ Our data report the revenue share of each input.

We can then compute a time-varying, plant-level markup by using the estimated output elasticity of a variable input and the revenue share of that input. Since we observe plant-level unit prices, we then recover marginal costs from the accounting identity that price equals markups times marginal costs: $MC_{it} = P_{it} / \mu_{it}$ where MC_{it} is the marginal cost of plant i in year t , and μ_{it} is the (multiplicative) markup.

A.2 Recovering Output Elasticities from Production Functions

The previous subsection showed that estimating markups requires the output elasticity of a variable plant-level input like materials. We estimate this output elasticity by using proxy methods to estimate production functions (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Akerberg, Caves, and Frazer, 2015). We focus on production functions with a scalar, Hicks-neutral productivity term and estimate elasticities separately by industry, assuming common technology across

firms and over time within an industry.⁴⁵ We show a Cobb-Douglas specification here to simplify exposition, though our results use a more flexible translog, gross-output production function:⁴⁶

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \omega_{it} + \epsilon_{it} \quad (8)$$

Throughout the paper, lowercase represents variables in logs. Here y_{it} represents a plant’s output quantity. We use output quantity rather than revenues here to avoid well-known bias in revenue-based productivity estimates.⁴⁷ Firms use three inputs: capital, labor, and materials (k_{it} , l_{it} , and m_{it}). Materials includes energy inputs in addition to other intermediate inputs used for production. The parameter vector which we estimate, $\beta \equiv (\beta_k, \beta_l, \beta_m)$, contains the output elasticities of these three inputs. The term ω_{it} represents productivity, which is known to the firm when making static input decisions but unobserved to the econometrician. The residual ϵ_{it} includes measurement error and unanticipated shocks to output.

Ordinary least squares estimates of equation (8) may suffer from omitted variables bias due to the unobserved productivity term ω_{it} (Marschak and Andrews, 1944). A firm observes its productivity, so input choices k_{it} , l_{it} , and m_{it} may depend on it, but productivity directly affects output, and data do not report it.

To address the possible omitted variable bias associated with OLS estimates of equation (8), we use control-function or proxy methods to control for the unobserved and omitted productivity term. Consider a general demand function for materials:⁴⁸ $m_{it} = m_t(k_{it}, l_{it}, \omega_{it})$. Assuming that $m_t(\cdot)$ is strictly monotonic in inputs, we invert this input demand equation to solve for productivity as a function of the observable inputs:

$$\omega_{it} = m_t^{-1}(m_{it}, k_{it}, l_{it})$$

This inversion provides a control function for productivity.⁴⁹

⁴⁵Output elasticities could in principle differ by industry and time period. Such flexible output elasticities are difficult to estimate with our data, however, since we have few years of data, require one lag to construct instruments, and have few observations for most industries.

⁴⁶Translog coefficients are the same across firms within an industry. Markups and output elasticities, however, differ across firms within an industry, because input demands differ across firms. This is an advantage of translog over Cobb-Douglas production functions, which would have the same output elasticity across firms within an industry. Under Cobb-Douglas, all variation in markups across firms within an industry would come from revenue shares.

⁴⁷Output quantity addresses the distinction between revenue and physical total factor productivity (Foster, Haltiwanger, and Syverson, 2008). Unobserved variation in input prices may also bias production function coefficients (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016). The homogeneity of our products potentially gives less scope for input price variation and associated bias (De Loecker and Goldberg, 2014). We have explored specifications that attempt to control for any remaining input price variation using a polynomial in the output price, and results are largely similar. Our dataset reports expenditures on inputs though not input quantities; we are unaware of any production function estimates in any setting using input quantity data for all inputs rather than price-deflated input expenditures.

⁴⁸We focus on materials as a variable input into the production function, where materials include both purchased intermediates as well as energy input expenditures. In theory, we could estimate a separate output elasticity for energy in the production function. In practice, adding a fourth input into a translog production function substantially increases the number of parameters to be estimated. With relatively small sample sizes and relatively few degrees of freedom, output elasticities become more sensitive and less robust. Moreover, many indirect energy input costs are embodied in material expenditures through feedstock purchases for example.

⁴⁹Inverting materials demand to recover productivity requires a one-to-one mapping between plant-level productivity and materials. This assumption fails if unobserved plant-level variables besides productivity drive changes in materials or if there is measurement error in materials. Alternative production function estimators, such as the dynamic panel methods developed by Blundell and Bond (2000) are not appropriate in our setting since we have few time periods to construct differences and lags. Some evidence suggests these may not be first-order concerns. Syverson (2004) finds robustness among producer TFP measures (and hence output elasticities) for one of our indus-

We apply this approach in two steps, following Akerberg, Caves, and Frazer (2015). The first step regresses plant output y_{it} on a function $\phi_t(\cdot)$ of observed inputs. This first step is designed to purge output data of measurement error and unanticipated shocks to output ϵ_{it} :

$$y_{it} = \phi_t(k_{it}, l_{it}, m_{it}) + \epsilon_{it}$$

We approximate $\phi_t(\cdot)$ using a polynomial expansion. We use estimates from this first step to calculate ϵ_{it} from

$$\hat{\epsilon}_{it} = y_{it} - \hat{\phi}_t(k_{it}, l_{it}, m_{it}) \quad (9)$$

where $\hat{\phi}$ contains the fitted values from this first step, and $\hat{\epsilon}_{it}$ are the residuals from this regression. Since ϵ_{it} contains measurement error and unanticipated shocks to production, we can use it to obtain a measure of output which is purged of both. After this first step, the only missing information needed to know the output elasticity vector β is the productivity vector ω_{it} . Given any candidate elasticity vector $\tilde{\beta}$, we can estimate productivity by manipulating equations (8) and (9) to get

$$\omega_{it}(\tilde{\beta}) = \hat{\phi}_{it} - \tilde{\beta}_k k_{it} - \tilde{\beta}_l l_{it} - \tilde{\beta}_m m_{it} \quad (10)$$

The second step selects the coefficient vector that best fits the data by relying on the law of motion for productivity. We follow Akerberg, Caves, and Frazer (2015) and assume that productivity follows a first-order Markov process.⁵⁰ We define productivity shocks ξ_{it} as the difference between productivity and the expectation of last period's productivity given last period's information set \mathbb{I}_{it-1} :

$$\xi_{it} = \omega_{it} - \mathbb{E}[\omega_{it} | \mathbb{I}_{it-1}]$$

where \mathbb{E} is the expectation operator. Equivalently, ξ_{it} represents the component of current productivity which was unexpected at time $t - 1$.

The second step estimates the production function coefficients using the assumption that this productivity innovation must be orthogonal to a set of current and lagged input demands d_{it} . We summarize these conditions as

$$\mathbb{E}[\xi_{it}(\beta)d_{it}] = 0 \quad (11)$$

With the translog production function we use for the empirical implementation, the vector d_{it} is

$$d_{it} = \{l_{it}, m_{it-1}, k_{it}, l_{it}^2, m_{it-1}^2, k_{it}^2, l_{it}m_{it-1}, l_{it}k_{it}, m_{it-1}k_{it}\}$$

These moments above are similar to those suggested by Akerberg, Caves, and Frazer (2015). They exploit the fact that capital and labor have adjustment costs, and that lagged capital and labor should not be correlated with the current productivity innovation. We use lagged rather than current materials to identify the materials coefficients since current material expenditures may

tries, ready-mixed concrete, with a specification incorporating idiosyncratic demand shocks. Van Biesebroeck (2007) also finds high TFP correlations across various measurement alternatives. Given the strong assumptions needed to estimate output elasticities, however, subsequent sections explore alternative methods to characterize incidence in the absence of information on output elasticities or markups and marginal costs.

⁵⁰We use the AR(1) process to derive a plausibly exogenous productivity shock ξ_{it} along the lines of Akerberg, Caves, and Frazer (2015). We have also allowed for the potential of additional lagged decision variables to affect current productivity outcomes (in expectation) in order to accommodate the concerns raised by De Loecker (2011) pertaining to the exogeneity of productivity process. For example, we have allowed productivity to depend on export status and the nonrandom exit of firms (De Loecker, 2011; Olley and Pakes, 1996). In practice, our output elasticity estimates are not particularly sensitive to these modifications.

react to contemporaneous productivity innovations. For lagged materials to be a valid instrument for current materials, input prices must to be correlated over time.

Finally, we use generalized method of moments to choose the production function coefficients β which minimize the moment conditions in equation (11). With translog production functions, the coefficients β combined with input data give the output elasticities $\hat{\theta}$.⁵¹

$$\hat{\theta}_{it} = \theta(\hat{\beta}, l_{it}, m_{it}, k_{it})$$

These output elasticities vary by plant and by year. This leads to plant-year variation in markups (i.e., different markups for each plant within an industry-year) that is driven by both changes in the plant-level revenue share for the variable input (e.g. materials) and changes in the mix of inputs used for production.

B Appendix Figures and Tables

Table B1: Pass-Through Rate of Marginal Costs into Unit Prices, by Product: Fuel-Mix Instrumental Variables

	(1)	(2)	(3)	(4)	(5)	(6)
	Boxes	Bread	Cement	Concrete	Gasoline	Plywood
	Fuel Shift-Share Instrument (with Year FE)					
Marginal Costs	0.946*** (0.031)	0.214 (0.443)	0.706*** (0.067)	0.743*** (0.098)	0.491*** (0.139)	0.825*** (0.076)
N	1414	248	229	3369	284	139
First Stage F-Statistic	12.60	4.25	18.50	15.22	4.73	77.18
Year FE	X	X	X	X	X	X
Plant FE	X	X	X	X	X	X
State-Trends FE	X	X	X	X	X	X

Notes: This table presents regression coefficients from 7 separate regressions; one per column. Each column represents a separate sample, where the sample is indicated in the column headings. An observation is a plant-year. The dependent variable is the plant-level unit-price, and the independent variable is plant-level marginal cost. Marginal cost is instrumented by the interactions between national fuel prices for industrial production and 2-year lagged industry energy expenditure shares. All regressions include plant fixed effects, year fixed effects, and state-specific trends. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census and Annual Survey of Manufacturers, MECS, EIA-SEDS.

⁵¹The estimated output elasticity for materials, for example, is $\hat{\theta}_{it} = \hat{\beta}_m + 2\hat{\beta}_{mm}m_{it} + \hat{\beta}_{lm}l_{it} + \hat{\beta}_{km}k_{it} + \hat{\beta}_{lmk}l_{it}k_{it}$.

Table B2: Pass-Through Rate of Marginal Costs into Unit Prices, in Levels, by Product: Instrumental Variables

	(1)	(2)	(3)	(4)	(5)	(6)
	Boxes	Bread	Cement	Concrete	Gasoline	Plywood
Panel A: Baseline - Electricity Price Instrument						
Marginal Costs	1.581*** (0.052)	0.029 (0.157)	1.585*** (0.219)	0.759*** (0.061)	0.461*** (0.132)	0.625*** (0.225)
N	1414	308	293	3369	345	163
First Stage F-Statistic	44.59	4.00	80.03	12.49	7.15	17.04
Plant FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
State-Trends FE	X	X	X	X	X	X
Panel B: Region-Year FE - Electricity Price Instrument						
Marginal Costs	1.597*** (0.111)	0.019 (0.096)	1.631*** (0.192)	0.717*** (0.064)	0.302** (0.123)	0.755*** (0.160)
N	1414	308	293	3369	345	163
First Stage F-Statistic	2.83	3.87	3.64	12.66	8.99	233.0
Plant FE	X	X	X	X	X	X
Region×Year FE	X	X	X	X	X	X
State-Trends FE	X	X	X	X	X	X

Notes: This table presents regression coefficients from 14 separate regressions; one per column in each of the two panels. Each column represents a separate sample, where the sample is indicated in the column headings. An observation is a plant-year. The dependent variable is the plant-level unit-price, and the independent variable is plant-level marginal cost. Marginal cost is instrumented by the interactions between national fuel prices for electricity generation and 5-year lagged electricity generation shares. Standard errors are in parentheses and are clustered by state. Regressions are weighted by Census sampling weights. ***, **, * denotes statistical significance at the 1, 5, and 10 percent levels, respectively. See text for details. Source: Census and Annual Survey of Manufacturers, MECS, EIA-SEDS.

Table B3: Plant-Level Incidence: Mean and Standard Deviation

	(1) Incidence Mean	(2) Incidence Standard Deviation	(3) Coefficient of Variation
Boxes	0.664	0.032	0.05
Bread	0.495	0.119	0.24
Cement	0.506	0.044	0.09
Concrete	0.605	0.092	0.15
Gasoline	0.364	0.034	0.09
Plywood	0.644	0.074	0.11

Notes: This table presents additional summary information pertaining to within-industry incidence heterogeneity. Column (1) presents the mean incidence for a given industry, where the mean is taken over plant-level incidence measures. Plant-level incidence measures were created using plant-level markup estimates, combined with plant-level pass-through rates and industry-level demand elasticities. Column (1) reflects the unweighted mean over these plant-level observations. Column (2) represents the standard deviation of these incidence estimates within an industry. Column (3) represents the coefficient of variation, which is defined as the ratio of the standard deviation of a sample relative to its mean. Thus, Column (3) is created by dividing Column (2) by Column (1).

Table B4: Demand Elasticity Estimates

	(1) Boxes	(2) Bread	(3) Cement	(4) Concrete	(5) Gas	(6) Plywood
Panel A: OLS						
Demand Elasticity (ϵ_D)	-0.377** (0.121)	-0.273 (0.211)	-0.387 (0.286)	-0.657 (0.505)	-0.0454 (0.0748)	0.00469 (0.196)
Panel B: Productivity IV Estimates						
Demand Elasticity (ϵ_D)	-2.762** (0.894)	-5.233 (9.187)	-2.902** (1.054)	-4.275* (1.980)	-0.131 (0.111)	-1.926* (0.820)
N	100	25	25	25	25	50
First Stage F-Statistic	11.71	0.267	9.312	5.209	8.673	8.181
Year Trend	X	X	X	X	X	X

Notes: This table presents 12 separate regressions, 6 per panel. An observation is the yearly change in an industry-year, where the dependent variable in all regressions is $\Delta \log(\text{quantity})$. The independent variable is $\Delta \log(\text{output price})$. Panel A presents OLS estimates, separately by industry. Panel B presents estimates where price is instrumented with changes in industry level total factor productivity. Total factor productivity is constructed using a quantity-based productivity index. The index is constructed by subtracting log inputs from log outputs using industry-level cost shares as proxies for output elasticities. We use capital, materials, labor, and energy inputs, where capital, materials, and energy are deflated by industry-year input price deflators, and labor is measured in production hours. Standard errors are computed using Newey and West (1987). Source: NBER-CES Manufacturing Database.