

Pass-Through in a Concentrated Industry: Empirical Evidence and Policy Implications*

Nathan H. Miller[†]
Georgetown University

Matthew Osborne[‡]
University of Toronto

Gloria Sheu[§]
Department of Justice

October 6, 2014

Abstract

We estimate pass-through with nearly 40 years of data from a concentrated industry. The empirical model accounts for oligopoly interactions and allows us to distinguish the effects of plant-specific and industry-wide cost changes, even though the available price data are aggregated to the regional level. Our results indicate that industry pass-through is complete, regardless of competitive conditions, but that the pass-through of firm-specific cost changes is incomplete and decreases in the degree of competition. The industry in question, portland cement, is a major focus of environmental and antitrust policy-makers. We use our pass-through estimates to (i) calculate that consumers would bear 80% of the burden of market-based CO₂ regulation; (ii) corroborate the simulation-based analysis of the EPA regarding its impending regulation of local pollutants; and (iii) predict the magnitude and location of price effects that would arise due to a merger currently under review by antitrust authorities.

Keywords: pass-through, portland cement, environmental policy, antitrust, merger enforcement, regulation

JEL classification: L1, L5, L6

*We thank Erika Lim, Conor Ryan, and Jiayi Zhang for excellent research assistance, and Glen Weyl for useful comments. The views expressed herein are entirely those of the authors and should not be purported to reflect those of the U.S. Department of Justice.

[†]Georgetown University, McDonough School of Business, 37th and O Streets NW, Washington DC 20057. Email: nhm27@georgetown.edu.

[‡]University of Toronto, Rotman School of Management, 105 St. George St., Toronto, ON, Canada M5S 3E6. Email: matthew.osborne@rotman.utoronto.ca

[§]Department of Justice, Antitrust Division, Economic Analysis Group, 450 5th St. NW, Washington DC 20530. Email: gloria.sheu@usdoj.gov.

1 Introduction

In recent decades, there has been a proliferation of theoretical research that extends the principles of incidence elucidated by Marshall (1890) to imperfectly competitive environments. In public economics, this effort has generated results on the distributional impacts of taxation (e.g., Delipalla and Keen (1992); Anderson, de Palma, and Keider (2001); Weyl and Fabinger (2013)). Within international trade, Krugman (1986) and Dornbusch (1987) show that market power can explain incomplete exchange rate pass-through, which in turn has motivated new theoretical work (e.g., Melitz and Ottaviano (2008); Atkeson and Burstein (2008); Berman, Martin, and Mayer (2012); Auer and Schoenle (2013)). Further, within industrial organization, it is now understood that pass-through is central to a wide range of economic analyses, from price discrimination (e.g., Aguirre, Cowan, and Vickers (2010)) to the unilateral effects of mergers (e.g., Jaffe and Weyl (2013)).¹

The empirical literature has not kept pace, with some notable exceptions, in providing estimates of pass-through that account for oligopoly interactions. The dearth of research is unfortunate because policy debates increasingly involve concentrated markets in which strategic interactions play an important role. An empirical understanding of pass-through in such markets, if grounded in theory, could inform policy decisions.

In this paper, we develop estimates of pass-through based on nearly 40 years of price data from a concentrated industry, and use the results to conduct counter-factual experiments of immediate relevance to environmental and antitrust policy. We develop a general empirical model of oligopoly interactions that allows us to disentangle the effects of firm-specific cost changes from those that are industry-wide, even though the available price data are aggregated to the regional level. We then apply new Bayesian regression techniques to estimate the model in a way that fully preserves its microfoundations, including autocorrelation and spatial correlations in plant-level pricing residuals. Our results indicate that industry pass-through is complete, regardless of competitive conditions, but that the pass-through of firm-specific cost changes is incomplete and decreases in the degree of competition. Cross pass-through effects – how firms adjust prices with competitors’ costs – account for this divergence in industry pass-through and own pass-through.²

¹Many of these findings, across fields, are consolidated within the theoretical frameworks developed in Weyl and Fabinger (2013).

²How competition affects pass-through has received substantial theoretical attention. Ambiguity arises here due to the LeChatelier Principle (e.g., Samuelson (1947); Milgrom and Roberts (1996)). For example, in Cournot models with strategic substitutes, own pass-through converges to zero as the number of firms grows large, while industry pass-through converges to unity from below or above, depending on the curvature of demand (Bergstrom and Varian (1985); ten Kate and Niels (2005)). However, with symmetric differenti-

The industry in question, portland cement, is a major focus of environmental policy-makers because it accounts for roughly five percent of global anthropogenic CO₂ emissions (Van Oss and Padovani (2003)) and is also a major source of local pollutants such as particulate matter and mercury (e.g., EPA (2009); EPA (2010)). These impacts have motivated academic research that models the effects of environmental regulation taking into account the market power held by cement firms (e.g., Ryan (2012); Fowlie, Reguant, and Ryan (2014)). The prospect that this market power could be enhanced through the recently proposed merger of Holcim and Lafarge – two of the industry’s largest manufacturers – also places the cement industry at the forefront for antitrust policy-makers. Against this backdrop, we use our pass-through estimates to analyze (i) the market-based regulation of carbon dioxide emissions, (ii) recent EPA action to reduce local hazardous air pollutants, and (iii) the magnitude and geographic location of merger price effects.

The setting conveys advantages in estimation. In the production of cement, fossil fuels are burned in order to create extreme kiln temperatures. The procurement of fuel comprises a substantial fraction of overall variables costs and revenues, and there are no viable substitutes for fuel in the production process. Identification is supported by plant-specific variation in fuel costs that arises from inter-temporal changes in fossil fuel prices paired with observable heterogeneity in kiln efficiency. Lastly, differentiation in the industry is predominately spatial in nature, due to regulatory standards that govern the production process and the substantial cost of transportation (e.g., Miller and Osborne (2014)). This allows us to evaluate how pass-through varies with the number and proximity of nearby competitors.

The empirical model is quite general and could be applied in other economic environments, including those where the sources of differentiation are not spatial. The starting point is a linear approximation to the equilibrium price of each plant. The approximation aggregates cleanly to the regional-level, such that pass-through can be obtained by regressing regional prices on fuel cost variables, provided that the fuel cost variables are constructed in a manner that preserves the plant-level microfoundations. We estimate the model with ordinary least squares (OLS), a feasible generalized least squares (FGLS) estimator that accounts for autocorrelation at the region-level, and a Bayesian estimator that allows for autocorrelation and spatial correlation in the underlying plant-level error terms. We incorporate controls for demand and cost conditions, as well as plant and year fixed effects. Identification survives aggregation in our application due to the substantial amount

ated Bertrand competition, oligopoly interactions increase both own and industry pass-through (Weyl and Fabinger (2013)).

of cross-sectional and time-series variation that exists among the region-year observations. The obtained regression coefficients are precisely estimated and robust across a range of specifications.

We apply the aforementioned estimation results to inform three policy questions. First, we examine how market-based regulation of CO₂ would affect producer and consumer surplus, following the principles of incidence derived in Weyl and Fabinger (2013). We calculate that consumers would bear 80 percent of the burden of regulation, given assumptions on margins and demand elasticities that are based on the empirical literature.³ We show how regulation affects plants differentially, and map the geographic dispersion of price effects across the United States. Our calculations have direct bearing on the political economy of regulation, and specifically on the question of whether cap-and-trade permits should be allocated (or “grandfathered”) to incumbent producers, free of charge. Because consumers appear likely to bear the greater burden of regulation, we conclude that it would be appropriate to auction permits from the outset and disburse the revenues broadly. The amount of monies in question is substantial: based simply on 2012 production levels, pricing CO₂ at \$40 per metric tonne would raise more than \$2.5 billion from the cement industry alone.⁴

Second, we evaluate recent regulation promulgated by the EPA to reduce emissions of hazardous air pollutants (HAPs), including particulate matter, mercury, hydrocarbons, and hydrogen chloride. The regulation is scheduled to take effect in September 2015 after more than two years of litigation and renegotiation. The EPA indicates that the monetized health benefits of regulation outweigh economics costs (EPA (2009); EPA (2010)), the latter of which can be first order in oligopoly models with market power (e.g., Buchanon (1969)). Our analysis corroborates the simulation results developed by the EPA. The pass-through estimates imply average price increases of \$4.49 across 20 local markets, relative a simulation average of \$4.66. Further, there is a high degree of correlation in the predictions market-by-market. We believe this is attributable to a fortuitous choice of functional forms in the EPA simulation model, and would not occur with simulations generally.

³We believe that our calculation are likely to *understate* the proportion of burden that falls on consumers because (i) we consider only short run effects, and (ii) we assume that regulation does not affect importers. We discuss both of these considerations in the text.

⁴The preliminary publications of the USGS indicate that the industry produced 74 million metric tonnes of portland cement in 2012. We use standard methods to obtain CO₂ emissions per metric tonne of clinker. Our work here complements the research of Fowlie, Reguant, and Ryan (2014), which examines the effects market-based regulation of CO₂ on abatement and welfare using a dynamic structural model. Our pass-through estimates inform directly how the burden of regulation is distributed across producers and consumers, whereas this split is largely predetermined in their structural model. This enables us to better address questions related to the appropriate allocation of revenues.

The third counterfactual analysis is of the recently proposed merger of Holcim and Lafarge, which currently is under review by the antitrust authorities. We believe our work is the first academic application of first order approximation to calculate merger price effects, following the theoretical insights of Jaffe and Weyl (2013). Our calculations indicate price elevations of 3%-7% at many plants, based on an industry snapshot in 2010, the final year of our data. However, accounting for the merging firms' recent plant divestitures and closures, which just postdate our sample, eliminates most of these effects. Remaining price increases are relatively modest, arise at only a handful of plants, and affect customers predominately in the Northeast and Great Plains. These price elevations likely could be eliminated through the divestiture of two plants, one for each geographic area.

Our research substantially advances the empirical literature on pass-through. The reduced-form research on pass-through can be classified as follows.⁵ To start, a handful of articles, arising from industrial organization, examine the relationship between firm-specific prices and costs, often in retail markets (e.g., Ashenfelter, Ashmore, Baker, and McKernan (1998); Peltzman (2000); Besanko, Dube, and Gupta (2005)) or in other settings with unusually rich data (e.g., Besanko, Dranove, and Shanley (2001); Fabra and Reguant (2014)). These articles typically do not allow for cross pass-through effects in estimation. This can create discord between the empirical estimate and theoretical notions of pass-through, with the estimate falling somewhere between own pass-through and industry pass-through.⁶ Indeed, this is precisely what happens in our application when we neglect cross pass-through effects. The effect of competition on pass-through is not a focus of these articles.⁷

Another set of articles examines the relationship between market prices and costs, exploiting variation from sources such as exchange rates (e.g., Campa and Goldberg (2005); Gopinath, Gourinchas, Hsieh, and Li (2011)), sales taxes (e.g., Barzel (1976); Poterba (1996); Besley and Rosen (1998); Marion and Muehlegger (2011)), input prices (e.g., Borenstein,

⁵A structural approach also is possible: a number of articles infer firm-specific pass-through from structural models of supply and demand (e.g., Villas-Boas (2007); Hellerstein (2008); Nakamura and Zerom (2010); Bonnet, Dubois, and Villas-Boas (2013); Golberg and Hellerstein (2013)). Because inference is conditional on the correct specifications of supply and demand, including both the first and second order properties, this methodology is most valuable when empirical price-cost variation is insufficient to identify pass-through.

⁶This statement has a simple econometric intuition: competitor's costs, which are omitted variables in these articles, tend to be positively correlated with the included cost measure. Thus the estimate overstates own pass-through but, unless the correlation is perfect, the estimate understates industry pass-through. We are aware of two articles that incorporate cross pass-through effects: Ashenfelter, Ashmore, Baker, and McKernan (1998) and Besanko, Dube, and Gupta (2005).

⁷However, Besanko, Dube, and Gupta (2005) finds that retail pass-through typically increases in the market share of the product, and Fabra and Reguant (2014) finds that the pass-through of electricity generation costs is higher during peak hours. Both of these results are consistent with our empirical result that competition reduces own pass-through.

Cameron, and Gilbert (1997) Genesove and Mullin (1998); Nakamura and Zerom (2010)), and interest rates (e.g., Scharfstein and Sunderam (2013)). These articles develop useful insights, but do not inform firm-specific pass-through rates, which are objects of obvious interest in concentrated markets. Again, the effect of competition on pass-through is not a point of focus.⁸

By contrast, several recent articles in the international trade literature emphasize the role of competition in determining pass-through. The standard result is that pass-through is negatively correlated with indicators of market power (e.g., Berman, Martin, and Mayer (2012); Amiti, Itskhoki, and Konings (2012) Hong and Li (2013)). However, similar to the reduced-form literature in industrial organization, this research does not account for oligopoly interactions in estimation, and therefore may conflate how competition affects own pass-through with how it affects industry pass-through. Auer and Schoenle (2013) is an exception, as they develop empirical results showing that competition affects industry pass-through through both own pass-through and cross pass-through, and that those two channels partially offset each other.⁹ This parallels our own findings.

The paper proceeds as follows. Section 2 sketches some relevant institutional details of portland cement markets and describes the data that support our empirical work. Section 3 presents the empirical model, details how we measure fuel costs, and provides summary statistics on selected regressors. Section 4 discusses the estimation strategy and identification. Section 5 presents the regression results, Section 6 provides the policy analysis, and Section 7 concludes.

2 The Portland Cement Industry

2.1 Production technology

Portland cement is a finely ground dust that forms concrete when mixed with water and coarse aggregates such as sand and stone. Concrete, in turn, is an essential input to many construction and transportation projects. The production of cement involves feeding limestone and other raw materials into rotary kilns that reach peak temperatures of 1400-1450° Celsius. Exposure to extreme heat transforms the raw materials into clinker – small semi-fused modules, usually 3-25 millimeters in diameter – that subsequently is mixed with a

⁸The exception here is Scharfstein and Sunderam (2013), which finds evidence of lower industry pass-through in concentrated financial markets.

⁹Auer and Schoenle (2013) refer to own pass-through as the “direct cost pass-through component” and refer to cross pass-through as the “indirect price complementarity component.”

Table 1: Plants and Kilns in the Cement Industry

	Number of Plants	Number of Kilns	Total Capacity	Capacity Share			
				Wet Kilns	Long Dry Kilns	Dry with Preheater	Dry with Precalciner
1975	157	396	79,938	57%	32%	11%	1%
1980	142	319	77,100	49%	27%	16%	8%
1990	109	208	72,883	32%	23%	19%	27%
2000	107	196	82,758	24%	20%	17%	39%
2010	101	153	103,482	8%	9%	14%	70%

Notes: Total capacity is in thousands of metric tonnes. All data are for the contiguous United States and are obtained from the PCA Plant Information Survey.

small amount of gypsum and ground to form portland cement. Kilns operate at peak capacity except for one or two maintenance periods a year, the duration of which can be adjusted according to demand conditions. The energy and labor costs associated with shutting down and restarting kilns are substantial.

Capital investments over the last forty years have increased the industry’s capacity and productive efficiency. Table 1 provides snapshots of the industry over 1975-2010. The number of plants falls from 157 to 101 and the number of kilns falls from 394 to 151. Total industry capacity, though, increases from 79 million metric tonnes per year to more than 100 million tonnes as older wet kilns are retired and replaced with higher-capacity dry kilns.¹⁰ Most cement now is produced in dry kilns equipped with gas-suspension preheaters and precalciners. This auxiliary equipment uses exhaust gases to heat the raw material before it enters the rotary kiln, allowing for calcination, one of the major chemical reactions required in clinker production, to occur partially or fully outside the rotary kiln. The process is supplemented with an additional combustion chamber if a precalciner is present.

The energy demands of cement production are substantial because plants burn fossil fuels to produce extreme kiln temperatures.¹¹ Figure 1 plots in Panel A the fraction of industry capacity that uses bituminous coal, petroleum coke, fuel oil, and natural gas as its primary fuel source. At the outset of the sample, capacity is roughly split between coal and natural gas, with much less capacity using oil. Coal displaces natural gas, and, by then end

¹⁰Wet and dry kilns differ in how the raw material is prepared. With wet kilns, the raw material is wet-ground with water to form a slurry whereas, with dry kilns, the raw material is dry-ground to form a powder. Extra fuel is required in the wet process to evaporate the added water.

¹¹Plants also use electrical power to grind raw materials into the kiln feed mix, to grind the clinker output into cement, and to power the fans that blow exhaust gases from the kiln through preheaters.

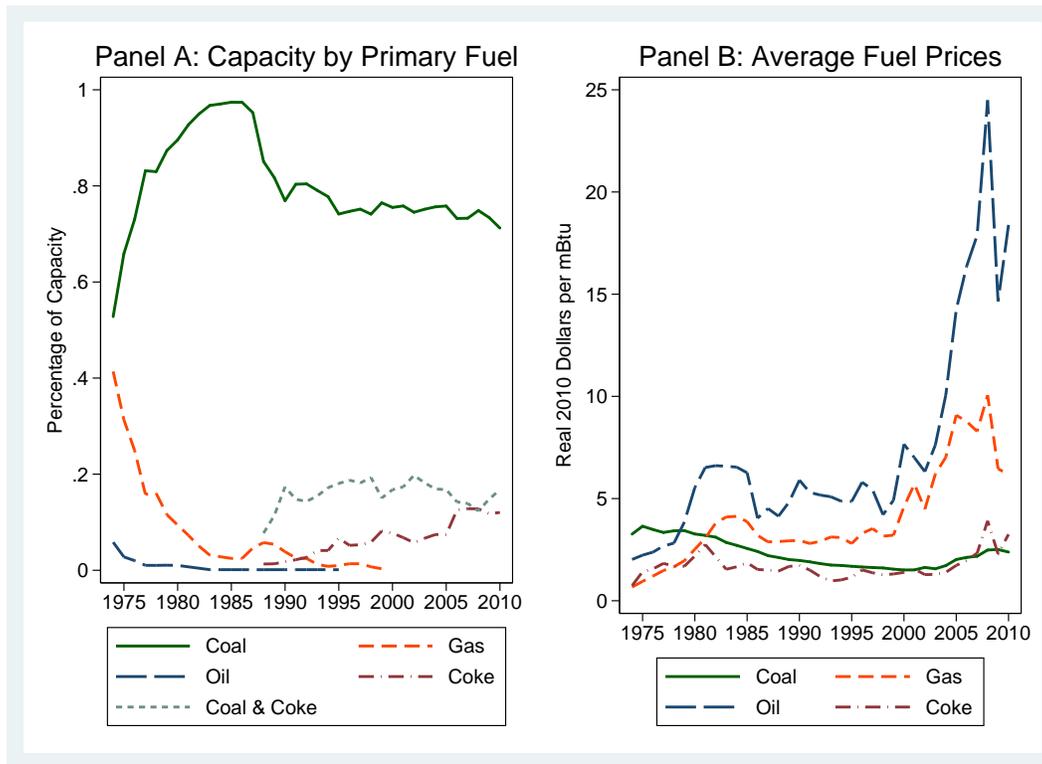


Figure 1: Primary Fuels and Fuel Prices

Notes: Panel A plots the fraction of kiln capacity that burns as its primary fuel (i) bituminous coal, (ii) natural gas, (iii) fuel oil, (iv) petroleum coke, and (v) bituminous coal and petroleum coke. Data are obtained from the PCA Plant Information Surveys. Panel B plots the average national prices for these fuel in real 2000 dollars per mBtu. Coal prices are obtained from the Coal Reports of the Energy Information Agency (EIA); the remaining prices are obtained from the State Energy Data System of the EIA.

of the sample, all operational kilns are heated with coal, petroleum coke, or a mix of the two. Panel B shows the cause of this shift: natural gas and oil prices increase relative to those of bituminous coal and petroleum coke in the late 1970s and, for the remainder of the sample period, coal and coke are more economical on a per mBtu basis. The variation in fuel choices and fuel prices, together with the heterogeneous kiln technologies, produces variation in fuel costs that we exploit in the empirical analysis.

Manufacturers of cement sell predominately to ready-mix concrete plants and large construction firms. Contracts are privately negotiated and relatively short term (often around one year in duration). They specify a free-on-board price at which cement can be obtained from the plant and discounts that reflect the ability of the customer to access the cement of competing manufacturers.¹² Most cement is trucked directly from the plant to the cus-

¹²While some cement manufacturers are vertically integrated into ready-mix concrete markets, Syverson and Hortaçsu (2007) determine that this has little impact on plant- and market-level outcomes.

Table 2: Cement Prices and Fuel Costs

	Cement Price	Fuel Costs				
		All Kilns	Wet Kilns	Long Dry Kilns	Dry with Preheater	Dry with Precalciner
1975	117.43	29.91	31.73	28.07	22.77	.
1980	130.33	27.88	31.52	25.47	22.27	18.67
1990	83.59	12.79	15.41	13.22	10.25	10.24
2000	100.86	8.91	11.09	9.63	7.54	7.29
2010	91.40	12.66	17.90	15.64	12.18	11.72

Notes: Cement prices and fuel costs are in real 2010 dollars per metric tonne. The cement price data are obtained from the USGS Minerals Yearbooks. Fuel costs are calculated by the authors and reflect the cost of marginal output.

to mer, though some cement is transported by barge or rail first to distribution terminals and only then trucked to customers. Transportation accounts for a substantial portion of purchasers' total acquisition costs because because portland cement is inexpensive relative to its weight. Miller and Osborne (2014) estimate transportation costs to be \$0.46 per tonne-mile, and determine that these costs create market power for spatially differentiated plants. Accordingly, the academic literature commonly models the industry using a number of geographically distinct local markets (e.g., Ryan (2012); Fowlie, Reguant, and Ryan (2014)). Aside from spatial considerations, cement is viewed as a commodity.

Table 2 shows average cement prices and fuel costs over 1975-2010, again using snapshots. The prices are obtained from published USGS data, while the fuel costs are based on calculations that we detail later. As shown, fuel costs decline from \$29.91 to \$12.66 per metric tonne, due to falling fossil fuel prices and the replacement of wet kilns with more fuel efficient dry kilns. Prices also trend down. Indeed, the national average prices and fuel costs over the sample period have a tight linear relationship, with a univariate correlation coefficient is 0.81, and a univariate OLS regression of the national average price on the national average fuels costs yields a coefficient of 1.48. This evidences high industry pass-through rates, though strict theoretical interpretation of the univariate coefficient is problematic.

2.2 Data sources

We draw data from numerous sources. Chief among these is the *Minerals Yearbook*, an annual publication of the United States Geological Survey (USGS), which summarizes a census of portland cement plants. The census response rate is typically well over 90 percent (e.g., 95

percent in 2003), and USGS staff imputes missing values for the few non-respondents based on historical and cross-sectional information. The *Minerals Yearbook* has long supported academic research on cement markets (e.g., McBride (1983); Jans and Rosenbaum (1997); Syverson and Hortaçsu (2007); Ryan (2012); Miller and Osborne (2014); Fowlie, Reguant, and Ryan (2014)). A defining characteristic of the *Minerals Yearbook* is that the data are aggregated to protect the confidentiality of census respondents. The price data we use reflect the average free-on-board price obtained by plants located in distinct geographic regions. The regions are not intended to approximate local markets, and the USGS frequently redraws boundaries to ensure that each region includes at least three independently owned plants.¹³ The *Minerals Yearbook* also contains non-price information, including production by region and consumption by state.

Our second source of data is the *Plant Information Survey*, an annual publication of the Portland Cement Association (PCA), which provides information on the plants and kilns in the United States. It also is featured frequently in academic research on cement markets. We obtain the location, owner, and primary fuel of each plant, as well as the annual capacity of each rotary kiln and the type of technology employed. Unlike the *Minerals Yearbook*, which summarizes economic activity over the year, the *Plant Information Survey* is a snapshot of the industry as it exists on December 31. In total, there are 4,416 plant-year observations over 1974-2010, of which 4,361 are active and 55 are idle. We also make use of the PCA's *U.S. and Canadian Portland Cement Labor-Energy Input Survey*, which is published intermittently and contains information on the energy requirements of clinker production and the energy content of fossil fuels burned in kilns. We have data for 1974-1979, 1990, 2000, and 2010.

We obtain data on the national average delivered bituminous coal price in the industrial sector over 1985-2010 from the annual *Coal Reports* of the Energy Information Agency (EIA). We backcast these prices to the period 1974-1984 using historical data on national average free-on-board prices of bituminous coal published in the 2008 *Annual Energy Review* of the EIA. We provide details on backcasting in Appendix A. We obtain national data on the prices of petroleum coke, natural gas, and distillate fuel oil, again for the industrial sector, from the State Energy Database System (SEDS) of the EIA.¹⁴ We obtain data on the national

¹³This “rule of three” prevents any one firm from backward engineering the business data of its competitors.

¹⁴The SEDS also includes data on coal prices, but no distinction is made between bituminous coal, sub-bituminous coal, lignite, and anthracite, despite the wide price differences that arise between those fuels. We also obtain state-level data on fossil fuel prices. There are many missing values at that level of aggregation, and we impute them as described in Appendix A. When included together in regressions, fuel cost variables based on national fossil fuel prices dominate fuel cost variables based on state-level prices. This could be a statistical artifact due to noise introduced by the imputation of missing values in the state-level data.

average price of unleaded gasoline over 1974-2010 from the Bureau of Labor Statistics, in order to better model the spatial configuration of the industry. We convert this series to an index that equals one in 2000. Lastly, to help control for demand, we obtain county-level data from the Census Bureau on construction employees and building permits. We provide details on data sources and related topics in Appendix A.

3 The Empirical Framework

Our econometric objective is to determine how the fuel costs of each portland cement plant affect both the prices obtained by the plant and the prices obtained by the plant’s competitors. The primary obstacle is that the available price data are aggregated to the regional level. We therefore build a model of regional prices that has realistic microfoundations. In this model, regional prices are a function of suitably aggregated plant fuel costs. We use data on kiln efficiency and fossil fuel prices to measure the fuel costs of each plant. The regression of regional prices on these fuel costs, once aggregated, obtains estimates of the average pass-through rates in the range of the data.

3.1 Modeling regional prices

We take as given that single-plant cement firms set free-on-board prices according to some pricing function that can be conceptualized as the equilibrium strategy for a consumer demand schedule and a competitive game. The product of each cement plant is differentiated due to geographic dispersion and transportation costs. Let there be $j = 1 \dots J_t$ cement plants in period t and let c_{jt} denote fuel costs per unit of output. A linear approximation to the equilibrium price of plant j is given by

$$p_{jt} = \rho_{jjt}c_{jt} + \sum_{k \neq j} \rho_{jkt}c_{kt} + x'_{jt}\gamma + \mu_j + \tau_t + \epsilon_{jt}, \quad (1)$$

where x_{jt} includes observable demand and cost variables, μ_j and τ_t are plant and year fixed effects, respectively, and ϵ_{jt} is a pricing residual that summarizes unobservable demand and cost conditions. The fuel cost coefficients are linear approximations to own and cross pass-through. Industry pass-through is $\rho_{jt}^M = \sum_k \rho_{jkt}$. We include among the controls nearby construction employment and building permits (which account for demand), indicators for the technology of the plant and the technology of nearby competitors (which account for non-fuel cost differences between kilns), and nearby competitor capacity.

Equation (1) is quite general but cannot be estimated, even with plant-level data, because the number of pass-through terms exceeds the number of observations. We impose restrictions on pass-through in order to facilitate estimation, leveraging the reasonable assumption that cross pass-through is greater between plants that are closer competitors.¹⁵ In particular, we construct a “distance metric” that summarizes the closeness of competition and impose that, for plants $j \neq k$, cross pass-through is given by

$$\rho_{jkt} = \begin{cases} \beta/d_{jkt} & \text{if } j \neq k \text{ and } d_{jkt} < \bar{d} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where d_{jkt} is the distance metric and \bar{d} is a distance threshold that determines the maximum distance at which one plant’s costs affect the other’s prices. This approach is attractive for the cement industry because the product is a commodity, aside from spatial considerations, so a distance metric can be constructed as the interaction of gasoline prices and the miles between plants. It is analogous to the assumption of Pinske, Slade, and Brett (2002) that the strategic complementarity of prices in wholesale gasoline markets decreases in the geographic distance between terminals. Further, the approach generalizes to markets with non-spatial differentiation provided that a reasonable Euclidean distance in attribute-space can be calculated (e.g., as in Langer and Miller (2013)).

Next, we let heterogeneity in own pass-through be determined by the degree of spatial differentiation, motivated by the theoretical result of ten Kate and Niels (2005) that own pass-through diminishes with the number of competitors in Cournot oligopoly models. In particular, we specify that

$$\rho_{jkt} = \alpha_0 + \alpha_1 \sum_{k \neq j, d_{jkt} < \bar{d}} 1/d_{jkt} \quad (3)$$

If α_1 is negative then the extent to which plants pass through plant-specific cost changes to customers diminishes with the number and proximity of competitors; the opposite effect arises if the parameter is positive. Together, restrictions (2) and (3) solve the dimensionality problem by reducing the number of pass-through parameters, while still allowing for the estimation of reasonable pass-through behavior.

The linear approximation in equation (1) makes aggregation to the regional level mathematically tractable. Suppose there exist $m = 1 \dots M$ geographic regions. The regions need

¹⁵Cross pass-through is intrinsically linked to the concept of strategic complementarity in prices, in the sense of Bulow, Geanakoplos, and Klemperer (1985), and in most standard demand systems the strength of strategic complementarity depends on the degree to which consumer view products as substitutes (e.g., Miller, Remer, and Sheu (2013)).

not comport to the local markets that commonly are used to model competition and instead should be conceptualized as sets of plants loosely defined based on geographic criteria, for data reporting purposes. Denote as \mathcal{J}_{mt} the set of plants that are in region m in period t . Then the average price that arises is $P_{mt} = \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} p_{jt}$, where ω_{jmt} is the fraction of the region's total production accounted for by plant j . We assume that production within regions is proportional to capacity, which yields proxies for the weights. This assumption, necessitated by the lack of plant-level production data, also is used by the EPA in its economic analysis of the industry (EPA (2010)).

Maintaining restrictions (2) and (3), a linear approximation to equilibrium prices at the regional-level then is given by

$$\begin{aligned}
P_{mt} &= \alpha_0 \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} c_{jt} + \alpha_1 \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} c_{jt} \sum_{k \neq j, d_{jkt} < \bar{d}} 1/d_{jkt} \\
&+ \beta \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \sum_{k \neq j, d_{jkt} < \bar{d}} c_{kt}/d_{jkt} \\
&+ \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} x'_{jt} \gamma + \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} (\mu_j + \tau_t) + \bar{\epsilon}_{mt}
\end{aligned} \tag{4}$$

where the region-year pricing residual is $\bar{\epsilon}_{mt} = \sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \epsilon_{jt}$. Many, but not all, of the plant fixed effects are separately identifiable in our application due to frequency with which the USGS redraws region boundaries.

Equation (4) provides the theoretical foundation for our reduced-form regression equation. The formulation, while novel in the empirical literature, remains quite general and could be applied to other markets with spatial or non-spatial differentiation. We highlight that the prices that arise in each region depend not only on the costs of plants in the region but also on the costs of plants outside the region, via the cross pass-through terms. Lastly, if autocorrelation or spatial correlations exist among the plant-level residuals, these manifest in the variance structure of the region-year residuals, and we use a number of estimation methodologies to account for that complication.

3.2 Measuring fuel costs

We calculate the fuel costs of each plant based on (i) the energy requirements of the plant's least efficient kiln, (ii) the primary fuel burned at the plant, and (iii) the price of the primary

fuel. Formally, the fuel costs per metric tonne of cement for plant j in year t equals

$$\text{Plant Fuel Cost}_{jt} = \text{Primary Fuel Price}_{jt} \times \text{Energy Requirements}_{jt} \div 1.05,$$

where the fuel price is in dollars per mBtu and the energy requirements are those of the least efficient kiln and are in mBtu per metric tonne of clinker. We scale down by five percent to reflect that a small amount of gypsum is ground together with clinker to form cement.¹⁶

We calculate energy requirements from the labor-energy input surveys of the PCA. There is no discernible change in the energy requirements of production, conditional on the kiln type, over 1990-2010. We calculate the average mBtu per metric tonne of clinker required in 1990, 2000, and 2010, separately for each kiln type, and apply these averages over 1990-2010. These requirements are 3.94, 4.11, 5.28, and 6.07 mBtu per metric tonne of clinker for dry precalciner kilns, dry preheater kiln, long dry kiln, and wet kilns, respectively. A recent survey of the USGS accords with our calculations (Van Oss (2005)). By contrast, technological improvements are evident over 1974-1990, conditional on kiln type. The labor-energy surveys indicate that in 1974 the energy requirements were 6.50 mBtu per metric tonne of clinker at dry kilns (a blended average across dry kiln types), and 7.93 mBtu per metric tonne of clinker at wet kilns. We assume that technological improvements are realized linearly over 1974-1990 and scale the energy requirements over that period appropriately.

We believe this to be the most reasonable methodology for calculating fuel costs, given the data available, but accept that it is impossible to measure perfectly the fuel costs at every kiln. Unobserved heterogeneity likely exists for a number of reasons, including variation in energy requirements within kiln types and variation in the fossil fuel prices paid by plants. We discuss two specific sources of possible measurement error, relating to supplementary waste fuels and the geographic heterogeneity of state-level fuel prices, in Appendix A.

Figure 2 explores the empirical distributions of regional prices and fuel costs over the sample period of 1974-2010. Panels A and B show the univariate distributions. The price distribution is nearly symmetric around the mean of \$102.39 per metric tonne. The fuel cost distribution is tighter and features two peaks around \$8 and \$18 per metric tonne. The relative tightness of the fuel cost distribution arises because fuel cost is one of many determinants of prices. The source of the asymmetry is heterogeneity in kiln technology, as illustrated in Panel C by the separate kernel density estimates for plants with wet and

¹⁶We focus on the least efficient kiln because it provides the most accurate measure of marginal fuel cost. The coal price data are in dollars per metric tonne, and we use the conversion factor of 23 mBtu per metric tonne, which we calculate to be the average energy content of bituminous coal obtained by cement plants, based on the labor-energy input surveys of the PCA.

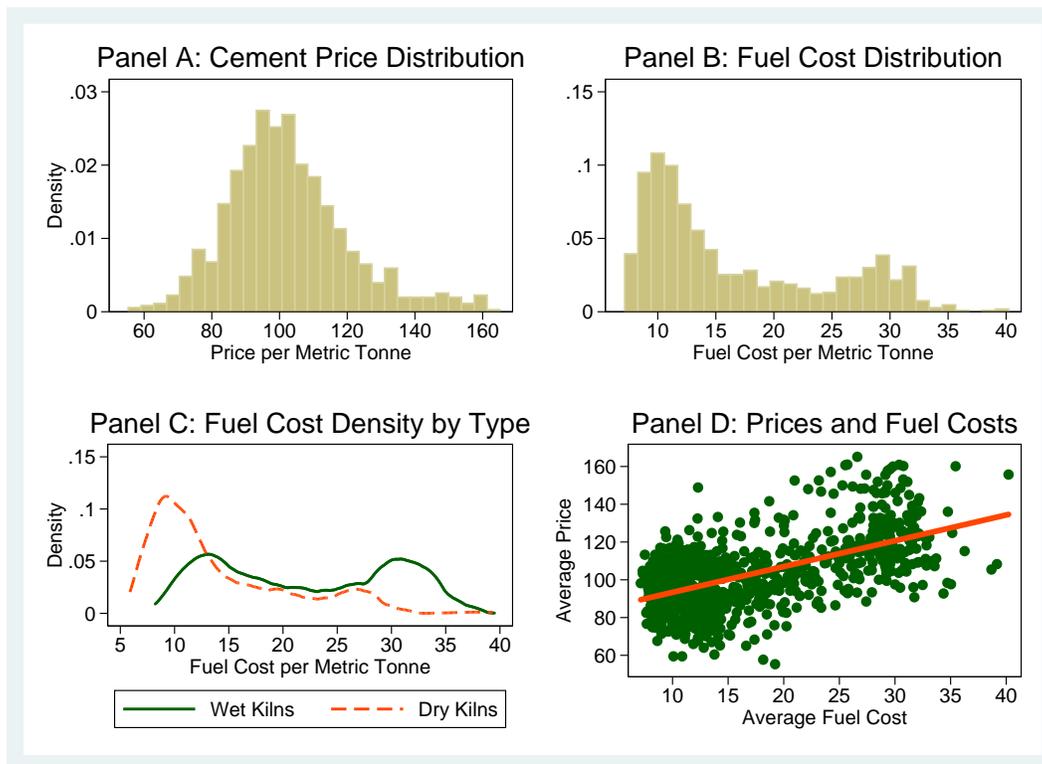


Figure 2: Regional Prices and Fuel Costs over 1974-2010

Notes: Panels A and B show the empirical distributions of cement price and fuel costs, and are based on 933 region-year observations. Panel C shows the kernel density of fuel costs, separately for plants with wet and dry kilns, and is based on 4,361 plant-year observations. Panel D shows a scatterplot of regional cement prices and fuel costs, as well as a line of best fit, and is based on 933 region-year observations. All prices and fuel costs are in real 2000 dollars per metric tonne of cement.

dry kilns. Panel D provides a scatterplot of the 933 region-year observations on prices and fuel costs. Observations with higher fuel costs also have higher prices, all else equal. A univariate regression of region-year prices on fuel costs yields a coefficient of 1.36 that, as with the national-level regression, is suggestive of high levels of pass-through.

3.3 Regressors

In constructing region-level regressors, we aggregate plant-level data such that the micro-foundations of the underlying model are preserved. Every regressor is constructed as a weighted average of a corresponding plant-level variable, consistent with the empirical model defined by equations (4). The weights are determined by region-specific capacity shares.

Table 3 defines these regressors explicitly, and provides summary statistics. We use two variables to control for demand. *Construction Employment* and *Building Permits* are

constructed by (i) calculating, for each plant, the total construction employment and building permits among all counties with centroids that are within the distance threshold, and (ii) aggregating to the region-level. Data on building permits and construction employment that we employ are highly predictive of portland cement consumption. A regression of state-level consumption data over 1974-2010, which is available from the *Minerals Yearbook* of the USGS, on building permits and construction employment aggregated to the state level, yields an R -squared value of 0.9354.¹⁷ Miller and Osborne (2014) use a similar data to calculate county-level “market sizes” in the estimation of a structural model of the industry.

We use two variables to control for competitive conditions. *Inverse Rival Distance* is constructed by calculating, for each plant, the count of competitors’ plants within some distance threshold. In this calculation, we divide competitors’ plants by their distance from the plant in question, so that closer competitors have greater influence. The variable increases in both the number and proximity of competitors. *Rival Capacity* is constructed by calculating, again for each plant, the total capacity at competitors’ plants within some distance threshold. We omit from Table 3 our controls for non-fuel costs, which are relatively straight-forward and seldom statistically significant in our regression analysis. The controls are based on plant-level indicator variables for the technology of the marginal kiln, i.e., whether the least efficient kiln at a plant is wet, long dry, dry with a preheater or dry with a precalciner. These plant-level variables are aggregated to the region-level, again using the capacity share weights. We also include as controls the count of competitor kiln types, within the distance threshold from each plant, aggregated to the region-level.

We use three main variables to capture pass-through, based on how restrictions (2) and (3) manifest in equation (4). *Fuel Costs* is constructed as the weighted average fuel cost of plants in the region, where plant fuel costs are calculated as described in the previous subsection and weights are by capacity share. $Fuel\ Costs \times Inverse\ Rival\ Distance$ is constructed based on the interactions of the fuel costs of each plant with the plant-level version of the *Inverse Rival Distance* variable, and allows for own pass-through to change based on the number and proximity of competitors. $Rival\ Fuel\ Costs \times Inverse\ Rival\ Distance$ captures cross pass-through and is constructed by calculating, for each plant, the sum of its competitors’ fuel costs normalized by distance. This captures cross pass-through. The total influence of cross pass-through, summing across competitors, varies with the number

¹⁷In the regression, there are 2,070 observations from the contiguous United States. The observations are at the state-year level, except for California, Pennsylvania and Texas each are split. The obtained regression coefficients are positive and statistically significant at better than the one percent level even when standard errors are clustered by state to account for auto-correlation.

and proximity of competitors. The latter two regressors are highly correlated, yet their coefficients are separately identifiable because plants often have different fuel costs than their nearby competitors. We extend the discussion of identification in Appendix B.

Our baseline specifications employ a distance metric defined by the interaction of the gasoline price index and the miles between plants, and a distance threshold of 400. This approach reflects the predominant role of trucking in distribution.¹⁸ Straight-line miles are highly correlated with both driving miles and driving time and, consistent with this, previously published empirical results on the industry are not sensitive to which of these measures is employed (e.g., Miller and Osborne (2014)). The baseline threshold follows prior findings that 80-90 percent of portland cement is trucked less than 200 miles (Census Bureau (1977); Miller and Osborne (2014)), so that plants separated by more than 400 miles are unlikely to compete for many customers. In robustness checks, we show that similar results are obtained with a distance metric defined by miles (i.e., not interacted with the gasoline price index), and with distance thresholds of 300 and 500.

We also have explored a more non-parametric approach to our specification of distance in the control variables. It is possible to construct the control variables defined above using many different distance thresholds, so that how impacts diminish with distance need not be linear in inverse distance. However, relaxing the specification along these lines complicates interpretation of the resulting coefficients, and it also has negligible effects on our pass-through results. Thus, in Section 5, we report results only from the simpler specification.¹⁹

¹⁸A fraction of cement is shipped to terminals by train (6% in 2010) or barge (11% in 2010), and only then is trucked to customers. Some plants may be closer than our metric indicates if, for example, both are located on the same river system.

¹⁹We use thresholds of 150, 200, 300, 400 and 500. When all of the resulting control variables are incorporated, the degree of autocorrelation in the pricing residuals explodes, creating instability in our FGLS and Bayesian regressions. To solve the problem, we use a Big Data shrinkage method, known as a “lasso,” that culls the number of explanatory variables. We implement the lasso in R using the `glmnet` package. Although the package produces coefficient estimates, we use it only as a variable selection method, i.e., we run our OLS, FGLS and Bayesian regression on those controls selected by the lasso. We refer interested readers to Hastie, Tibshirani, and Friedman (2009) for an overview of the procedure.

Table 3: Definitions and Summary Statistics for Selected Regressors

Regressor	Definition	Mean	St. Dev.	Description
<i>Control variables</i>				
Construction Employment	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \sum_{d_{ajt} < \bar{d}} EMP_{at}$	790.21	(533.39)	Total construction employment in nearby counties, aggregated to the region level.
Building Permits	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \sum_{d_{ajt} < \bar{d}} PER_{at}$	199.89	(143.90)	Total building permits in nearby counties, aggregated to the region level.
Inverse Rival Distance	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \sum_{k \neq j, d_{jkt} < \bar{d}} 1/d_{jkt}$	0.24	(0.26)	The count of competitors normalized by distance, aggregated to the region level.
Rival Capacity	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \sum_{d_{jkt} < \bar{d}} CAP_{kt}$	11.09	(6.34)	Total nearby competitor capacity, aggregated to the region level.
<i>Own pass-through variables</i>				
Fuel Costs	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} c_{jt}$	16.62	(7.83)	Fuel costs of the plant as defined in the text, aggregated to the region level.
Fuel Costs \times Inverse Rival Distance	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} c_{jt} \sum_{k \neq j, d_{jkt} < \bar{d}} 1/d_{jkt}$	3.93	(5.60)	Fuel costs times the count of competitors normalized by distance, aggregated to the region level.
<i>Cross pass-through variables</i>				
Rival Fuel Costs \times Inverse Rival Distance	$\sum_{j \in \mathcal{J}_{mt}} \omega_{jmt} \sum_{k \neq j, d_{jkt} < \bar{d}} c_{kt} / d_{jkt}$	3.97	(5.64)	Summation of competitors' fuel costs normalized by distance, aggregated to the region level.

Notes: Aggregation to the region level is conducted with the weights ω_{jmt} , which are approximated with capacity shares. In all equations, c_{jt} is the fuel cost of plant j in period t , d_{jkt} is the distance between plants j and k in period t , d_{ajt}^c is the distance between county a and plant j in period t , PER_{at} and EMP_{at} are building permits and construction employment in county a in period t , respectively, and CAP_{kt} is the capacity of plant k in period t . The summary statistics are calculated using a distance metric based on the gasoline index times the miles between plants. The distance threshold employed is 400.

4 Estimation

4.1 Conceptual discussion

We estimate the model defined by equation (4). The objective is to obtain the average relationships between costs and prices that arises in the sample period. We maintain, throughout, the assumption that the region-year pricing residuals are orthogonal to the regressors. We believe this assumption is appropriate. For example, while bias could arise if fossil fuel prices are correlated with unobserved components of cement demand, there are many reasons this is unlikely: (i) year fixed effects control directly for unobservable nation-wide changes in demand; (ii) the cement industry accounts for a small fraction of the fossil fuels consumed in the United States; (iii) the data indicate that, indeed, bituminous coal and petroleum coke prices do not follow the strongly pro-cyclical pattern of cement consumption. If anything, we expect unobserved costs to dominate the residuals, rather than unobserved demand, due to the predictive accuracy of our demand-side control variables. Unobserved costs should be uncorrelated with fuel costs because we include fixed effects for kiln technology.

Our focus on the estimation of average pass-through warrants discussion for two reasons. First, the literature emphasizes that pass-through is constant only for certain demand systems (e.g., Bulow and Pfleiderer (1983); Fabinger and Weyl (2014)). Absent constant pass-through, estimates of average pass-through can diverge from theoretical notions of pass-through, especially if the cost distribution is asymmetric, as it is in our data (MacKay, Miller, Remer, and Sheu (2014)). The data provide some support for constant pass-through. We construct an additional variable, based on the quadratic of plant fuel costs aggregated to the region level. The resulting coefficient is statistically insignificant, and the implied pass-through behavior of plants is similar.

Second, even if the underlying economic environment generates constant pass-through, it is possible that changes in demand and supply conditions create structural breaks within the sample period. While documenting average effects is of academic interest and advances substantially the existing empirical literature, more recent pass-through behavior is of greater interest for the policy applications that we develop. The coarseness of the data aggregation limits our ability to test for inter-temporal changes in pass-through (e.g, there are only 270 region-year observations over 2000-2011). That said, we develop some empirical evidence indicating that recent pass-through is similar to average pass-through: (i) when we interact fuel costs with an indicator variable for the 2000-2011 period, the variable has little explanatory power and its coefficient is statistically insignificant, and (ii) sub-sample regressions

using the 2000-2011 period produce coefficients that are similar to what is produced using the full sample.²⁰ We return to this discussion with the results.

We conclude by making explicit our assumption, necessary for the analysis, that the capacity and location of kilns are exogenous in their relationship to pricing decisions. We observe 352 kiln retirements in our data, and the median kiln age at retirement is 37 years. In contrast, prices adjust much more rapidly due to the prevalence of short-term supply contracts, and we therefore consider it unlikely that capacity and the geographic configuration of plants would be strongly correlated with the region-year pricing residual in equation (4), especially in the presence of plant and year fixed effects.

4.2 Estimation methodologies

We employ three distinct econometric methodologies in estimation: ordinary least squares (OLS), feasible generalized least squares (FGLS), and Bayesian regression. Each methodology has its own strengths: OLS is transparent and has desirable small sample properties; FGLS offers potential efficiency gains by accounting for autocorrelation within regions; and Bayesian regression allows for estimation that is fully consistent with the microfoundations of the model, including autocorrelation and spatial correlation at the plant-level, and also has desirable small sample properties.

We start with OLS, which we motivate with the linearity of the empirical model and the desirable small sample properties of the estimator. The coefficient estimates are unbiased and consistent under the assumption of orthogonality between the regressors and the region-year pricing residual, even in the presence of fixed effects (e.g., Lancaster (2000); Baltagi (2005)). Following Breusch (1978) and Godfrey (1978), we test for the presence of autocorrelation within regions, using our baseline specification, by regressing the residuals on lagged residuals. The procedure finds support for modest first-degree autocorrelation.²¹ Accordingly, we report standard errors that are clustered at the region level to account for within-region autocorrelation. We follow Wooldridge (2010) in implementing this correction, in order to ensure consistency in the presence of fixed effects.²² We caution that, by treating

²⁰The latter result requires that the control variables are omitted from the regressions because the empirical variation that exists in the 2000-2011 subsample is insufficient to separately identify pass-through from other demand and cost considerations.

²¹The coefficients from the baseline specification are shown in column (i) of Table 4. In the Breusch-Godfrey regression, the coefficient on the lagged residuals is 0.29, and the t -statistic is 9.15. Additional lags, if included, do not produce large or statistically significant coefficients.

²²Our correction, which we describe in Appendix C, differs slightly from the standard correction used for fixed effects models, which rely on a within estimator. Since we include plant-level fixed effects, and those

the residuals of different regions as independent, the correction accounts for neither spatial correlations that exist among regions in the same year, nor autocorrelation across regions that arises due to the (frequent) redrawing of regional boundaries.

Second, we estimate parameters with FGLS, under the assumptions that the residuals within each region are characterized by first-degree autocorrelation, but that residuals across regions are independent. This structure follows the results of the Breusch-Godfrey test described above, and has the practical effect of placing greater weight on regions with fewer observations, relative to OLS. It has the potential to improve the efficiency of the parameter estimates. We again follow Wooldridge (2010) in implementing FGLS in order to ensure consistency in the presence of fixed effects.²³ As with OLS, independence is maintained for observations in different regions, which simplifies estimation but is at odds with the underlying empirical model.

Lastly, we use Bayesian regression techniques to estimate the model in a way that is fully consistent with its microfoundations, including autocorrelation and spatial correlations in the plant-level residuals, and also has desirable small sample properties. In modeling the spatial correlations, we follow existing methodologies outlined in Gelfand (2012) and Bakar and Sahu (2011). We let the plant-year pricing residual be given by

$$\epsilon_{it} = \eta_{it} + \varepsilon_{it} \tag{5}$$

The term ε_{it} is a “nugget effect” that is i.i.d across plants and years, and has a mean zero normal distribution with variance σ_ε^2 . We let η_{it} generate first-degree autocorrelation and spatial correlations among the plant-level pricing residuals:

$$\eta_{it} = \kappa\eta_{i,t-1} + \nu_{it} \tag{6}$$

The term ν_{it} generates the spatial correlations. We assume that it is normally distributed with a variance matrix variance matrix of $\Sigma_\nu = \sigma_\nu^2 S_\nu$, such that the (i, j) element of S_ν equals

$$\psi(d_{ij}; \phi, \varphi) = \max \left\{ \frac{\varphi - d_{ij}}{\varphi}, 0 \right\}^\phi \tag{7}$$

fixed effects are weighted by plant capacities within a region, we multiply the dependent and independent variables by the matrix that projects them onto a space that is orthogonal to the capacity shares. This transformation removes the plant fixed effects and ensures that the estimated error variances are consistent.

²³As with the construction of the clustered standard errors our FGLS estimator is based on projection of the independent and dependent variables on a matrix orthogonal to plant shares. Our procedure is described in Appendix C.

This function is bounded weakly by zero and one. When the distance between plants (i.e., d_{ij}) exceeds φ , the spatial correlations are zero. The parameter ϕ determines how quickly the spatial correlation approaches zero as the distance between plant increase: the larger is ϕ , the faster correlation falls.²⁴ In total, the Bayesian regression requires the estimation of five addition parameters (κ , σ_ε^2 , σ_ν^2 , ϕ , φ). We assume standard priors wherever possible and, for most parameters, we choose the values of priors to be consistent with the default values used in the univariate regression R code provided with Rossi, Allenby, and McCulloch (2005). We develop the mathematical details of estimation in Appendix D. The Bayesian approach, while more complicated than the simple OLS and FGLS estimators we use, provides a conceptual advantage in that it correctly models the distribution of the error term. Classical approaches (such as maximum likelihood) could be used to estimate the model where the region-level errors are aggregates of the plant level errors, but they would be much more cumbersome to implement. In particular, they would require us to specify the distribution of the aggregate error term $\bar{\varepsilon}_{mt}$ conditional on past errors, $\bar{\varepsilon}_{m,t-1}, \dots, \bar{\varepsilon}_{m1}$, which could get quite complicated.²⁵ In the Bayesian approach we draw the plant-level η_{it} 's and condition on them when we take draws on the other parameters.

5 Regression Results

Table 4 provides results generated from OLS, FGLS, and Bayesian regression. With each, we use a distance metric based alternately on (i) miles times the gasoline price index and (ii) miles. The baseline specification is used throughout. Regressors include the main pass-through variables, controls variables for demand, cost, and competitive conditions, and fixed effects for plants and years. We report coefficients and standard errors for OLS and FGLS, and averages and standard deviations of the posterior distributions for Bayesian regression.

We focus particularly on the pass-through regressors. The *Fuel Cost* parameter is precisely estimated and near unity in all six regressions. Strictly interpreted, this corresponds to complete own and industry pass-through for plants with no competitors within the distance threshold. The parameter on the interaction *Fuel Costs* \times *Inverse Rival Distance* is negative and, while less precisely measured, still meets the usual standards for statistical significance. It follows that own pass-through typically is incomplete, as plants usually have competitors

²⁴We restrict φ and ϕ to be positive. We use a distant metric based on the miles between plants, regardless of the metric used to define the regression variables, because using a metric based on miles time the gasoline price index creates instability in the estimation procedure.

²⁵Note that we cannot apply the standard transformation with AR(1) errors, because $\bar{\varepsilon}_{mt} \neq \kappa \bar{\varepsilon}_{m,t-1} + \bar{\nu}_{it}$ due to the fact that capacity weights change over time.

Table 4: Regression Results with the Baseline Specification

	OLS		FGLS		Bayesian	
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
<i>Pass-through variables</i>						
Fuel Costs	0.99 (0.23)	1.01 (0.23)	1.02 (0.15)	1.16 (0.24)	1.1 (0.17)	1.31 (0.16)
Fuel Costs \times Inverse Rival Distance	-5.49 (1.71)	-4.14 (1.70)	-6.95 (0.67)	-5.09 (0.97)	-3.1 (0.95)	-3.75 (1.01)
Rival Fuel Costs \times Inverse Rival Distance	5.07 (2.07)	3.52 (2.18)	6.93 (0.77)	4.55 (1.15)	3.1 (1.03)	3.62 (1.09)
<i>Control variables</i>						
Construction Employment	0.014 (0.004)	0.039 (0.001)	0.018 (0.002)	0.039 (0.003)	-0.000 (0.004)	0.010 (0.005)
Building Permits	0.020 (0.016)	0.012 (0.015)	0.015 (0.008)	0.028 (0.006)	0.023 (0.011)	0.024 (0.010)
Inverse Rival Distance	9.75 (10.27)	-0.90 (23.47)	-3.48 (3.73)	-10.19 (7.26)	0.09 (5.71)	0.82 (6.24)
Rival Capacity	-0.54 (0.20)	0.04 (0.50)	-0.62 (0.09)	-0.15 (0.25)	-0.25 (0.21)	-1.03 (0.29)
Distance Metric	Miles \times Gas	Miles	Miles \times Gas	Miles	Miles \times Gas	Miles

Notes: Regression results obtained with 933 region-year observations over 1974-2010. The dependent variable is the cement price. All regressions include plant and year fixed effects, and aggregated indicators for kiln type and competitor kiln types. The distance threshold is 400, in miles times the gasoline price index for columns (i), (iii), and (v), and in miles for columns (ii), (iv) and (vi). OLS standard errors are calculated with a clustering correction for observations within the same region. The FGLS regressions results account for first-degree autocorrelation within regions. The Bayesian regressions account for first-degree autocorrelation and spatial correlations, both at the plant-level.

within the distance threshold, and that own pass-through diminishes in the number and proximity of competitors. The parameter on the interaction *Rival Fuel Costs* \times *Inverse Rival Distance*, which determines the cross pass-through terms, is positive and statistically significant. Because the magnitudes of the interaction terms are similar to each other within each regression, industry pass-through is largely unaffected by competition. Further, given the magnitude of the *Fuel Cost* parameter, industry pass-through should be roughly complete across a range of competitive conditions.²⁶ Lastly, we highlight that the magnitude of the interaction terms varies noticeably across the regressions. This has the implication, that we develop next, that own pass-through is more sensitive to specification choices and estimation techniques than is industry pass-through.

Table 5 provides derived pass-through statistics that we calculate by applying the regression coefficients to the 4,361 plant-year observations in the sample. The median industry pass-through varies from 0.92 to 1.29 across the six regressions, though the differences do not appear to be statistically significant. The median industry pass-through in the full sample is similar to that in 2010 alone because changing competitive conditions (there are fewer plants in 2010) matter little for industry pass-through. Median own pass-through ranges from 0.01 to 0.76 across the regressions, reflecting the varying magnitudes of the *Fuel Costs* \times *Inverse Rival Distance* parameter. Further, the confidence intervals are wide due to the (relative) imprecision with which that parameter is estimated. All of the regression models imply some number of negative own pass-through rates, which is inconsistent with standard economic theory. In our counterfactual exercises, we use the Bayesian regression results of column (v), which we believe provides the most reasonable pass-through results.²⁷

In Table 6, we explore alternative specifications to gain additional insight and explore the robustness of our results. We estimate with OLS for simplicity.²⁸ First, in column

²⁶The pass-through patterns we develop are reconciled easily with economic theory. Consider the case of Cournot competition among firms with constant (but possibly heterogeneous) marginal costs facing some market demand schedule. It can be shown (e.g., ten Kate and Niels (2005)) that

$$\rho_{jj} = \frac{1}{N + 1 - z} \quad \text{and} \quad \rho^M = \frac{N}{N + 1 - z} \quad (8)$$

where ρ_{jj} is own pass-through, ρ^M is industry pass-through, N is the number of firms, and z is positive with convex demand, negative with concave demand, and zero with linear demand. Specifically, $z = -\left(Q \frac{\partial^2 P}{\partial^2 Q}\right) / \left(\frac{\partial P}{\partial Q}\right)$, where Q and P are the market quantity and price, respectively. Own pass-through converges to zero as the number of firms grows large, while industry pass-through converges to unity from below or above, depending on the curvature of demand.

²⁷We normalize negative own pass-through rates to zero before conducting the counter-factual exercises to be consistent with economic theory.

²⁸We obtain similar results with FGLS, and have yet to run the corresponding Bayesian regressions.

Table 5: Derived Pass-Through Statistics

	OLS		FGLS		Bayesian	
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
<i>Industry Pass-Through</i>						
Median	0.94 (0.47,1.37)	0.92 (0.42,1.38)	1.02 (0.74,1.30)	1.08 (0.60,1.56)	1.13 (0.83,1.45)	1.29 (0.99,1.61)
Median, 2010	0.98 (0.51,1.39)	0.95 (0.46,1.39)	1.02 (0.75,1.31)	1.11 (0.63,1.58)	1.13 (0.84,1.46)	1.29 (0.99,1.61)
<i>Own Pass-Through</i>						
Median	0.19 (-0.46,0.79)	0.41 (-0.22,0.98)	0.01 (-0.29,0.30)	0.42 (-0.06,0.81)	0.67 (0.34,1.06)	0.76 (0.43,1.12)
Negatives	39%	28%	49%	30%	20%	21%
Median, 2010	0.69 (0.21,1.12)	0.58 (0.05,1.08)	0.64 (0.37,0.92)	0.63 (0.25,1.02)	0.96 (0.67,1.29)	0.91 (0.61,1.23)
Negatives, 2010	21%	23%	23%	23%	15%	20%

Notes: Pass-through in each column is calculated for 4,361 plant-year observations based on the corresponding regression results. In parentheses are 95% confidence intervals. For OLS and FGLS, these are obtained by drawing out of the asymptotic distribution of the regression coefficients. For the Bayesian regressions, the confidence intervals are the 2.5% and 97.5% order statistics obtained by applying the posterior distributions.

Table 6: Additional OLS Regression Results

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
Fuel Costs	0.48 (0.29)	0.98 (0.28)	1.06 (0.27)	0.83 (0.38)	1.01 (0.29)	1.55 (0.09)	1.40 (0.26)
Fuel Costs ²				0.004 (0.005)			
Fuel Costs × 1{Year ≥ 2000}					-0.06 (0.64)		
Fuel Costs × Inverse Rival Distance		-5.00 (1.59)	-5.24 (1.52)	-5.60 (1.53)	-5.49 (1.50)	-4.94 (1.00)	-3.92 (4.44)
Rival Fuel Costs × Inverse Rival Distance		4.71 (1.86)	4.83 (1.78)	5.11 (1.78)	5.07 (1.76)	4.34 (0.96)	2.66 (4.25)
Distance Threshold	400	300	500	400	400	400	400
Control Variables	yes	yes	yes	yes	yes	no	no
Fixed Effects	yes	yes	yes	yes	yes	no	no
Sample Period	Full	Full	Full	Full	Full	Full	2000-2010

Notes: OLS regression results. The sample is composed of 933 region-year observations over 1974-2010 in columns (i)-(vi) and 270 region-year observations over 2000-2010 in column (vii). The dependent variable is the cement price. The distance metric in all regressions is miles times the gasoline price index. The standard errors are calculated with a clustering correction for observations within the same region.

(i), we show that the *Fuel Costs* parameter drops to 0.48 if the fuel cost interaction terms are excluded from the regression, roughly midway between the median own pass-through and the industry pass-through that we calculate when the interaction terms are included as regressors (see the same column in Table 5). This is precisely what econometric theory would predict. The fuel cost of competitors, which is an omitted variable here, is positively correlated with the fuel cost regressor. Thus, the estimate should overstate own pass-through but, barring perfect correlation in plant costs, it should understate industry pass-through. In general, where the coefficient falls between own and industry pass-through should depend on the strength of the correlation between the costs of each plant and the costs of its competitors, and the degree to which prices are strategic complements. Our result highlights how accounting for oligopoly interactions in the empirical model can be necessary to obtain results that can be mapped to theoretical concepts.

The remainder of Table 6 shows the results of several robustness checks.²⁹ Columns (ii) and (iii) make clear that our results change little if we redefine regressors using a distance threshold of 300 or 500. Column (iv) adds a regressor constructed by squaring the fuel costs of each plant and then aggregating to the region level. The resulting parameter is not statistically significant, and does not materially affect inferences on pass-through. Thus, the data provide some support for constant pass-through, at least in the range of the data. This results holds when we also incorporate nonlinear interaction terms. Column (v) shows the results obtained when the *Fuel Costs* are interacted with an indicator variable for the years 2000-2010. There is no statistical support for a structural break in pass-through, and again this result holds if we also interact the other fuel cost regressors.³⁰ The empirical variation in the data is insufficient to support estimation based only on 2000-2010. However, as we show in columns (vi) and (vii), if the control variables and fixed effects are removed from the model, then the pass-through parameters obtained from the full sample are similar to those obtained from the subsample. This again provides little support for a structural break.

Lastly, we use our estimates to evaluate an implicit pass-through assumption that is made in recent articles on the portland cement industry (e.g., Ryan (2012); Fowlie, Reguant, and Ryan (2014)). The structural models used in those articles feature (i) Cournot competition among firms in local markets and (ii) constant elasticity market demand curves. Pass-through is determined by the number of firms and the elasticity of demand. In Table 7, we list the theoretical industry pass-through implied by the model, for selected local markets delineated by the EPA and used in Fowlie, Reguant, and Ryan (2014), over a range of elasticities considered in that article.³¹ We also show our empirical estimates of industry pass-through. The similarity between the theoretical predictions and the empirical estimates is apparent. In our mind, this reinforces the robustness of the structural models in how they project CO₂ regulation would affect dynamic investment decisions in the industry.

²⁹Due to space constraints, we are unable to enumerate the full list of robustness checks that we have conducted. However, our work thus far indicates that the results shown in this section are not driven by specific modeling choices, variable selections, outliers, or other peculiarities in the data.

³⁰We are interesting in testing for structural breaks because, in our counter-factual analyses, current pass-through behavior is what is most relevant.

³¹The authors can provide results for all 20 EPA markets upon request.

Table 7: Industry Pass-Through in Selected EPA Markets

	N	Theoretical Predictions			Empirical Estimate	95% Confidence Interval
		$\epsilon^D = 1.0$	$\epsilon^D = 1.5$	$\epsilon^D = 2.0$		
Atlanta	6	1.20	1.13	1.09	1.11	(0.80,1.44)
Birmingham	5	1.25	1.15	1.11	1.13	(0.84,1.46)
Chicago	4	1.33	1.20	1.14	1.13	(0.84,1.46)
Cincinnati	3	1.50	1.29	1.20	1.13	(0.84,1.46)
Detroit	2	2.00	1.50	1.33	1.14	(0.84,1.47)

Notes: Theoretical predictions are derived from a model of Cournot competition among firms with constant but heterogeneous marginal costs and a constant elasticity market demand schedule. We denote the number of firms with active plants in the EPA market in 2010 as N and the market elasticity as ϵ^D . The empirical estimate of industry pass-through is a capacity-weighted average of industry pass-through, as derived from the regression results, among plants in the EPA market.

6 Counterfactual Exercises

6.1 Market-based regulation of CO₂ emissions

We combine our pass-through results with estimates of margins and industry-wide demand elasticities, which we cull from the existing literature on portland cement, to evaluate how market-based regulation of CO₂ would affect producer and consumer surplus. Our calculations have direct bearing on the political economy of regulation and specifically on the question of how revenues obtained from regulation should be redistributed. Throughout, we model market-based regulation as a uniform carbon tax, which is economically equivalent to a cap-and-trade program in which permits are allocated with a uniform price auction. Our quantitative focus is on the short run and, in our calculations, we assume that the carbon tax is imposed only on domestic producers – this best utilizes our pass-through estimates, which relate domestic costs to domestic prices. We then discuss qualitatively how results would change in the long run and if importers are subject to the tax.

To motivate our calculations, first consider a model of symmetric oligopoly, and let total producer surplus be given by $PS = Q \times (P - C)$, where Q is total industry output, P is the industry price, and C is a constant marginal cost. Denote industry pass-through as ρ^M , the industry elasticity of demand as ϵ^D , and the price-cost margins of firms as m . Normalize the demand elasticity to be positive. The change in producer surplus due to an arbitrarily small output tax t is given by

$$\frac{\partial \pi}{\partial t} = [\rho^M (1 - m\epsilon^D) - 1] Q \quad (9)$$

This equation is derived in Atkin and Donaldson (2014) and appears as a “principle of incidence” in Weyl and Fabinger (2013). It is useful because it expresses the change in producer surplus in terms of industry pass-through, which we estimate, together with margins and the industry elasticity of demand, which have been estimated elsewhere in the literature. Assuming substitutes, it must be that $m\epsilon^D \in [0, 1]$ with zero representing price-taking behavior and one representing monopoly.³² We translate the output tax into a CO₂ tax using standard conversion methods.³³

We report short run results based on margins of 0.20, 0.30, 0.35, 0.40 and 0.50. The middle of this range is the most consistent with the available evidence. The demand estimates of Miller and Osborne (2014) imply to a margin of 31% when applied to single-plant firms. Further, a recent analysis conducted by the EPA constructed kiln-specific variables costs for each of 20 local markets; the costs imply an average margin of 43% when paired with the reported market prices (EPA (2009)). Similarly, we use a range of demand elasticities, reporting results for a domestic elasticity of demand of 0.60, 0.80, 1.00, 1.20, 1.40 and 1.60. Estimates in the literature range from roughly 0.87 to 2.03. Jans and Rosenbaum (1997) report an estimate of 0.87, Miller and Osborne (2010) report an estimate of 1.11, and Fowlie, Reguant, and Ryan (2014) report estimates ranging between 0.89 and 2.03. We find the lower portion of this range to be most plausible because cement comprises only a small fraction of overall construction costs, and for most projects, alternatives such as lumber, steel or asphalt are not economical.³⁴ Indeed, the existing evidence indicates that consumer

³²The product $m\epsilon^D$ is mathematically equivalent to the Rothschild Index (Rothschild (1942)), a measure of monopoly power based on the ratio of the industry elasticity to the firm-specific elasticity. The Rothschild index equals $1/N$ with Cournot competition, where N is the number of firms, so that calculating the change in producer surplus does not require knowledge of margins or the industry demand elasticity. We prefer to treat margins and elasticities independently because it allows for general inferences that are untethered from any specific model of competition. However, when we apply the Cournot framework and average over the 20 EPA local markets discussed in Section 5, we obtain results that are nearly identical to those obtained with our preferred methodology, assuming a margin of 0.35 and a domestic elasticity of 0.80. This conveys an additional robustness to our results.

³³We calculate that CO₂ emissions, in metric tonnes per metric tonne of cement, are 1.05, 0.98, 0.87, and 0.86 for wet, long dry, dry preheater and dry precalciner kilns, respectively. Our methodology is consistent with the Cement CO₂ Protocol, developed by leading cement firms for the Cement Sustainability Initiative of the World Business Council for Sustainable Development. We assume that 0.51 metric tonnes of CO₂ per metric tonne of clinker are produced from the chemical conversion of calcium carbonate, contained in limestone, into lime and carbon dioxide. We scale this up to 0.525 to account for CO₂ emitted during the calcination of cement kiln dust. We add to this the CO₂ emitted from the burning of coal, based on a emissions factor of 0.095 metric tonnes per mBtu and the kiln energy requirements reported in Section 3.2. Finally, we scale down by five percent to convert units of clinker to units of cement. Similar calculations underly the analysis in Fowlie, Reguant, and Ryan (2014). The capacity-weighted industry average in 2010 is 0.88 metric tonnes of CO₂ per metric tonne of cement.

³⁴Syverson (2004) documents that ready-mix concrete accounts for two percent of total construction costs

substitution away from domestic cement is captured predominately by cement importers (Miller and Osborne (2014)).

Table 8 shows the changes in short run producer surplus, per dollar of carbon tax, that we calculate over the ranges of margins and demand elasticities considered. Panel A uses an industry pass-through rate of 1.10, which is selected based on the median 2010 industry pass-through that we estimate from Bayesian regression. Panels B and C use an industry pass-through of 0.90 and 1.30, respectively, reflecting the statistical error that arises with the Bayesian regression. Producer surplus loss increases mechanically with margins and the elasticity of demand, and decreases mechanically with industry pass-through.³⁵ With margins of 0.35, an elasticity of 1.00, and industry pass-through of 1.10, the loss is \$17.12 million. While this is small relative to industry revenues (about \$8 billion in 2010), the loss becomes appreciable for larger carbon taxes. For example, the loss becomes \$685 million with a \$40 dollar carbon tax, assuming a constant pass-through rate. Official estimates of the social cost of carbon range from \$12 to \$129 per metric tonne for the year 2020, depending on the social discount rate (Working Group on Social Cost of Carbon (2013)).

The loss of consumer surplus typically is much larger. Following the methodology of Weyl and Fabinger (2013), we calculate consumer surplus to be \$66 million per dollar of carbon tax with industry pass-through of 1.10. If instead industry pass-through is 0.90 or 1.30, the loss of consumer surplus is \$54 million and \$78 million, respectively. Table 9 shows the resulting incidence, i.e., the ratio of consumer surplus loss to producer surplus loss. Incidence decreases mechanically with margins and the elasticity of demand, and increases mechanically with industry pass-through. With margins of 0.35, an elasticity of 1.00, and industry pass-through of 1.10, the burden on consumers is 3.86 times larger than the burden on producers. Equivalently, consumers bear 79% of the tax burden. While the magnitude of incidence varies over the plausible ranges of margins and elasticities, consumer burden always exceeds producer burden, and the discrepancy is substantial in nearly every instance.

The result that the burden of market-based regulation largely would fall on consumers is immediately relevant to policy discussions regarding the proper disbursement of taxation revenues. We highlight that our calculations likely *understate* the consumer burden for reasons alluded to above. First, reductions in demand due to the carbon tax predominately reflect substitution to imported cement (“leakage”). Market-based regulation, if it is to be politically palatable, would almost certainly be designed in a manner that mitigates leakage.

based on the 1987 Benchmark Input-Output Tables.

³⁵For some combinations of margins, elasticities and pass-through, producer surplus increases with the carbon tax (see Panel C). This is recognized as a theoretical possibility, but one that cannot be true globally.

Table 8: Change in Producer Surplus (\$MM) Per Dollar of Carbon Tax

Panel A: Industry Pass-through of 1.10						
	Domestic Elasticity of Demand					
Margins	0.60	0.80	1.00	1.20	1.40	1.60
0.20	-1.92	-4.56	-7.21	-9.85	-12.49	-15.14
0.30	-5.89	-9.85	-13.81	-17.78	-21.74	-25.71
0.35	-7.87	-12.49	-17.12	-21.74	-26.37	-30.99
0.40	-9.85	-15.14	-20.42	-25.71	-30.99	-36.28
0.50	-13.81	-20.42	-27.03	-33.64	-40.24	-46.85

Panel B: Industry Pass-through of 0.90						
	Domestic Elasticity of Demand					
Margins	0.60	0.80	1.00	1.20	1.40	1.60
0.20	-12.49	-14.66	-16.82	-18.98	-21.14	-23.31
0.30	-15.74	-18.98	-22.22	-25.47	-28.71	-31.95
0.35	-17.36	-21.14	-24.93	-28.71	-32.50	-36.28
0.40	-18.98	-23.31	-27.63	-31.95	-36.28	-40.60
0.50	-22.22	-27.63	-33.04	-38.44	-43.85	-49.25

Panel C: Industry Pass-through of 1.30						
	Domestic Elasticity of Demand					
Margins	0.60	0.80	1.00	1.20	1.40	1.60
0.20	8.65	5.53	2.40	-0.72	-3.84	-6.97
0.30	3.96	-0.72	-5.41	-10.09	-14.78	-19.46
0.35	1.62	-3.84	-9.31	-14.78	-20.24	-25.71
0.40	-0.72	-6.97	-13.21	-19.46	-25.71	-31.95
0.50	-5.41	-13.21	-21.02	-28.83	-36.64	-44.45

Notes: Calculations are based on a general model of symmetric oligopoly. Units are in millions of real 2010 dollars. We aggregate to the industry level based on the 2011 industry output of 67.90 million metric tonnes. We use the industry average ratio of 0.88 metric tonnes of CO₂ per metric tonne of cement to convert from an output tax to a carbon tax. Margins refer to $(P - C)/P$ where P is price and C is marginal cost. The domestic elasticity of demand is the percentage change in total domestic cement output with respect to a one percent increase in the domestic price. The ranges shown for margins and domestic elasticity reflect the existing literature on the portland cement industry.

Table 9: Incidence of a Carbon Tax

Panel A: Market Pass-through of 1.10						
	Domestic Elasticity of Demand					
Margins	0.60	0.80	1.00	1.20	1.40	1.60
0.20	34.38	14.47	9.17	6.71	5.29	4.37
0.30	11.22	6.71	4.78	3.72	3.04	2.57
0.35	8.40	5.29	3.86	3.04	2.51	2.13
0.40	6.71	4.37	3.24	2.57	2.13	1.82
0.50	4.78	3.24	2.44	1.96	1.64	1.41

Panel B: Market Pass-through of 0.90						
	Domestic Elasticity of Demand					
Margins	0.60	0.80	1.00	1.20	1.40	1.60
0.20	4.33	3.69	3.21	2.85	2.56	2.32
0.30	3.44	2.85	2.43	2.12	1.88	1.69
0.35	3.11	2.56	2.17	1.88	1.66	1.49
0.40	2.85	2.32	1.96	1.69	1.49	1.33
0.50	2.43	1.96	1.64	1.41	1.23	1.10

Panel C: Market Pass-through of 1.30						
	Domestic Elasticity of Demand					
Margins	0.60	0.80	1.00	1.20	1.40	1.60
0.20	-9.03	-14.13	-32.50	108.33	20.31	11.21
0.30	-19.70	108.33	14.44	7.74	5.28	4.01
0.35	-48.15	20.31	8.39	5.28	3.86	3.04
0.40	108.33	11.21	5.91	4.01	3.04	2.44
0.50	14.44	5.91	3.71	2.71	2.13	1.76

Notes: Calculations based on a general model of symmetric oligopoly. We use the industry average ratio of 0.88 metric tonnes of CO₂ per metric tonne of cement to convert from an output tax to a carbon tax. Margins refer to $(P-C)/P$ where P is price and C is marginal cost. The domestic elasticity of demand is the percentage change in total domestic cement output with respect to a one percent increase in the domestic price. The ranges shown for margins and domestic elasticity reflect the existing literature on the portland cement industry.

This would limit the demand losses of domestic firms, and it also would increase the relevant pass-through rate, as the costs of importers increase with the costs of domestic firms. Both effects shift burden from producers to consumers. Second, our quantitative analysis captures only short run effects. In the long run, profit loss likely could be reduced by the elimination of excess kiln capacity, which would increase market power and consumer prices. Predicting the magnitude of this long run effect requires a dynamic model along the lines of Fowlie, Reguant, and Ryan (2014). Indeed, as we develop in Section 5, our results corroborate the model employed there, along with its implicit assumptions on pass-through.

We now analyze the differential effects of market-based regulation, relaxing the assumption of symmetry used to derive equation (9). We focus on markup and price effects, rather than producer and consumer surplus. The plant-specific demand elasticities that would be required for surplus statements are not readily available in the literature.³⁶

Table 10 shows summary statistics regarding the change in markups that arise per dollar of carbon tax. Markups increase with the carbon tax on average because, in our baseline Bayesian regression, industry pass-through just exceeds unity. Plants that utilize less efficient kiln technology see smaller markup increases, though the differences are not large. Thus, unless inefficient plants face more elastic demand than other plants, our calculations provide little support for the notion that market-based regulation impacts substantially the distribution of producer surplus among technology classes. There also is some heterogeneity within technology classes. The inefficient plants that experience markup decreases are near efficient competitors, and the precalciner plants that experience the largest markup increases are near inefficient competitors.

Figure 3 maps the county-level price changes that arise per dollar of carbon tax. These are not informed directly from our regression results, so we approximate the geographic dispersion of effects by calculating the weighted average of the plant-level price changes, with weights that are proportional to the inverse miles between the plant and the county centroid. Counties with larger price increases are shown with deeper shades of blue. The distribution of price increases exhibits a modest degree of dispersion – nearly all counties experience increases between \$0.90 and \$1.30. The counties with larger price increases typically are near relatively inefficient cement plants. Smaller price increases arise in the southwest and southeast, where cement is produced in kilns that utilize modern technology.

³⁶In principle, one could obtain plant-specific elasticities by applying the structural estimates of Miller and Osborne (2014), which are obtained based on data from the U.S. Southwest over 1983-2003, to the entire country based on the geographic configuration in 2010.

Table 10: Change in Markup Per Dollar of Carbon Tax

Kiln Type	Change in Markup					
	Mean	5%	25%	50%	75%	95%
Wet	0.07	-0.19	0.10	0.11	0.13	0.14
Long Dry	0.08	-0.12	0.10	0.12	0.13	0.13
Dry with Preheater	0.11	0.07	0.11	0.12	0.13	0.13
Dry with Precalciner	0.13	0.10	0.11	0.12	0.12	0.22

Notes: Calculations are obtained from a general model of symmetric oligopoly that is calibrated to a industry pass-through of 1.326 and the 2011 industry output of 67.90 million metric tonnes. The industry pass-through is the in 2010 generated in the baseline regression specification. Markup refer to $(P - C)$ where P is price and C is marginal cost.

6.2 NESHAP Amendments

We turn now to an economic analysis of recent regulation promulgated by the EPA that reduces dramatically the legally permissible emissions of hazardous air pollutants (HAPs) including particulate matter, mercury, hydrocarbons, and hydrogen chloride. EPA analysis indicates that monetized health benefits, which it predicts exceed \$7-\$18 billion, far outweigh economic costs (EPA (2009); EPA (2010)). We revisit the price predictions of the EPA using our pass-through estimates.

The EPA relies on a Cournot model of competition to simulate the effect of regulation in each of 20 local markets based on conditions in 2005. The model incorporates a constant elasticity market demand curve and, for markets that are adjacent to a customs office, a constant elasticity import supply curve. It is calibrated to elasticity estimates in the existing literature. Details on the model and calibration are provided in Appendix F. We are able to fully replicate the EPA modeling results, up to the restriction that the EPA makes compliance costs public only at the market-average level. We then update the analysis to 2010, the most recent year of our sample, and compare the price predictions to an alternative based on our pass-through estimates. Because the economic costs of concern result from the pass-through of compliance costs to customers, empirical estimates of pass-through usually provide more accurate short run predictions than model-based simulations (Miller, Remer, Ryan, and Sheu (2013)). While the EPA approach is grounded in modeling techniques and functional forms that are standard in the literature of industrial organization (e.g., Fowlie, Reguant, and Ryan (2014)), the drawback is that pass-through is fully determined by functional forms and the first order properties of the system. Relying on empirical estimates of pass-through

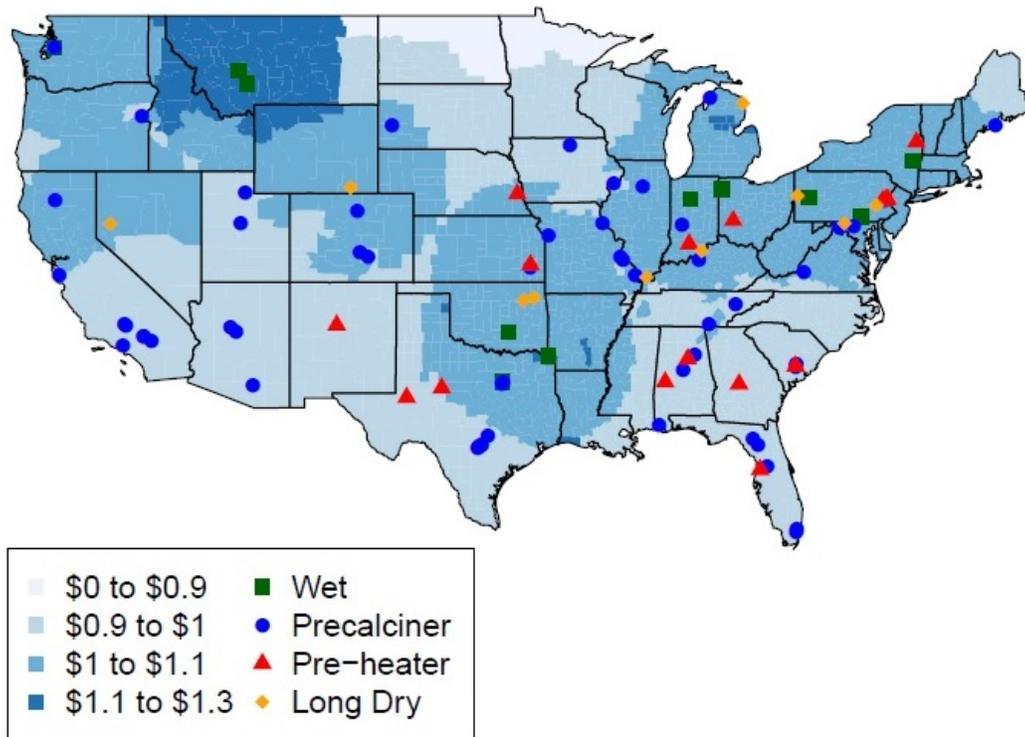


Figure 3: County-Level Price Changes per Dollar of Carbon Tax

Notes: County-level price changes are calculated as a weighted average of the plant-level price changes, with weights that are proportional to the inverse miles between the plant and the county centroid.

relaxes these assumptions and allows the data to inform predictions more directly.

Figure 4 provides a scatter-plot of market-specific predictions from the EPA’s Cournot model (on the vertical axis) and the predictions from our pass-through estimates (on the horizontal axis). Across the 20 local markets, the Cournot model yields average price increases of \$4.66 per metric tonne and the pass-through calculations yield average increases of \$4.49. The predictions are highly correlated, with a univariate correlation statistic of 0.88.³⁷ The similarity between the two methodologies arises because the industry pass-through that is implicit in the EPA model is close to the industry pass-through that estimate (e.g., see Table 7). To our knowledge, the quality of this match between the EPA model and empirical pass-through is coincidental. Interpreted in that light, our results allow us to confirm a previously untestable assumption on demand curvature that has first order implications for pass-through and the price effects of regulation.

³⁷The exceptions are Pittsburgh, for which the EPA under-predicts by \$3.30 relative to the pass-through calculation, and Cincinnati, for which the EPA over-predicts by \$2.70 relative to the pass-through calculation.

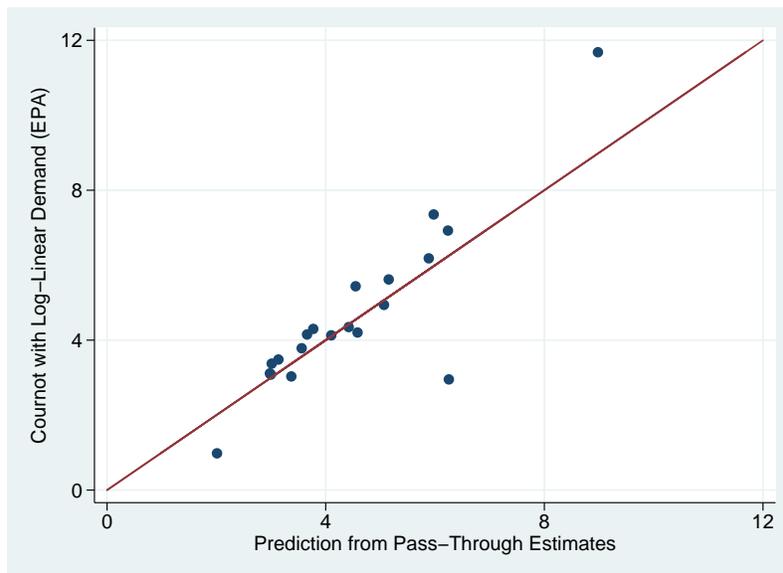


Figure 4: Price Effects of NESHAP Amendments for Portland Cement
 Notes: Each dot represents the price predictions based on (i) the EPA model of Cournot competition between firms facing a constant elasticity demand curve and (ii) our estimates of pass-through. Also shown is a 45-degree line.

6.3 Analysis of the Holcim/LaFarge merger

Here we predict the price effects of the proposed merger of Holcim and Lafarge, currently under review by the antitrust authorities. We believe this represents the first academic application of first order approximation (FOA) in the study of merger price effects. The methodology is grounded in oligopoly theory. The core logic is that horizontal mergers generate opportunity costs because, for each merging firm, a lower price requires it to forgo some profit that otherwise would be earned by its merging partner (Farrell and Shapiro (2010)). It follows that price changes can be approximated by multiplying these opportunity costs by a relevant notion of pass-through (Jaffe and Weyl (2013)), and Monte Carlo evidence supports the accuracy of this calculation in the merger context (Miller, Remer, Ryan, and Sheu (2013)).³⁸

In the final year of our data, the merging parties were the first and third largest cement

³⁸We provide mathematical details on FOA in Appendix E. In our application, the approach has notable advantages over most simulation methodologies. Modeling the industry based on Cournot competition in local markets would be inappropriate because mergers are not profitable, except to monopoly, unless additional complicating factors are invoked. The alternative of modeling the industry based on Bertrand competition with spatial differentiation, as in Miller and Osborne (2014), is computationally difficult and requires an assumption on the functional form of demand that partially determines the magnitude of the merger effects. By contrast, our present calculations are both simple and consistent with profitable mergers, spatial differentiation, and arbitrary demand functions.

firms in the United States, by clinker capacity. In the wake of the Great Recession, however, both firms divested and closed unprofitable plants. Thus, we provide predictions based on (i) an industry snapshot in 2010, where we have complete data; and (ii) a 2014 snapshot in which the status of Holcim and Lafarge plants is updated from press releases, but the status of other plants is left as in 2010. These two sets of results indicate interesting contrasts in the spatial distribution of price effects due to the merger.

The requisite inputs for FOA are pass-through and what is known as “upward pricing pressure” or “UPP” among antitrust economists. UPP equals the opportunity cost created by the merger and can be calculated from diversion – the fraction of sales lost by each merging firm, due to a price increase, that shift to the other merging firm – and the margins of the merging firms. Diversion and margins often are available to antitrust agencies through confidential business data, but we must rely on publicly available data and informed assumptions. Namely, we let diversion be proportional to the inverse distances between plants, we set margins to 30 percent, and we obtain measures of pre-merger prices based on the USGS regions in which plants are located, following Fowlie, Reguant, and Ryan (2014).³⁹

Table 11 shows results for each the 16 Holcim and Lafarge plants with a nonzero predicted price change in either the pre-divestiture and post-divestiture samples. Our calculations indicate that, but for the post-recession divestitures and closures, the merger would have resulted in substantial price elevations, on the order of 5%-7% at many Holcim plants and 3%-7% at many Lafarge plants. Accounting for changes in plant status, the predicted effects are more modest, at 3%-4%, and these exist only for five plants.

In Figure 5, we map the county-level distribution of price effects, both pre-divestiture (map A) and post-divestiture (map B). We calculate these county-level price changes based on a weighted-average of the plant-level price changes, with weights that are proportional to the inverse distance between the plants and the county centroids. While this a crude correspondence, we nonetheless consider it a useful way to examine the geographic dispersion of effects. The pre-divestiture map shows substantial price elevations in the Northeast, Southeast, and Great Plains. These effects arise due to Holcim and Lafarge plants, shown in orange circles and red diamonds, that are in close proximity to each other. The post-divestiture map, however, shows that price elevations are confined to the Northeast and the Great Plains, and that these elevations are smaller in magnitude. Although we do not

³⁹To illustrate the diversion assumption, consider a simple model with three firms. The distances between firm A and its competitors, firms B and C, is 50 and 150 miles, respectively. Our working assumption is that 75% of firm A’s customers view firm B as their next best option and 25% view firm C as their next best option. The same diversion rates emerge if distance instead is measured in miles times the gasoline price because the gasoline price affects both distances proportionally.

Table 11: Predicted Price Effects of a Holcim/Lafarge Merger

City	State	Capacity	Pre-Merger Price	Pre-Divestiture Price	Pre-Divestiture Price Effects	Post-Divestiture Price	Post-Divestiture Price Effects
<i>Holcim Plants</i>							
Bloomsdale	MO	3,720	82.50	5.42	6.6%	3.88	4.70%
Midlothian	TX	2,126	91.00	0.06	0.0%	.	.
Holly Hill	SC	1,860	86.78	5.48	6.3%	.	.
Theodore	AL	1,503	83.00	6.84	8.2%	.	.
Catskill	NY	572	83.00	7.42	8.1%	.	.
Ada	OK	524	97.38	7.02	7.2%	.	.
Hagerstown	MD	512	84.16	3.83	4.5%	3.57	4.2%
Mason City	IA	363	103.13	0.08	0.0%	.	.
<i>Lafarge Plants</i>							
Ravena	NY	1,680	91.72	6.82	7.4%	2.30	2.5%
Calera	AL	1,403	83.00	3.09	3.7%	.	.
Grand Chain	IL	1,014	89.07	2.72	3.1%	2.67	3.0%
Harleyville	SC	978	86.78	2.38	2.7%	.	.
Buffalo	IA	975	103.13	6.84	6.6%	3.08	3.0%
Sugar Creek	MO	943	82.50	3.30	4.0%	.	.
Whitehall	PA	702	95.00	1.27	1.3%	0.61	0.1%
Tulsa	OK	580	97.38	4.75	4.9%	.	.

Notes: Predicted price effects are obtained from first order approximation. Prices are in real 2010 dollars. Capacity is in thousands of metric tonnes of clinker per year. Plants for which no price change is predicted do not appear, including Holcim plants in Florence CO, Morgan UT and Three Forks MT, and Lafarge plants in Alpena MI, Paulding OH and Seattle WA. The divestitures are based on Holcim's and Lafarge's consummated plant sales over 2011-2013.

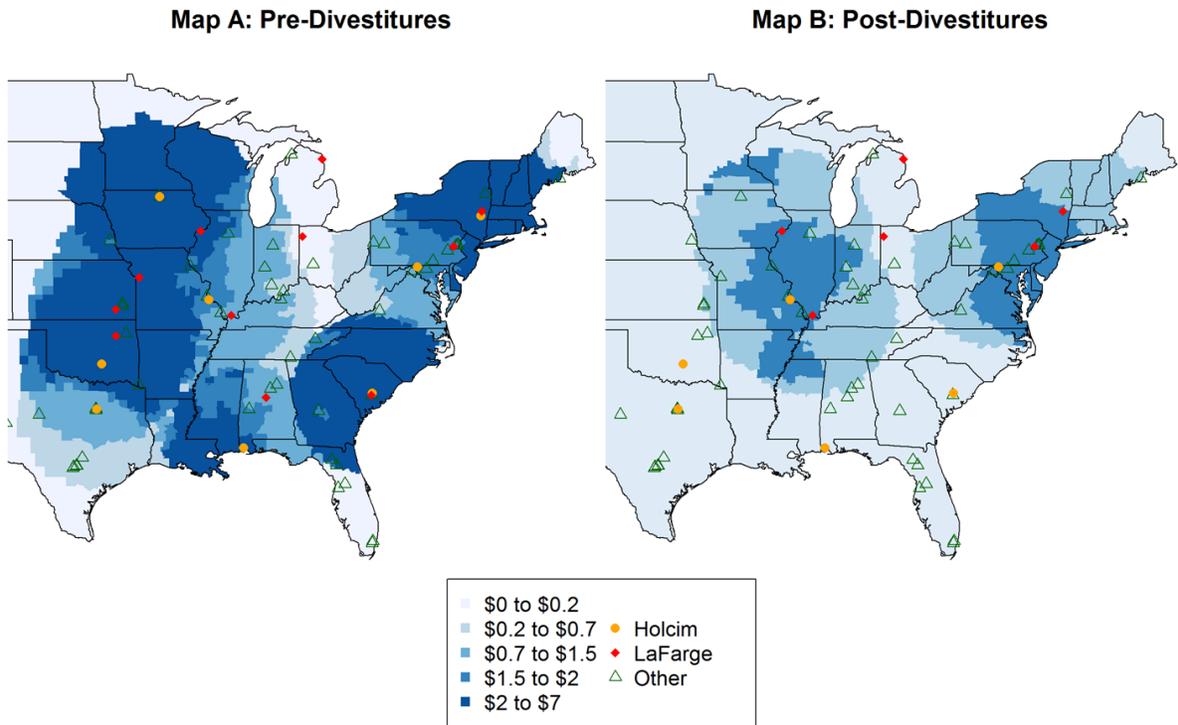


Figure 5: County-Level Predicted Price Effects of a Holcim/Lafarge

Notes: County-level price changes are calculated as a weighted-average of the plant-level price changes obtained from first order approximation. The weights are proportional to inverse distance. The distance metric employed is the log of one plus the miles between the plant and the county centroid. The distance threshold is 400. Price changes are in real 2010 dollars.

investigate the matter formally, we suspect that each remaining pocket of harm could be remedied with a single divestiture.⁴⁰

7 Conclusion

Our objectives in conducting the research described herein are twofold: First, we have intended to demonstrate that the estimation of pass-through is feasible, even without access to large quantities of price data at the firm/product level, and in a manner consistent with the oligopoly interactions of concentrated markets. Second, we have intended to reinforce theoretical findings about how pass-through can be used to better understand markets. In our view, the value of empirical research on pass-through is great, and we hope that our

⁴⁰Our calculations also do not inform whether these price effects could be mitigated by cost efficiencies or other factors, and we leave such considerations to the antitrust authorities.

own work helps spur endeavors elsewhere. With that in mind, we offer some caveats that are relevant to our own research, and that are likely to generalize to other settings.

Estimates of pass-through typically are obtained with reduced-form regressions of price on a cost shifter. These regressions are vulnerable to bias from measurement error and omitted variables. Care must be taken in constructing the regressors, and in considering factors that correlate with both price and the cost shifter. In our application, the electricity price stands out as one such factor, but bringing it into the regression does not affect inference. Regression coefficients, even if unbiased, provide information on the average short run pass-through that arises in the data. These can diverge from long run pass-through, which typically is more interesting from a policy standpoint, if menu costs exist or if firms use simple rule-of-thumb pricing rules. Further, because equilibrium pass-through depends on higher order properties of the cost and demand functions, whether average pass-through corresponds to a theoretical notion of pass-through, at any equilibrium point, is unclear. It is possible to investigate this latter point, as we do in our application, but tests can be limited by the amount of empirical variation present.

Caution also must be taken when applying pass-through to analyze counter-factual scenarios. It is an open question whether historical pass-through are helpful in evaluation events that increase marginal costs well above historical levels. Yet this concern should not be overly limiting. In counter-factual scenarios, some assumptions must be made, and the existing Monte Carlo evidence indicates that using pass-through to inform predictions typically improves accuracy (Miller, Remer, Ryan, and Sheu 2013). Finally, the theoretical ambiguities that exist with respect to pass-through make external validity challenging, absent careful consideration of institutional details. With these caveats in mind, we reiterate our belief in the value of empirical pass-through research.

References

- Aguirre, I., S. Cowan, and J. Vickers (2010). Monopoly price discrimination and demand curvature. *American Economic Review* 100(4), 1601–1515.
- Amiti, M., O. Itskhoki, and J. Konings (2012). Importers, exporters, and exchange rate disconnect. NBER Working Papers 18615, National Bureau of Economic Research, Inc.
- Anderson, S. P., A. de Palma, and B. Keider (2001). Tax incidence in differentiated product oligopoly. *Journal of Public Economics* 81(2), 173–192.
- Ashenfelter, O., D. Ashmore, J. B. Baker, and S.-M. McKernan (1998). Identifying the firm-specific cost pass-through rate. *FTC Working Paper*.
- Atkeson, A. and A. Burstein (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review* 98(5), 1998–2031.
- Atkin, D. and D. Donaldson (2014). Who’s getting globalized? The size and nature of intranational trade costs.
- Auer, R. A. and R. S. Schoenle (2013). Market structure and exchange rate pass-through.
- Bakar, K. and S. Sahu (2011). sptimer: Spatio-temporal bayesian modeling using R. Working Paper.
- Baltagi, B. H. (2005). *Econometric Analysis of Panel Data* (3rd ed.). John Wiley & Sons, Ltd.
- Barzel, Y. (1976). An alternative approach to the analysis of taxation. *Journal of Political Economy* 84(6), 1177–1197.
- Bergstrom, T. C. and H. R. Varian (1985). Two remarks on Cournot equilibria. *Economics Letters* 19, 5–8.
- Berman, N., P. Martin, and T. Mayer (2012). How do different exporters react to exchange rate changes? *Quarterly Journal of Economics* 127(1), 437–492.
- Besanko, D., D. Dranove, and M. Shanley (2001). Exploiting a cost advantage and coping with a cost disadvantage. *Management Science* 47(2), 221–235.
- Besanko, D., J.-P. Dube, and S. Gupta (2005, Winter). Own-brand and cross-brand retail pass-through. *Marketing Science* 1(1), 123–137.
- Besley, T. J. and H. Rosen (1998). Sales taxes and prices: An empirical analysis. *National Tax Journal* 52(2), 157–178.

- Bonnet, C., P. Dubois, and S. B. Villas-Boas (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.
- Borenstein, S., C. Cameron, and R. Gilbert (1997). Do gasoline prices respond asymmetrically to crude oil price changes? *Quarterly Journal of Economics* 112(1), 305–339.
- Breusch, T. S. (1978). Testing for autocorrelation in dynamic linear models. *Australian Economic Papers* 17(31), 334–355.
- Broda, C., N. Limao, and D. E. Weinstein (2008). Optimal tariffs and market power: The evidence. *American Economic Review* 98(5), 2032–2065.
- Buchanan, J. M. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review* 59(1), 174–177.
- Bulow, J. I., J. D. Geanakoplos, and P. D. Klemperer (1985). Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy* 93(3), pp. 488–511.
- 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.
- Campa, J. M. and L. S. Goldberg (2005). Exchange rate pass-through into import prices. *Review of Economics and Statistics* 87(4), 679–690.
- 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.
- Dornbusch, R. (1987). Exchange rates and prices. *American Economic Review* 77, 93–106.
- EPA (1998). Regulatory impact analysis of cement kiln dust rulemaking.
- EPA (2009). Regulatory Impact Analysis: National Emission Standards for Hazardous Air Pollutants from the Portland Cement Manufacturing Industry.
- EPA (2010). Regulatory Impact Analysis: Amendments to the National Emission Standards for Hazardous Air Pollutants (NESHAP) and New Source Performance Standards (NSPS) for the Portland Cement Manufacturing Industry.
- Fabinger, M. and E. G. Weyl (2014). A tractable approach to pass-through patterns with applications to international trade.
- Fabra, N. and M. Reguant (2014). Pass-through of emissions costs in electricity markets. *American Economic Review* 104(9), 2872–2899.

- Farrell, J. and C. Shapiro (2010). Antitrust evaluation of horizontal mergers: An economic alternative to market definition. *B.E. Journal of Theoretical Economics: Policies and Perspectives* 10(1).
- Fowlie, M., M. Reguant, and S. P. Ryan (2014). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*.
- Gelfand, A. (2012). Hierarchical modeling for spatial data problems. *Spatial Statistics* 1, 30–39.
- Genesove, D. and W. P. Mullin (1998). Testing static oligopoly models: Conduct and cost in the sugar industry, 1980-1914. *RAND Journal of Economics* 29(2), 355–377.
- Godfrey, L. G. (1978). Testing for higher order serial correlation in regression equations when the regressors include lagged dependent variables. *Econometrica* 46(6), 1303–1310.
- Golberg, P. K. and R. Hellerstein (2013). A structural approach to identifying the source of local-currency price stability. *Review of Economic Studies* 80(1), 175–210.
- Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh, and N. Li (2011). International prices, costs, and markup differences. *American Economic Review* 101, 1–40.
- Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning*. Springer.
- Hellerstein, R. (2008). Who bears the cost of a change in the exchange rate: Pass-through accounting for the case of beer. *Journal of International Economics* 76, 14–32.
- Hong, G. H. and N. Li (2013). Market structure and cost pass-through in retail.
- Jaffe, S. and E. G. Weyl (2013). The first order approach to merger analysis. *American Economic Journal: Microeconomics* 5(4), 188–218.
- Jans, I. and D. Rosenbaum (1997). Multimarket contact and pricing: Evidence from the U.S. cement industry. *International Journal of Industrial Organization* 15, 391–412.
- Krugman, P. (1986). Pricing to market when the exchange rate changes. NBER Working Papers 1926, National Bureau of Economic Research, Inc.
- Lancaster, T. (2000). The incidental parameter problem since 1948. *Journal of Econometrics* 95, 391–413.
- Langer, A. and N. H. Miller (2013). Automakers’ short-run responses to changing gasoline prices. *Review of Economics and Statistics* 95(4), 1198–1211.

- MacKay, A., N. H. Miller, M. Remer, and G. Sheu (2014). Bias in reduced-form estimates of pass-through. *Economics Letters* 123(2), 200–202.
- Marion, J. and E. Muehlegger (2011). Fuel tax incidence and supply conditions. *Journal of Public Economics* 95, 1202–1212.
- Marshall, A. (1890). *Principles of Economics*. New York: Macmillan.
- McBride, M. (1983, December). Spatial competition and vertical integration: Cement - concrete revisited. *American Economic Review* 73, 1011–1022.
- Melitz, M. J. and G. I. P. Ottaviano (2008). Market size, trade and productivity. *Review of Economic Studies* 75, 295–316.
- Milgrom, P. and J. Roberts (1996). The LeChatelier principle. *American Economic Review* 86(1), 173–179.
- Miller, N. H. and M. Osborne (2010). Competition among spatially differentiated firms: An estimator with an application to cement. *EAG Discussion Series Paper 10-2*.
- Miller, N. H. and M. Osborne (2014). Spatial differentiation and price discrimination in the cement industry: Evidence from a structural model. *RAND Journal of Economics*.
- Miller, N. H., M. Remer, C. Ryan, and G. Sheu (2013). On the first order approximation of counterfactual price effects in oligopoly models.
- Miller, N. H., M. Remer, and G. Sheu (2013). Using cost pass-through to calibrate demand. *Economic Letters* 118, 451–454.
- Nakamura, E. and D. Zerom (2010). Accounting for incomplete pass-through. *Review of Economic Studies* 77(3), 1192–1230.
- Peltzman, S. (2000). Prices rise faster than they fall. *Journal of Political Economy* 108, 466–502.
- Pinske, J., M. E. Slade, and C. Brett (2002). Spatial price competition: A semiparametric approach. *Econometrica* 70(3), 1111–1153.
- Poterba, J. M. (1996). Retail price reactions to changes in state and local sales taxes. *National Tax Journal* 49(2), 165–176.
- Rossi, P., G. Allenby, and R. McCulloch (2005). *Bayesian Statistics and Marketing*. Wiley.
- Rothschild, K. W. (1942). The degree of monopoly. *Economica* 9, 24–39.
- Ryan, S. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica* 80(3), 1019–1062.

- Samuelson, P. A. (1947). *The foundations of economic analysis*. Cambridge, MA: Harvard University Press.
- Scharfstein, D. and A. Sunderam (2013). Concentration in mortgage lending, refinancing activity, and mortgage rates.
- Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy* 112, 1181–1222.
- Syverson, C. and A. Hortaçsu (2007). Cementing relationships: Vertical integration, foreclosure, productivity, and prices. *Journal of Political Economy* 115, 250–301.
- ten Kate, A. and G. Niels (2005). To what extent are cost savings passed on to consumers? an oligopoly approach. *European Journal of Law and Economics* 20, 323–337.
- UBS (1997). Global cement outlook. UBS Global Research.
- Van Oss, H. G. (2005). Background facts and issues concerning cement and cement data. Open-File Report 2005-1152, U.S. Department of the Interior, U.S. Geological Survey.
- Van Oss, H. G. and A. Padovani (2003). Cement manufacture and the environment, Part II: Environmental challenges and opportunities. *Journal of Industrial Ecology* 7(1), 93–126.
- Villas-Boas, S. B. (2007). Vertical relationships between manufacturers and retailers: Inference with limited data. *Review of Economic Studies* 74(2), 625–652.
- Weyl, E. G. and M. Fabinger (2013). Pass-through as an economic tool. *Journal of Political Economy* 121(2), 528–583.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). MIT Press.
- Working Group on Social Cost of Carbon (2013). Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under executive order 12866. Technical Report, United States Government.

Appendix

A Details on the Data Collection

We discuss details of the data collection process here in order to assist replication. We start with the *Plant Information Survey* (PIS) of the PCA. Our sample includes annual observations over 1974-2010. The PIS is published annually over 1974-2003 and also semi-annually in 2004, 2006, 2008 and 2010. We make use of all of the publications with the exception of 1978 and 1981. We impute values for the capacity, technology, and primary fuel of each kiln in the missing years based on the preceding and following data. In most instances, imputation is trivial because capacity, technology and fuel are persistent across years. When the data from the preceding and following years differ, we use the data from the preceding year. We are able to identify kilns that are built in the missing years because the PIS provides for each kiln the year of construction. We remove from the analysis 198 kiln-year observations for which the kiln is identified in the PIS as being idled. These occur mostly in the late 1980s and over 2009-2010. There are 55 plant-year observations – out of 4,416 – for which all kilns at a plant are observed to be idled. A handful of kilns drop out of the PIS and then reappear in later years. We treat those observations on a case-by-case basis, leveraging detailed qualitative and quantitative information provided in the *Minerals Yearbook* of the USGS. We detail the available evidence and the selected treatment in our annotated Stata code. Lastly, we remove from the analysis a small number of kilns that produce white cement, which takes the color of dyes is used for decorative purposes. Production requires higher kiln temperatures and iron-free raw materials, and the resulting cost differential makes it a poor substitute for gray cement in most instances.

We obtain data on delivered bituminous coal prices for the industrial sector from the annual *Coal Reports* of the Energy Information Agency (EIA). Averages are available at the national, regional and state levels over 1985-2012. We convert prices from dollars per short ton to dollars per metric tonne using the standard conversion factor. Roughly 40% of the state values are withheld and must be imputed. We first use linear interpolation to fill in missing strings no longer than three years in length. We then calculate the average percentage difference between the observed data of each state and the corresponding national data, and use that together with the national data to impute missing values. For 14 states, all or nearly all of the state-level data are withheld, and we instead set the state price equal to the regional price.⁴¹ We backcast the coal price data to the period 1974-1984 using data on

⁴¹These states are Connecticut, Delaware, Louisiana, Massachusetts, Maine, Mississippi, Montana, North

the national average free-on-board (FOB) price of bituminous coal over 1974-2008 published in the 2008 *Annual Energy Review* of the EIA. Backcasting is based on (1) the state-specific average percentage differences between the delivered state and national prices; and (2) the percentage differences between the delivered national prices and the FOB national prices over the 1985-1990. The coal price data are reported in dollars per metric tonne. We convert to dollars per mBtu using the conversion factor of 23 mBtu per metric tonne of bituminous coal, which we calculate based on the labor-energy input surveys of the PCA.

The USGS *Minerals Yearbook* publishes average prices per region. In total, there are 61 regions, fully contained in the contiguous United States, that appear at least once.⁴² In Table A.1, we list the number times we observe each region over the sample period 1974-2010. Only five regions are observed in every year – Alabama, Illinois, Maine/New York, Missouri, and Ohio. Regions more commonly are observed for a portion of the sample. The median number of observations for a region is nine. The regions exhibit numerous features that make it difficult to interpret them as local markets. We highlight two here. First, regions are not always contiguous. An example is Georgia, which in 14 years is grouped with Virginia and West Virginia but not with South Carolina. Second, the regions exhibit little constancy over the sample period. An example is Nevada, which in 19 years is grouped with Idaho, Montana and Utah and in nine years is grouped with Arizona and New Mexico. Nonetheless, the data provide useful information on prices throughout the United States and serve to motivate our empirical framework, which we develop to accommodate such data.

We obtain state-level data on the prices of petroleum coke, natural gas, and distillate fuel oil, again for the industrial sector, from the State Energy Database System (SEDS) of the EIA. The imputation of missing values is required only for petroleum coke. To perform the imputation, we first calculate average percentage difference between the observed data of each state and the corresponding national data, and use that together with the national data to impute missing values. In five states with active kilns, all or nearly all of the state-level data are withheld so we base imputation instead on the average petroleum coke prices that arise in adjacent states and nationwide.⁴³ The SEDS data are in dollars per mBtu.

Dakota, New England, New Jersey, New Mexico, Nevada, Oregon and Vermont.

⁴²We do not include regions that incorporate states and territories outside the contiguous United States. For example, we exclude Oregon/Washington/Alaska/Hawaii, which exists over 1983-1985.

⁴³We use the national price here because the prices in many adjacent states similarly are withheld. We impute the price of Maine using the national price because data for adjacent states are withheld (there are no kilns in adjacent states). We impute the price of Iowa using the arithmetic mean of the Illinois price and the national price. We impute the price of Nevada and Arizona using the arithmetic mean of the California price and the national price. We impute the price of Kansas using the arithmetic mean of the Oklahoma price, the Missouri price, and the national price.

Table A.1: Number of Observations and Active Plants by USGS Region

Region	Observations	Plants	Region	Observations	Plants
AL	37	5.41	ID/MT/WY	9	4.00
IL	37	3.84	MD/WV	9	4.00
ME/NY	37	5.19	OK	9	3.00
MO	37	5.41	GA	8	3.00
OH	37	3.54	LA/MS	8	4.00
FL	36	4.74	OR/NV	8	3.00
East PA	36	7.47	TN	9	5.33
West PA	36	3.72	TX	8	18.00
North CA	35	4.26	AR/MS/LA	7	3.57
South CA	35	7.34	KY/VA/WV	6	3.00
KS	34	4.29	SD	10	1.00
IN	30	4.03	ID/MT	5	3.00
SC	30	3.00	ID/MT/UT	5	5.60
North TX	29	5.62	KY/NC/VA	5	3.00
South TX	29	5.52	MD/VA/WV	6	5.00
AR/OK	28	4.32	NE/WI	5	2.00
CO/WY	27	3.89	IN/KY/WI	4	6.25
MI	26	3.69	VA/NC/SC	4	5.00
MD	20	3.00	AR/MS	2	3.00
AZ/NM	19	3.05	GA/SC	3	5.00
IA/NE/SD	19	5.00	IN/KY	3	3.00
ID/MT/NV/UT	19	6.47	KS/NE	3	7.00
KY/MS/TN	19	3.89	UT	3	2.67
IA	18	4.11	CA	2	8.00
GA/VA/WV	14	3.79	GA/MD/VA/WV	2	5.00
OR/WA	13	3.00	MN/SD/NE	2	3.50
WA	12	3.83	SD/NE	2	3.00
MI/WI	11	3.30	CO/NE/WY	1	5.00
AZ/NM/NV	9	4.00	PA	1	8.00
CO/AZ/UT/NM	9	8.11			
GA/TN	9	4.00			

Notes: The table provides the number of observations and the mean number of active plants for each USGS region over the period 1974-2010. In total there are 61 regions and 933 region-year observations. We do not include regions that incorporate states and territories outside the continental United States. The mean number of plants is calculated based on the *Plant Information Survey* of the PCA and includes only plants with active active kiln capacity.

Plants sometimes list multiple primary fuels in the Plant Information Survey. There is little data available on the mix of primary fuels in those instances, however, and we allocate such plants based on a simple decision rule. We calculate fuel costs with the price of coal if coal is among the primary fuels. If not, we use petroleum coke prices if coke is among primary fuels. Otherwise we use natural gas prices if natural gas is among the multiple fuels. We use oil prices only if oil is the only fossil fuel listed. The exception to the above decision rule is when plants use a mix of coal and petroleum coke – there we assign equal weights to coal and petroleum coke prices. We have experimented with more sophisticated methodologies, leveraging data published in the *Minerals Yearbook* of the USGS on the total amounts of each fossil fuel burned by cement plants nationally. These methodologies are not fully satisfactory because, among other reasons, the USGS numbers include fuel burned (especially natural gas) to reheat kilns after maintenance periods. Our regression results are not sensitive to methodology on this subject and, given this, we prefer the simple rule.

Our methodology does not incorporate secondary fuels, the most popular of which are waste fuels such as solvents and used tires. The labor-energy input surveys of the PCA indicate that waste fuels account for around 25% of the energy used in wet kilns and 5% of the energy used in dry kilns. We do not have data on the prices of waste fuels but understand them to be lower on a per-mBtu basis than those of fossil fuels. Accordingly, we construct an alternative fuel cost measure in which we scale down the fossil fuel requirements of wet and dry kilns in accordance with the survey data. Whether this adjustment better reflects the fuel costs of marginal output depends in part on (i) the relative prices of waste and fossil fuels and (ii) whether the average fuel mix reported in the survey data reflect the marginal fuel mix. On the latter point, if marginal clinker output is fired with fossil fuels then our baseline measurement should reflect marginal fuel costs more closely than the alternative measurement. Regardless, our regression results are not very sensitive to the adjustment for waste fuels.

We obtain county-level data from the Census Bureau on construction employees and building permits to help control for demand. Construction employment is part of the County Business Patterns data. We identify construction as NAICS Code 23 and (for earlier years) as SIC Code 15. The data for 1986-2010 are available online.⁴⁴ The data for 1974-1985 are obtained from the University of Michigan Data Warehouse. The building permits data are maintained online by the U.S. Department of Housing and Urban Development.⁴⁵ We base the permits variable on the number of units so that, for example, a 2-unit permits

⁴⁴See <http://www.census.gov/econ/cbp/download/>, last accessed April 16, 2014.

⁴⁵See <http://socds.huduser.org/permits/>, last accessed April 16, 2014.

counts twice as much as a 1-unit permit. For both the construction employment and building permits, it is necessary to impute a small number of missing values. We calculate the average percentage difference between the observed data of each county and the corresponding state data, and use that together with the state data to fill in the missing values.

B Identification

We highlight here the sources of empirical variation that separately identify the own and cross pass-through parameters. We highlight here the sources of empirical variation that separately identify the own and cross pass-through parameters. The empirical variation we use to disentangle the own pass-through heterogeneity parameters (i.e., α_1) from the cross pass-through parameter (i.e., β), is straightforward – plants often have different fuel costs than their nearby competitors – and needs no further explanation. Instead, we focus on the empirical variation that distinguishes the baseline fuel cost parameter (i.e., α_0) from the cross pass-through parameter. There we can identify four distinct sources of identification (i) time-series variation in the distance metric, (ii) heterogeneity of capacity shares within a region, (iii) variation in fuel costs of plants in neighboring regions, and (iv) variation in the spatial composition of regions. We illustrate each source with simple examples below.

First, suppose that data consist of a single region and two plants with equal capacity. The linear approximation to regional prices then can be expressed

$$P_t = (\alpha_0 + \beta/d_{12t})\frac{c_{1t} + c_{2t}}{2} + \bar{\varepsilon}_t, \quad (\text{B.1})$$

where we have normalized $\alpha_1 = \gamma = 0$, without loss of generality. Absent inter-temporal variation in the distance metric, the coefficients α_0 and β are not separately identifiable. This remains true if more firms are incorporated, provided that plant capacity is homogeneous. However, time-series variation in the distance metric is sufficient for identification. Periods with greater effective plant dispersion (i.e., a bigger d_{12t}) exhibit lower rates of industry pass-through due to more muted cross pass-through. We introduce time-series variation in the distance metric by interacting the miles between plants with the gasoline price index.

Second, suppose that the distance metric is constant over time, but that capacities differ for the two plants in the single region. Regional prices then take the form

$$P_t = \alpha_0(\omega_1 c_{1t} + \omega_2 c_{2t}) + \beta/d_{12}(\omega_2 c_{1t} + \omega_1 c_{2t}) + \bar{\varepsilon}_t, \quad (\text{B.2})$$

The higher-capacity plant exercises greater influence on the own pass-through regressor, while the lower-capacity plant exercises greater influence on the cross pass-through regressor.⁴⁶ This is sufficient for identification, provided non-collinearity in the plants' fuel costs, which exists in regions containing plants that utilize different kiln technology. Identification through this channel becomes stronger with the inter-temporal changes in capacity weights that occurs with the retirement and introduction of kilns.

Third, the fuel costs of a plant can affect prices in a region even if the plant is not located in that region. Suppose that capacity shares of our two plants are equal, and the distance measure does not vary over time. Suppose further that we observe costs and distance for a third plant, denoted as plant 3, which is outside the region in the data. In this case, regional prices take the form:

$$P_t = \alpha_0(c_{1t} + c_{2t}) + \beta((1/d_{12})c_{1t} + (1/d_{12})c_{2t} + (1/d_{13} + 1/d_{23})c_{3t}) + \bar{\varepsilon}_t \quad (\text{B.3})$$

The third plant's fuel costs affect the cross pass-through regressor but not the own pass-through regressor, and this is sufficient for identification if the fuel costs of the third plant are not collinear with the fuel costs of the first two plants. Identification through this channel becomes stronger, the closer is the third plant to the first and second plants.

Turning to the final source of variation in the data, identification is assisted by having multiple regions in the data. Consider a case with two regions and four plants. Plants 1 and 2 are in region A and plants 3 and 4 are in region B. Stripping away all other sources of identifying variation, assume that capacity is homogeneous and constant, there is no inter-temporal variation in the distance metric, plants do not affect prices outside their region, and the fuel costs of all plants are equal and collinear. Regional prices then take the form

$$\begin{bmatrix} P_{At} \\ P_{Bt} \end{bmatrix} = \begin{bmatrix} \alpha_0 + \beta/d_{12} \\ \alpha_0 + \beta/d_{34} \end{bmatrix} c + \begin{bmatrix} \bar{\varepsilon}_{At} \\ \bar{\varepsilon}_{Bt} \end{bmatrix}. \quad (\text{B.4})$$

Identification is possible if $d_{12} \neq d_{34}$, as regions with greater plant dispersion exhibit lower rates of industry pass-through. Having multiple regions also amplifies the identifying variation available through the other channels enumerated above.

⁴⁶If capacity shares are equal then the two data vectors will be $0.5c_{1t} + 0.5c_{2t}$ and $(0.5c_{1t} + 0.5c_{2t})/d_{12}$, respectively, and collinearity causes identification to fail.

C Estimation Details for OLS and FGLS

Our approach to estimating clustered standard errors follows Wooldridge (2010), with small modifications to account for the fact that we use plant rather than region fixed effects. Note that if we used region fixed effects, then a standard within estimator could be used to consistently estimate the variance of the non fixed-effects coefficients. The logic of the within estimator is that one can apply a transformation to the data which removes the fixed effects, and then use standard techniques to consistently estimate a heteroskedasticity and autocorrelation consistent variance matrix for the remaining parameters. To see how this works denote ω as the matrix of plant-level capacity shares, where each row corresponds to an observation in our region level data. Similarly denote \mathbf{P} as the matrix of prices and \mathbf{X} as the matrix of non-fixed effects regressors from our regression equation (4). We apply the transformation

$$\mathbf{P}_\omega = \mathbf{I} - \omega(\omega'\omega)^{-1}\omega'$$

to the \mathbf{P} and \mathbf{X} matrices to obtain

$$\begin{aligned}\ddot{\mathbf{X}} &= \mathbf{P}_\omega \mathbf{X} \\ \ddot{\mathbf{Y}} &= \mathbf{P}_\omega \mathbf{Y}.\end{aligned}$$

A regression of $\ddot{\mathbf{Y}}$ on $\ddot{\mathbf{X}}$ will yield the OLS estimates of the \mathbf{X} coefficients by the Frisch-Waugh-Lovell theorem. One can construct robust standard errors using $\ddot{\mathbf{X}}$ and $\ddot{\mathbf{Y}}$ in the same manner one does using the within estimator. FGLS proceeds in a similar manner to that outlined in Chapter 10 of Wooldridge (2010); rather than using the within projection matrix one uses \mathbf{P}_ω .

D The Bayesian Regression Model

Our approach for Bayesian estimation of the spatial error structure is based on previous work outlined in Gelfand (2012) and Bakar and Sahu (2011). To start we assume a plant level model that looks like the following:

$$p_{it} = \mathbf{X}'_{it}\theta + \eta_{it} + \varepsilon_{it},$$

where

$$\eta_{it} = \kappa\eta_{i,t-1} + \nu_{it},$$

and θ is a vector of all the regression parameters from equation (4). Following the spatial literature, the *nugget effect*, ε_{it} , is i.i.d across plants and time and has a normally distribution with variance σ^2 . As in our prior work the ν_{it} is a spatially correlated error term. The spatio-temporal effects ν_{it} are assumed to be normally distributed and have a variance matrix of $\Sigma_\nu = \sigma_\nu^2 \mathbf{S}_\nu$. The elements of \mathbf{S}_ν are specified as

$$\psi(d_{ij}; \phi, \varphi) = \max \left\{ \frac{\varphi - d_{ij}}{\varphi}, 0 \right\}^\phi,$$

where ϕ and φ are restricted to be positive. This function will always be between 0 and 1, and when the distance is above φ is will be 0. The parameter φ determines when the correlation becomes 0. The parameter ϕ determines how quickly the correlation approaches zero as distance rises: the larger is ϕ , the more quickly the correlation drops off.

The regression equations we will be estimating will be the region level aggregates of the plant level equations. In particular the region level nugget effect will be

$$\bar{\varepsilon}_{mt} = \boldsymbol{\omega}'_{mt} \boldsymbol{\varepsilon}_{mt},$$

and the spatial effect will be

$$\bar{\boldsymbol{\eta}}_{mt} = \boldsymbol{\omega}'_{mt} \boldsymbol{\eta}_{mt},$$

where the bolded $\boldsymbol{\eta}$'s and $\boldsymbol{\varepsilon}$'s contain the errors of all the plants in region m , and $\boldsymbol{\omega}$ is the vector of capacities.

We allow period 0 spatial errors to follow a different process than errors in periods 1 through T , and specify the period 0 variance matrix as $\sigma_0^2 \mathbf{S}_\nu$.

An attractive feature of Bayesian estimation relative to classical approaches is that we can treat the $\boldsymbol{\eta}$'s as random effects and draw them alongside our other parameters.⁴⁷ The steps will be

1. Draw θ given the data and draws on σ_ε^2 , and η_{it} 's.
2. Draw κ given σ_ν^2 , ϕ , φ and η_{it} 's.
3. Draw σ_ε^2 given data, β , and η_{it} 's.

⁴⁷Note that if we took a classical approach such as Maximum Likelihood we would have to work with the distributions of $\bar{\varepsilon}$, which would be quite messy.

4. Draw σ_ν^2 given κ, ϕ, φ and η_{it} 's.
5. Draw η_{it} given other period η_{it} 's, $\sigma_\nu^2, \phi,$ and φ .
6. Draw σ_0^2 given η_0 .
7. Draw η_0 given $\sigma_0^2, \kappa,$ and η_1
8. Draw ϕ given everything, using Metropolis-Hastings.
9. Draw φ given everything, using Metropolis-Hastings.

We assume standard priors wherever possible (normal for β and κ , and inverse gamma for σ_ε^2 and σ_ν^2). For most parameters, we choose the values of priors to be consistent with the default values used in the univariate regression R code provided with the book Rossi, Allenby, and McCulloch (2005).

Below we describe the conditional posteriors for each parameter and describe how we draw from them.

D.1 Posterior for θ

Here we will use standard Bayesian regression. We assume a prior mean of $\bar{\theta} = 0$ and variance of δ_θ^2 is set to be 0.01 times the dimensionality of θ . The posterior distribution is $N(\Delta\chi, \Delta)$ where

$$\begin{aligned}\Delta^{-1} &= \sum_{m=1}^M \sum_{t \in T_m} \frac{\bar{\mathbf{X}}'_{mt} \bar{\mathbf{X}}_{mt}}{\boldsymbol{\omega}'_{mt} \boldsymbol{\omega}_{mt}} \frac{1}{\sigma_\varepsilon^2} + \mathbf{I}_N / \delta_\theta^2 \\ \chi &= \bar{\theta} / \delta_\theta^2 + \sum_{m=1}^M \sum_{t \in T_m} \frac{\bar{\mathbf{X}}'_{mt} (\bar{p}_{mt} - \bar{\eta}_{mt})}{\boldsymbol{\omega}'_{mt} \boldsymbol{\omega}_{mt}} \frac{1}{\sigma_\varepsilon^2}\end{aligned}$$

Note that we compute $\bar{\eta}_{it}$ from plant level draws on η_{it} .

D.2 Posterior for κ

A Bayesian regression can also be used to construct a posterior for κ . To do this, we stack up the plant level η_{it} 's into vectors $\boldsymbol{\eta}_t$. Assuming prior parameters on κ are $\bar{\kappa}$ and δ_κ^2 , the posterior will be $N(\Delta\chi, \Delta)$ where

$$\begin{aligned}\Delta^{-1} &= \sum_{t=1}^T \boldsymbol{\eta}'_t \boldsymbol{\Sigma}_\nu^{-1} \boldsymbol{\eta}_t + \mathbf{I}_N / \delta_\kappa^2 \\ \chi &= \bar{\kappa} / \delta_\kappa^2 + \sum_{t=2}^T \boldsymbol{\eta}'_{t-1} \boldsymbol{\Sigma}_\nu^{-1} \boldsymbol{\eta}_t\end{aligned}$$

We set the prior mean and variance for κ to 0.5 and 0.01, respectively.

D.3 Posterior for σ_ε^2

Assuming an inverse gamma prior with parameters a and b we obtain an inverse gamma posterior

$$\pi(1/\sigma_\varepsilon^2; \dots) = G\left(\frac{N}{2} + a, b + \frac{1}{2} \sum_{m=1}^M \sum_{t \in T_m} \frac{(p_{mt} - \bar{\eta}_{mt} - \bar{\mathbf{X}}'_{mt} \theta)^2}{\boldsymbol{\omega}'_{mt} \boldsymbol{\omega}_{mt}}\right).$$

The prior parameters a and b are both set to 0.001.

D.4 Posterior for σ_ν^2

This case is a lot like the previous one, except we treat the η 's like data. Assuming prior parameters of a and b (both set to 0.001) we obtain an inverse gamma posterior

$$\pi(1/\sigma_\nu^2; \dots) = G\left(\frac{N}{2} + a, b + \frac{1}{2} \sum_{t=1}^T (\boldsymbol{\eta}_t - \kappa \boldsymbol{\eta}_{t-1})' \mathbf{S}_\nu^{-1} (\boldsymbol{\eta}_t - \kappa \boldsymbol{\eta}_{t-1})\right).$$

D.5 Posterior for η_{it} 's

To derive this posterior we will begin by writing down the part of the joint posterior likelihood that depends on η_{it} . Note that we can write the aggregate variables in period t as a matrix of capacity shares, $\boldsymbol{\omega}_t$, times the variable (which could be plant level prices, \mathbf{p}_t , spatial effects, $\boldsymbol{\eta}_t$, etc.). This will be helpful when we compute the kernel of the distribution for $\boldsymbol{\eta}_t$. The part of the posterior that depends on η_{it} can then be written

$$\begin{aligned}
L &= -\frac{1}{2\sigma_\varepsilon^2} \sum_{t=1}^T (\mathbf{p}_t - \mathbf{X}'_t \theta - \boldsymbol{\eta}_t)' \boldsymbol{\omega}'_t \mathbf{D}_{\omega,t} \boldsymbol{\omega}_t (\mathbf{p}_t - \mathbf{X}'_t \theta - \boldsymbol{\eta}_t) \\
&\quad - \frac{1}{2\sigma_\nu^2} \sum_{t=1}^T (\boldsymbol{\eta}_t - \kappa \boldsymbol{\eta}_{t-1})' \mathbf{S}_\nu^{-1} (\boldsymbol{\eta}_t - \kappa \boldsymbol{\eta}_{t-1}) - \dots
\end{aligned}$$

The matrix $\mathbf{D}_{\omega,t}$ is a diagonal matrix where each diagonal entry corresponds to the inverse of the sum of squared capacity shares for region m in period t . To compute the posterior for η_{it} we have to consider three cases. The first, and simplest case is when $t = T$, the last period. To solve for the posterior for η_{iT} we will complete the square on the posterior as follows:

$$(v - W\boldsymbol{\eta}_t)'(v - W\boldsymbol{\eta}_t) = (\boldsymbol{\eta}_t - \hat{\boldsymbol{\eta}}_t)' W' W (\boldsymbol{\eta}_t - \hat{\boldsymbol{\eta}}_t) + \dots$$

where

$$v = \begin{bmatrix} \frac{1}{\sigma_\varepsilon} \mathbf{D}_{\omega,t}^{1/2} \boldsymbol{\omega}_t (\mathbf{p}_t - \mathbf{X}'_t \theta) \\ \kappa \boldsymbol{\Sigma}_\nu^{-1/2} \boldsymbol{\eta}_{t-1} \end{bmatrix} \quad W = \begin{bmatrix} \frac{1}{\sigma_\varepsilon} \mathbf{D}_{\omega,t}^{1/2} \boldsymbol{\omega}_t \\ \boldsymbol{\Sigma}_\nu^{-1/2} \end{bmatrix}.$$

The mean $\hat{\boldsymbol{\eta}}_t$ can be expressed as

$$(W'W)^{-1} W'v = \left(\frac{1}{\sigma_\varepsilon^2} \boldsymbol{\omega}'_t \mathbf{D}_{\omega,t} \boldsymbol{\omega}_t + \kappa^2 \boldsymbol{\Sigma}_\nu^{-1} \right)^{-1} \left(\frac{1}{\sigma_\varepsilon^2} \boldsymbol{\omega}'_t \mathbf{D}_{\omega,t} \boldsymbol{\omega}_t (\mathbf{p}_t - \mathbf{X}'_t \theta) + \kappa \boldsymbol{\Sigma}_\nu^{-1} \boldsymbol{\eta}_{t-1} \right)$$

and the variance of the $\boldsymbol{\eta}_t$ will be

$$(W'W)^{-1} W'v = \left(\frac{1}{\sigma_\varepsilon^2} \boldsymbol{\omega}'_t \mathbf{D}_{\omega,t} \boldsymbol{\omega}_t + \kappa^2 \boldsymbol{\Sigma}_\nu^{-1} \right)^{-1}.$$

The second case, for $t < T$, is more involved because $\boldsymbol{\eta}_t$ shows up more than once in the lag equation. In this case the variance should be

$$(W'W)^{-1} W'v = \left(\frac{1}{\sigma_\varepsilon^2} \boldsymbol{\omega}'_t \mathbf{D}_{\omega,t} \boldsymbol{\omega}_t + (1 + \kappa^2) \boldsymbol{\Sigma}_\nu^{-1} \right)^{-1},$$

and the mean should be

$$\left(\frac{1}{\sigma_\varepsilon^2}\boldsymbol{\omega}'_t\mathbf{D}_{\omega,t}\boldsymbol{\omega}_t + (1 + \kappa^2)\boldsymbol{\Sigma}_\nu^{-1}\right)^{-1}\left(\frac{1}{\sigma_\varepsilon^2}\boldsymbol{\omega}'_t\mathbf{D}_{\omega,t}\boldsymbol{\omega}_t(\mathbf{p}_t - \mathbf{X}'_t\boldsymbol{\theta}) + \kappa\boldsymbol{\Sigma}_\nu^{-1}\boldsymbol{\eta}_{t-1} + \kappa\boldsymbol{\Sigma}_\nu^{-1}\boldsymbol{\eta}_{t+1}\right).$$

The third case is $t = 0$. The variance of $\boldsymbol{\eta}_0$ will be equal to $\boldsymbol{\Sigma}_\nu/(1/\sigma_0^2 + 1/\sigma_\nu^2)$, and the mean will be $\kappa\boldsymbol{\eta}_1(1/\sigma_\nu^2)/(1/\sigma_0^2 + 1/\sigma_\nu^2)$.

D.6 Posteriors for ϕ and φ

We draw these parameters using the Metropolis-Hastings algorithm, assuming uninformative priors for each parameter. We use a random walk Metropolis-Hastings algorithm with a normal proposal distribution. Given last iteration's draw on ϕ , which we denote ϕ_0 , we draw a candidate $\phi_1 \sim N(\phi_0, \sigma_\phi^2)$. Then we compute the posterior log likelihood at both ϕ_0 and ϕ_1 , where this log-likelihood is (suppressing terms that don't depend on ϕ or φ):

$$\begin{aligned} L(\phi, \varphi) = & -\frac{1}{2} \sum_{t=1}^T \log |\boldsymbol{\Sigma}_\nu(\phi, \varphi)| - \frac{1}{2\sigma_\nu^2} \sum_{t=1}^T (\boldsymbol{\eta}_t - \kappa\boldsymbol{\eta}_{t-1})' \mathbf{S}_\nu(\phi, \varphi)^{-1} (\boldsymbol{\eta}_t - \kappa\boldsymbol{\eta}_{t-1}) \\ & - \frac{1}{2} \log |\sigma_0^2 \mathbf{S}_\nu(\phi, \varphi)| - \frac{1}{2} \frac{1}{\sigma_0^2} \boldsymbol{\eta}'_0 \mathbf{S}_\nu(\phi, \varphi) \boldsymbol{\eta}_0. \end{aligned}$$

We then accept the candidate draw with probability

$$\min \left\{ 1, \frac{\exp(L(\phi^1, \varphi))}{\exp(L(\phi^0, \varphi))} \right\}.$$

The parameter σ_ϕ^2 is set so that the acceptance rate is about 50%. We perform a similar Metropolis-Hastings step for φ . Additionally, since ϕ and φ are constrained to be positive we reject any draws that are below zero.

E First Order Approximation

We sketch in this appendix the mathematics of first order approximation (FOA) as it pertains to merger price effects. Greater detail is provided in Jaffe and Weyl (2013) and Miller, Remer, Ryan, and Sheu (2013). The starting point for FOA is the first order condition that characterizes profit-maximization. Let cement firms set free-on-board prices to maximize

profit, taking as given the prices of other firms. Then the first order conditions of any firm i can be expressed

$$f_i(P) \equiv - \left[\frac{\partial Q_i(P)^T}{\partial P_i} \right]^{-1} Q_i(P) - (P_i - MC_i) = 0, \quad (\text{E.1})$$

where P_i is a vector of firm i 's plant prices, P is a vector of all prices, $Q_i(P)$ is a demand schedule and MC_i is a vector of firm i 's plant marginal costs. The post-merger first order conditions then can be expressed

$$h_i(P) \equiv f_i(P) + g_i(P) = 0, \quad (\text{E.2})$$

where, for a merger of firms j and k ,

$$g_j(P) = - \underbrace{\left(\frac{\partial Q_j(P)^T}{\partial P_j} \right)^{-1} \left(\frac{\partial Q_k(P)^T}{\partial P_j} \right)}_{\text{Diversion from } j \text{ to } k} \underbrace{(P_k - MC_k^1)}_{\text{Markup of } k}, \quad (\text{E.3})$$

the form of $g_k(P)$ is analogous, and $g_i(P) = 0$ for $i \neq j, k$. The g function captures the opportunity costs, or “upward pricing pressure,” created by the merger.⁴⁸ Notice that it enters the post-merger first order conditions in the same way as a cost shock. To a first order approximation, the resulting price changes equal

$$\Delta P = - \left(\frac{\partial f(P)}{\partial P} + \frac{\partial g(P)}{\partial P} \right)^{-1} \Bigg|_{P=P^0} g(P^0), \quad (\text{E.4})$$

where P^0 is the vector of pre-merger equilibrium prices. In this expression, the Jacobian of the post-merger first order conditions – “merger pass-through” – depends on the first and second derivatives of demand. Given knowledge of the first derivatives, it is possible to infer the second derivatives from cost pass-through based on the formula

$$\rho = - \left(\frac{\partial f(P)}{\partial P} \right)^{-1} \quad (\text{E.5})$$

⁴⁸Each firm in the merger, when making a sale, forgoes with some probability a sale by the other firm. The diversion matrix represents the fraction of sales lost by firm j 's products that shift to firm k 's products due to an increase in firm j 's prices. When multiplied by firm k 's markups, this yields the value of diverted sales; the more these sales are worth, the greater incentive a firm has to raise price following a merger.

Merger pass-through then can be calculated with the first and second demand derivatives. In our application, the large number of plants makes it numerically difficult to identify the second derivatives. We instead use cost pass-through to proxy merger pass-through. Price predictions then are based on the matrix multiplication of the pass-through matrix and the vector of upward pricing pressure. This simplification is proposed in Jaffe and Weyl (2013) and shown in Miller, Remer, Ryan, and Sheu (2013) to cause little loss of predictive accuracy.

F EPA Analysis of NESHAP Ammendments

The EPA relies on a Cournot model of competition to simulate the effect of regulation in each of 20 local markets based on conditions in 2005. The model incorporates a constant elasticity market demand curve and, for markets that are adjacent to a port, a constant elasticity import supply curve. It is calibrated to elasticity estimates in the existing literature. We provide details on the model here. After the implementation of regulation, the first order conditions of firm i can be expressed

$$dMC_i = dP \left[1 + \frac{s_i}{\eta} \right] + dq_i \left[\frac{P}{\eta} \frac{1}{Q} \right] - dQ \left[\frac{P}{\eta} \frac{q_i}{Q^2} \right], \quad (\text{F.1})$$

P is the market price, s_i is the share of sales for plant i , q_i is the quantity sold by plant i , Q is market consumption including imports, MC is marginal cost, and η is the elasticity of consumption with respect to price. Thus the object dMC is the compliance cost of regulation. Equation F.1 governs how compliance costs, represented by dMC_i , affect output and, in turn, market price. Imports are supplied according to an elasticity ϕ , such that

$$dI = \phi \left(\frac{dP}{P} \right) I, \quad (\text{F.2})$$

where I is the quantity of imports. Total consumption in a market (again including imports) evolves according to

$$dQ = \eta \left(\frac{dP}{P} \right) Q. \quad (\text{F.3})$$

Finally, the model is closed with supply equaling demand,

$$dQ = \sum_i dq_i + dI. \quad (\text{F.4})$$

The EPA calibrates the model with a price elasticity of consumption of 0.88, based on EPA (1998), an import elasticity of 2.0, based on Broda, Limao, and Weinstein (2008). Prices and plant-level production are calculated by manipulating the region-level data published in the *Minerals Yearbook* of the USGS, following a methodology that is detailed in Section A.1 of EPA (2009). We are able to replicate the calibration process exactly so that discrepancies between our predictions and those of the EPA are due solely to the decision of the EPA not to publish plant-level compliance costs.