

Spatial differentiation and price discrimination in the cement industry: evidence from a structural model

Nathan H. Miller*

and

Matthew Osborne**

We estimate a structural model of the cement industry that incorporates spatial differentiation and price discrimination, focusing on the US Southwest over 1983–2003. We leverage the structure of the model to obtain consistent estimates of the underlying parameters using data on market outcomes that are substantially aggregated. Our results indicate that transportation costs around \$0.46 per tonne-mile rationalize the data. This friction enables relatively isolated plants to obtain higher prices from nearby customers. We further find that disallowing price discrimination would create \$30 million in consumer surplus annually and show how the model can identify suitable divestitures in merger analysis.

1. Introduction

■ In many industries, firms are differentiated in geographic space, and transportation is costly. Seminal theoretical contributions demonstrate that these conditions can soften the intensity of competition, facilitate markups above marginal cost, and induce firms to discriminate among consumers based on location (Hotelling, 1929; Salop, 1979; Anderson and de Palma, 1988; Vogel, 2008). The empirical literature of industrial organization, however, only recently has grappled with the estimation of economic models that capture realistically the salient features of competition that arise with spatial differentiation.

We estimate a structural model of the cement industry that incorporates spatial differentiation and spatial price discrimination, focusing on the US Southwest over 1983–2003. The cement

*Georgetown University; nhm27@georgetown.edu.

**University of Toronto; matthew.osborne@rotman.utoronto.ca.

We thank Jim Adams, Simon Anderson, Allan Collard-Wexler, Abe Dunn, Masakazu Ishihara, Robert Johnson, Scott Kominers, Ashley Langer, Russell Pittman, Chuck Romeo, Jim Schmitz, Adam Shapiro, Charles Taragin, Raphael Thomadsen, Glen Weyl and seminar participants at the Bureau of Economic Analysis, the Federal Trade Commission, George Washington University, Georgetown University, the University of British Columbia — Sauder School of Business, University of Virginia, and the US Department of Justice for valuable comments. Carl Fisher, Sarolta Lee, Parker Sheppard, Gian Wrobel, and Vera Zavin provided research assistance.

industry has been analyzed frequently in the industrial economics literature in part because high shipping costs engender localized competition that is amenable to empirical analysis (e.g., Syverson and Hortaçsu, 2007; Salvo, 2010a; Ryan, 2012).¹ We build on this literature by explicitly modelling transportation costs and allowing firms to discriminate based on customer location; the existing empirical literature does not address spatial discrimination despite the long and litigious history of discrimination in the industry. Our results characterize how transportation costs affect economic outcomes such as mill prices, plant production, and the trade flows that arise within and across geographic regions.

Our results indicate that transportation costs around \$0.46 per tonne-mile rationalize the observed data given the structure of the model. We find an average shipping distance of 122 miles for the year 2003, which is small relative to the distances in the US Southwest (e.g., 372 miles separate Los Angeles and Phoenix). This market friction enables plants that are relatively isolated geographically to charge higher mill prices to nearby customers. Our model estimates imply that the mill prices and margins of these plants decline quickly with distance to the customer. By contrast, plants located nearby many other plants have less localized market power and do not appear to discriminate based on geography.

We conduct two counterfactual experiments. First, we find that disallowing spatial discrimination would increase consumer surplus by nearly \$30 million annually, relative to a volume of commerce of \$1.3 billion. The effects of such a ban would vary widely across the US Southwest, with regions near cement plants benefitting and more isolated customers being harmed. Although the theoretical literature has long recognized that spatial price discrimination can increase or decrease consumer surplus (e.g., Gronberg and Meyer, 1982; Katz, 1984; Hobbs, 1986; Anderson, de Palma, and Thisse 1989), to our knowledge we provide the first empirical evidence on the topic. Second, we evaluate a hypothetical merger between the two largest portland cement manufacturers in the US Southwest in 1986 to illustrate how our approach could be used to illuminate the competitive effects of mergers in industries characterized by high transportation costs. We find that the hypothetical merger would have increased prices in Southern California and Arizona by 4.9% and 9.8%, respectively. By contrast, a standard Cournot model predicts price increases of 1% in Southern California and 25% in Arizona, demonstrating that incorporating spatial considerations can have a meaningful impact on counterfactual price predictions.

To generate these results, we develop an estimation strategy that is implementable with data on market outcomes, such as mill prices and production, that are substantially aggregated. The strategy allows us to proceed despite not observing the prices that are charged to specific consumer locations—a data availability problem that makes standard structural estimation techniques inapplicable. The estimation strategy involves modelling cement demand at the county level, where measures of market size are available, and using familiar minimum distance techniques to choose supply and demand parameters that rationalize the data. For each candidate parameter vector we compute the equilibrium prices that plants charge in each county and then aggregate the equilibrium predictions to the level of the data. This makes it possible to identify the parameter vector that brings the predicted aggregate moments as close as possible to the observed aggregate moments.² The equilibrium price vector is high dimensional because a typical year in our data has 14 plants and 90 counties, and we ease the computational burden of repeatedly solving for an equilibrium using the recently developed numerical techniques of La Cruz, Martíñex, and Raydan

¹ Earlier articles include Rosenbaum and Reading (1988), Rosenbaum (1994), Jans and Rosenbaum (1997), and Newmark (1998).

² Parallels can be drawn between our estimator and existing estimators for models in which firms are differentiated in product space (e.g., Berry, Levinsohn, and Pakes, 1995; Nevo, 2001). With product differentiation, the challenge is to recover structural parameters when prices and quantities are observed but nonprice product characteristics are imperfectly observed. With spatial differentiation, by contrast, the challenge is that prices and quantities are imperfectly observed. In either case, numerical techniques allow the recovery of the unobservable: the contraction mapping of Berry (1994) obtains the unobserved product characteristic with product differentiation, just as computing equilibrium obtains prices and quantities with spatial differentiation.

(2006). In a companion article, we provide conditions under which the obtained estimates are consistent and asymptotically normal (Miller and Osborne, 2014).

The estimation strategy we employ utilizes the information contained in the structure of the economic model, as is standard for structural research in industrial economics. It follows that how the competition is modelled can affect results. For instance, we impose the assumption that demand in each county is characterized by the nested logit model. This facilitates estimation by easing the burden of repeatedly computing equilibrium and by smoothing the objective function—but it also ensures that competition is global (i.e., plants ship some cement to even distant counties) and has implications for how the model predicts market shares and prices to vary across the counties of the US Southwest. In robustness checks, we evaluate how changing the variance of consumers' idiosyncratic tastes (the "logit error") affects implied shipping distances and prices and determine that these predictions are materially similar for a range of assumptions. Nonetheless, the trade-off is fundamental and our results should be understood in context.

The articles closest to ours estimate models of spatial differentiation in nondiscriminatory settings, including fast-food restaurants (Thomadsen, 2005), movie theaters (Davis, 2006), coffee shops (McManus, 2007), and retail gasoline (Houde, 2012). The main methodological distinction is that their estimation strategies require the observation of all relevant prices whereas ours does not. Such prices are rarely available for industries, such as the cement industry, characterized by business-to-business sales and privately negotiated contracts. The estimation strategy we introduce could extend the reach of researchers to cover these settings. Our work also relates to that of Pinske, Slade, and Brett (2002), which introduces a reduced-form estimator that can evaluate whether competition is localized but does not recover the underlying structural parameters of an economic model. Finally, our work can be related to the literature on auctions when producers have transportation costs, including the research of Porter and Zona (1999), which examines the spatial pattern of bids to identify collusion in milk markets.³

The article proceeds as follows. In Section 2, we outline the institutional details of the portland cement industry in the US Southwest, describe the available data, and provide reduced-form evidence about the role of transportation costs in creating localized market power. In Section 3, we formalize a model of spatial price discrimination that is tailored to the salient features of the cement industry, including capacity constraints and import competition. In Section 4, we derive the estimator and showcase the empirical variation that drives the parameter estimates. We present the results of estimation in Section 5, with a particular focus on geographic patterns of market shares and mill prices. The results of the two counterfactual experiments appear in Section 6. Section 7 concludes.

2. The portland cement industry

□ **Institutional details.** Portland cement is a finely ground dust that forms concrete when mixed with water and coarse aggregates such as sand and stone.⁴ Concrete, in turn, is an essential input to many construction and transportation projects. The production of cement involves feeding limestone into coal-fired rotary kilns that reach peak temperatures of 1400–1450° Celsius. The output of the kilns, "clinker," is mixed with a small amount of gypsum and ground to form portland cement. Kilns operate at peak capacity except for an annual maintenance period, the duration of which can be adjusted according to demand conditions. The five main variable costs are due to materials, coal, electricity, labor, and maintenance (EPA, 2009).

Spatial price discrimination has a long history in the industry. Cement producers used basing-point pricing from 1902 to 1948, when the Supreme Court determined in *FTC vs. Cement*

³ Porter and Zona (1999) finds that milk producers not engaged in collusion bid lower at nearby districts and higher at distant districts, consistent with transportation costs, but that colluding producers bid higher at nearby districts (which were targeted by the cartel).

⁴ We draw heavily from the publicly available documents and publications of the United States Geological Survey and the Portland Cement Association to support the analysis in this section.

Institute that this facilitated coordinated conduct among competitors in violation of the antitrust laws.⁵ Under basing-point pricing, delivered prices depend on prices at some publicized location (the base) adjusted for shipping costs from the base to the customer. Cement producers often used the location of a competing center of production as the base, yielding higher prices for customers with less attractive outside options. That prices that were sometime lower to customers farther away from plants was one count in the complaint of *FTC vs. Cement Institute*.⁶

Cement producers now privately negotiate contracts with their customers. These contracts specify a mill price (or a “free-on-board” price) and discounts that reflect the willingness and ability of customers to substitute toward cement produced by competitors. This enables producers to price discriminate among its customers without running afoul of the antitrust laws as interpreted by the courts.⁷ Purchasers are responsible for the transportation of cement. The bulk of portland cement is moved by truck, although some is sent by train or barge to distribution terminals and only then trucked to customers.⁸ Transportation accounts for a substantial portion of purchasers’ total acquisition costs because portland cement is inexpensive relative to its weight.⁹

In some cases, the contracting process is made unnecessary by the vertical integration of cement and ready-mix concrete plants. Syverson and Hortaçsu (2007) document two distinct waves of vertical mergers and acquisitions in the industry, arising in the early to mid 1960s and over 1982–1992, respectively, and separated by a lengthy period of antitrust scrutiny by the Federal Trade Commission (FTC). In 1997, vertically integrated cement producers accounted for 55% of cement sales. Syverson and Hortaçsu (2007) determine that vertical integration has little causal impact on plant- and market-level outcomes, however, and we abstract from such relationships in the analytical framework we introduce below.¹⁰

□ **Cement in the US Southwest.** We focus the empirical application on California, Arizona, and Nevada, which we refer to collectively as the US Southwest. This eases the computational burden of repeatedly solving for equilibrium, which increases quickly in the number of plants and counties under consideration. As we develop below, the cement industry in the US Southwest is insulated from competition from other domestic areas, making the region well suited for our analysis.

Figure 1 maps the geographic configuration of the industry in the US Southwest, circa 2003, based on *Plant Information Survey*, an annual publication of the Portland Cement Association. Most plants are located along an interstate highway, nearby one or more population centers. Foreign imports enter through four customs offices, located in San Francisco, Los Angeles,

⁵ *FTC vs. Cement Institute*, 37 FTC 87 (1943). See Karlson (1990) for a detailed account of basing point pricing in the cement industry. Prices were initially based on distance from Lehigh Valley, the first location known to have rock deposits suitable for making cement. As the industry expanded geographically the number of basis points proliferated and, by 1940, there were over 50 basis points in the United States.

⁶ This case and contemporaneous antitrust actions motivated economists to investigate theoretically the merits of basing-point pricing. Although a review of the resulting literature is beyond the scope of this article, we point interested readers to Kaysen (1949), Haddock (1982), DeCanio (1984), and Thisse and Vives (1988) as useful starting points.

⁷ The relevant case law focuses on whether spatial price discrimination evidences coordinated conduct. See, for example, *Cement Mfrs. Protective Assn. vs. United States*, 268 US 588 (1925); *Maple Flooring Mfrs. Assn. vs. United States*, 268 US 563 (1925); *Sugar Institute vs. United States*, 297 US 553 (1936); *FTC vs. Cement Institute*, 37 FTC 87 (1943); *Triangle Conduit & Cable Co. vs. FTC*, 168 F.2d 175 (7 Cir., 1948); *Boise Cascade vs. FTC*, 637 F.2d 573, 581–82 (9 Cir., 1980); *FTC vs. Ethyl et al.*, 101 FTC 425, final order March 1983; *E.I. DuPont de Nemours & Co. vs. FTC*, 729 F. 2d 128 (2d Cir., 1984).

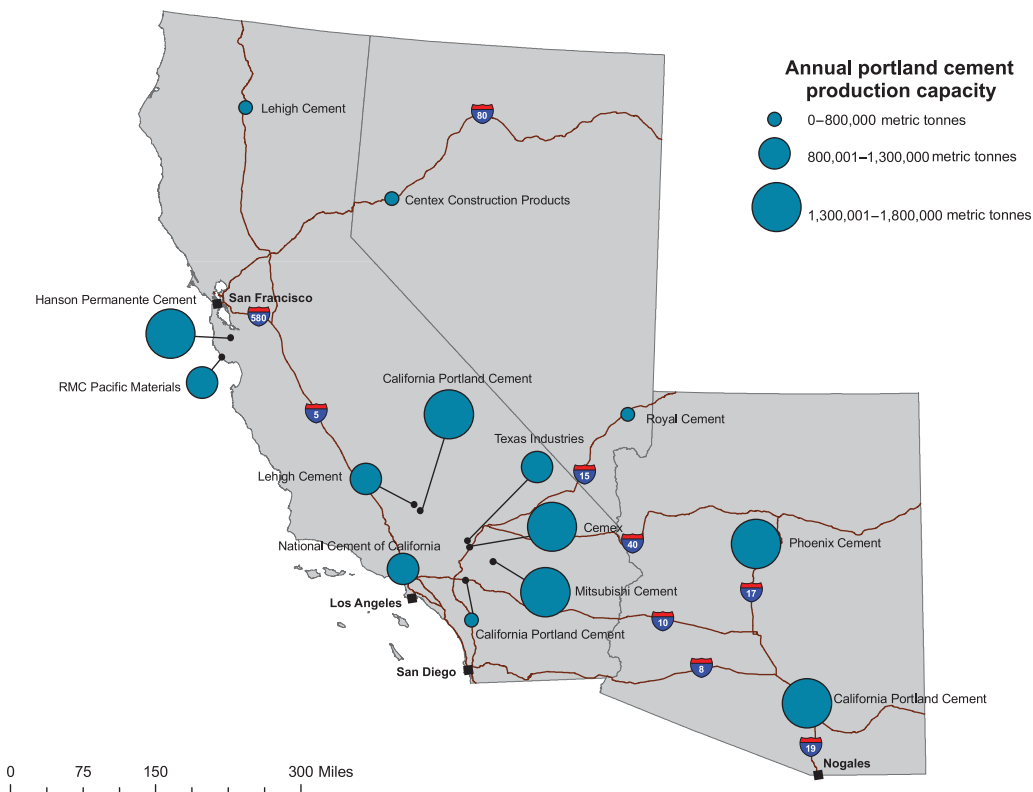
⁸ Barge transport is not feasible in the US Southwest due to the lack of navigable rivers.

⁹ Scherer et al. (1975) calculates that transportation would account for roughly one third of total customer expenditures on a hypothetical 350-mile route between Chicago and Cleveland, and a US Census Bureau study (1977) reports that more than 80% is transported within 200 miles. More recently, Salvo (2010b) presents evidence consistent with the importance of transportation costs in the Brazilian portland cement industry.

¹⁰ The subsequent research of Enghin, Syverson, and Hortaçsu (2012) examines comprehensively the shipping practices of vertically integrated firms in manufacturing sectors and calls into question whether vertical integration typically is used to facilitate the transfer of goods along the production chain.

FIGURE 1

PORTLAND CEMENT PLANTS IN THE US SOUTHWEST, CIRCA 2003



San Diego, and Nogales. Foreign imports are mostly produced by large, efficient plants located in Southeast Asia. Exports from the US Southwest to foreign markets are negligible because domestic plants are not competitive in the international market.


Figure 2 plots domestic consumption and production in the US Southwest over 1983–2003. Consumption and production are highly procyclical because demand is tied to construction. That consumption is *more* procyclical than domestic production is due to the costliness of capacity adjustments as documented in Ryan (2012). The figure also shows that foreign imports match, nearly exactly, the gap between consumption and domestic production (“apparent imports”), consistent with no meaningful trade flows between the US Southwest and other domestic regions.¹¹ Cement can be shipped economically into the area from Asia but not from other domestic areas due to the cost discrepancy between freighter and truck transportation and the relative efficiency of the large foreign plants.

□ **Summary statistics and reduced-form evidence.** Our primary data source is the *Minerals Yearbook* of the US Geological Survey (USGS). The USGS conducts an annual establishment-level census of portland cement producers and publishes aggregated statistics on consumption, production, and revenues in its *Minerals Yearbook*. We also make use of data on cross-region

¹¹ Other statistics published by the USGS corroborate this interpretation. More than 98% of cement produced in southern California was shipped within the US Southwest over 1990–1999, and more than 99% of cement produced in California was shipped within the region over 2000–2003. Outflows from Arizona and Nevada are unlikely because consumption routinely exceeds production in those states. As net trade flows between the US Southwest and other domestic regions are insubstantial, these data points imply that gross domestic inflows must also be insubstantial.

FIGURE 2

CONSUMPTION, PRODUCTION, AND IMPORTS OF PORTLAND CEMENT



Notes: Apparent imports are defined as consumption minus production. Observed imports are total foreign imports shipped into San Francisco, Los Angeles, San Diego, and Nogales.

TABLE 1 Consumption, Production, and Prices

Description	Mean	Standard	Minimum	Maximum
<i>Consumption</i>				
Northern California	3513	718	2366	4706
Southern California	6464	1324	4016	8574
Arizona	2353	650	1492	3608
Nevada	1289	563	416	2206
<i>Production</i>				
Northern California	2548	230	1927	2894
Southern California	6316	860	4886	8437
Arizona-Nevada	1669	287	1050	2337
<i>Domestic prices</i>				
Northern California	85.81	11.71	67.43	108.68
Southern California	82.81	16.39	62.21	114.64
Arizona-Nevada	92.92	14.24	75.06	124.60
<i>Import prices (excludes duties and grinding costs)</i>				
US Southwest	50.78	9.30	39.39	79.32

Notes: Statistics are based on observations at the region-year level over the period 1983–2003. Production and consumption are in thousands of metric tonnes. Prices are per metric tonne, in real 2000 dollars. Import prices exclude duties. The region labeled “Arizona-Nevada” incorporates information from Nevada plants only over 1983–1991.

shipments from the *California Letter*, another annual publication of the USGS, and data on plant locations and kiln capacities from the *Plant Information Survey*. Finally, we collect data from the Energy Information Agency (EIA) on coal, electricity, and diesel prices. Our sample period of 1983–2003 reflects the availability of the *Plant Information Survey*.¹² We provide additional details on the data sample in Appendix A.

Table 1 provides sample statistics on the consumption, production, and average prices of cement in the US Southwest over the sample period.¹³ Consumption is available separately

¹² We were unable to obtain freely the *Plant Information Survey* for years after 2003.

¹³ We calculate average prices as the ratio of revenues to production.

TABLE 2 Reduced-Form Evidence on Localized Market Power

Description	(1)	(2)
<i>Measure of region isolation</i>		
Distance to customs office		0.114** (0.06)
<i>Region indicator variables</i>		
Northern California	3.01** (1.29)	1.68 (1.34)
Arizona-Nevada 1983–1991	6.47*** (1.99)	-7.75 (6.66)
Arizona-Nevada 1992–2003	12.84*** (1.43)	4.41 (4.25)
Year fixed effects	yes	yes
Adjusted- R^2	0.9182	0.9265

Notes: Results of OLS regressions based on 63 observations at the region-year level. The dependent variable is the natural log of the average domestic price. Distance to customs office is the average number of miles between the region's counties and the nearest customs offices (located in San Francisco, Los Angeles, San Diego, and Nogales), interacted with the price of diesel fuel. Standard errors are adjusted for heteroskedasticity and are shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

for northern California, southern California, Arizona, and Nevada. The level of aggregation for production and prices follows the policy of the USGS to include at least three independently owned plants in each reporting region. This “rule of three” protects the confidentiality of responses to the establishment-level census because the production and prices of one plant cannot be inferred from the *Minerals Yearbook* and knowledge of another plant's production and prices. We report statistics separately for northern California, southern California, and a composite Arizona-Nevada region. This composite region includes information from Nevada only over 1983–1991 because the USGS combines Nevada with three states outside the US Southwest starting in 1992.¹⁴ The price of imported clinker, which we calculate as a weighted average of the prices at the four customs offices of the US Southwest, does not incorporate import duties or the cost of grinding the clinker into cement.¹⁵

The variation in these aggregated data, even without support from an economic model, is sufficient to support that transportation costs convey localized market power to cement plants. Table 2 shows the results of two reduced-form regressions based on the available region-year observations. In each, the average price of domestic cement (in logs) is regressed on region and period fixed effects. The second also features as a regressor the average number of miles between the region's counties and the nearest customs office (located in San Francisco, Los Angeles, San Diego, and Nogales), interacted with the price of diesel fuel. This provides a simple measure of how insulated a region is from import competition. Its coefficient is identifiable in the presence of region and year fixed effects because the regressor varies across both regions and years.

The results of the first regression, shown in column 1, confirm that an impediment to arbitrage exists—otherwise, the estimated price discrepancies between regions would not exist.¹⁶ The results of the second regression identify transportation costs as the impediment. As shown, the coefficient on the measure of region isolation is positive and statistically significant, consistent with imports providing a stronger competitive constraint on prices when (i) the region is near the

¹⁴ The new reporting region includes Nevada, Idaho, Montana, and Utah. The Arizona-Nevada region includes a small plant located in New Mexico. We scale production and revenues by kiln capacity to minimize the influence of this plant.

¹⁵ Imports are in the form of clinker rather than ground cement because clinker is less prone to absorbing water from the air. Imported clinker is ground into cement after it clears customs.

¹⁶ Prices in northern California are estimated to be 3% higher than those in southern California, on average, and that the difference is statistically significant. Similarly, the prices in Arizona-Nevada are estimated to be 6.5% higher than southern California over 1983–1991 and 12.8% higher over 1992–2003.

import point; and (ii) the expense of transporting the imported cement is lower. Further, when this regressor is incorporated, the price differences across regions are smaller and not statistically different than zero, consistent with transportation costs being the primary impediment to arbitrage. We turn now to how one can leverage the structure of an economic model to generate more specific inferences about the role of transportation costs in the industry.

3. Economic model

□ **Supply.** We start with a standard oligopoly model of competition that incorporates price discrimination, capacity constraints, and a competitive fringe of foreign importers. Cement in the US Southwest is produced by a mix of multiplant and single-plant firms. We take as given the ownership structure and the location of plants.¹⁷ We assume that firms set different mill prices from each plant to each of the 90 counties in the US Southwest. This mill price does not include the transportation cost, which is paid by consumers to third parties. In principle, more sophisticated discrimination could be incorporated through the use of finer geographic partitions (e.g., census tracts or zip codes rather than counties). We focus on counties because two useful measures of construction activity are available at that level. We further assume that consumers do not conduct arbitrage across counties, consistent with the reduced-form evidence and our understanding that there is no operational secondary market for cement.

To formalize the model, let each firm f operate some subset \mathbb{J}_f of these plants and ship from any plant $j \in \mathbb{J}_f$ to any county n . Each firm chooses a vector of prices, $\mathbf{p}_f = (p_{jn}; j \in \mathbb{J}_f, n = 1, \dots, 90)$, to maximize its short run profits conditional on the prices chosen by all other firms. The profit function of firm f is

$$\pi_f(\mathbf{p}_f, \mathbf{p}_{-f}; \cdot) = \sum_{j \in \mathbb{J}_f} \sum_n p_{jn} q_{jn}(\mathbf{p}_n; \mathbf{x}_n, \boldsymbol{\beta}) - \sum_{j \in \mathbb{J}_f} \int_0^{Q_j(\mathbf{p}; \mathbf{x}, \boldsymbol{\beta})} c(Q; \mathbf{w}_j, \boldsymbol{\alpha}) dQ. \quad (1)$$

The quantity demanded from plant j by consumers in area n , denoted $q_{jn}(\cdot)$, is a function of all the prices in the county (\mathbf{p}_n). Total production at plant j is $Q_j(\cdot) = \sum_n q_{jn}(\cdot)$. The vectors \mathbf{x} and \mathbf{w} include demand and cost shifters, respectively, and $c(\cdot)$ is a marginal cost function. The vectors $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$ are the underlying structural parameters. We denote the equilibrium price vector \mathbf{p}^* . For 2003, this price vector is of length 1260 because there were 14 plants operating in the US Southwest during that year and 90 counties.

The marginal cost function allows for the incorporation of nonlinear production factors, such as capacity constraints, that are common in many industrial settings. We specify a marginal cost function that depends on the level of capacity utilization:

$$c(Q_j(\cdot); \mathbf{w}_j, \boldsymbol{\alpha}, \gamma, \nu, \mu) = \mathbf{w}'_j \boldsymbol{\alpha} + \gamma \mathbf{1} \left\{ \frac{Q_j(\cdot)}{CAP_j} > \nu \right\} \left(\frac{Q_j(\cdot)}{CAP_j} - \nu \right)^\mu, \quad (2)$$

where CAP_j is total plant capacity. This treatment of capacity constraints, an innovation of Ryan (2012), imbeds the intuition that production near capacity creates shadow costs due to foregone kiln maintenance. Thus, marginal costs increase in production once utilization exceeds ν , and the penalty due to production at capacity is $\gamma(1 - \nu)^\mu$. We find that it is difficult to estimate both γ and μ and we normalize the latter to 1.5.

¹⁷ Figure 1 provides the ownership structure and plant locations for 2003. We consider the treatment of plant locations as predetermined to be reasonable because we observe only two plant closures, one plant entry, and three substantive kiln upgrades in the US Southwest over 1983–2003, consistent with the substantial sunk costs of kiln construction documented in Ryan (2012).

We include the price of coal and a time trend as the linear cost shifters.¹⁸ Coal is the single largest input cost for many plants. Our specification abstracts from any heterogeneity in the fuel efficiency of cement plants for tractability. The plants in the US Southwest rely on dry kilns, which are fuel efficient relative to wet kilns.¹⁹ Any heterogeneity likely arises predominately from whether fuel-saving technologies are employed, such as preheaters and precalciners. The time trend is intended to capture changes in marginal costs that are unrelated to the procurement of coal. The baseline specification employs a time trend in logs. For robustness, we also estimate the model using a linear time trend and using a trend based on the total factor productivity of cement plants, as reported in the NBER-CES Manufacturing Industry Database (Bartlesman, Becker, and Gray, 2000). The latter approach may best capture the impact of unobserved factors affecting plant productivity, such as renegotiation of union contracts Dunne, Klimek, and Schmitz, (2009).²⁰

We assume that domestic plants compete against a competitive fringe of foreign importers. We denote this fringe as “plant” $J + 1$, and assign the fringe to four locations in the US Southwest based on the customs offices through which cement can enter (San Francisco, Los Angeles, San Diego, and Nogales). Consumers pay the door-to-door cost of transportation from these customs offices. We rule out spatial price discrimination on the part of the fringe, consistent with perfect competition among importers, and assume that the import price is set exogenously based on the marginal costs of the importers or other considerations. Thus, the supply specification is capable of generating the stylized fact developed above that foreign importers provide substantial quantities of portland cement to the US Southwest when demand is strong.

□ **Demand.** We use the nested logit demand system to model the behavior of consumers within each county. We specify the indirect utility that consumer i in county n receives from plant j as

$$u_{ij} = \beta^c + \beta^p p_{jn} + \beta^d DISTANCE_{jn} + \beta^i IMPORT_j + \epsilon_{ij}, \quad (3)$$

where $DISTANCE_{jn}$ is a measure of the distance between the plant and the centroid of county, $IMPORT_j$ is an indicator for imported cement, and ϵ_{ij} is a preference shock that is i.i.d. across consumers. We normalize the mean utility of the outside option to zero. The preference shock can be motivated as capturing the ability of the plant to meet the specific requirements of the consumer (e.g., related to the timeliness of production), the relationships between consumers and plants, and other considerations that are plausibly specific to the consumer.²¹ As we develop below, the distributional assumption that generates the nested logit demand system is a key modelling ingredient that both makes estimates feasible and has meaningful implications for the competitive outcomes.

The ratio β^d/β^p captures consumers’ willingness-to-pay for proximity to the plant. We interpret this as the cost of transportation. The two concepts are not exactly equivalent if distance

¹⁸ We have experimented with other cost shifters, such as electricity prices, the state-level wage rate for durable goods manufacturing, and a price index for crushed stone (a proxy for the price of limestone). Electricity prices are highly correlated with coal prices and their effects are difficult to identify separately. The estimated wage and stone coefficients tend to be near zero, which we suspect is due to measurement error in those variables: the durable goods manufacturing wage likely is a poor proxy for the wages of cement plant workers, and the crushed stone price imperfectly captures variation in the price of limestone.

¹⁹ We are aware of only two plants that operated wet kilns in the sample period. The Calaveras plant in San Andreas, California, operated three wet kilns before it shuttered in 1987. The Calaveras plant in Tehachapi, California, operated one wet kiln until it replaced the wet kiln in 1991 with a high-capacity dry kiln equipped with precalciner technology.

²⁰ The productivity measure is calculated at the national level and thus does not capture productivity specifically for the plants for the US Southwest, which is one reason we also rely on specifications with linear and log-linear time trends.

²¹ There is also an error-in-variables interpretation for the preference shock. The distance between the consumer and the plant is imperfectly captured by $DISTANCE_{jn}$ because consumers are dispersed throughout each county, whereas the regressor measures distance based on the county centroid. The consumer-specific deviations in distance would be orthogonal to the mean distance by construction.

affects consumer preferences for other reasons (e.g., reduced reliability). In our baseline specification, we measure distance using the driving distance interacted with a diesel price index, which should be reasonably correlated with the cost of truck transport. We also estimate the model using alternative measures based on straight-line distances and driving time, respectively, again interacted with a diesel price index.²²

We use a nesting structure that places the inside options (i.e., the domestic plants and foreign imports) in a different nest than the outside option. This allows the model to fit industries with inelastic aggregate demand and elastic plant-level demand—which is important here because materials such as steel, asphalt, and lumber are poor substitutes for portland cement in most construction projects but buyers plausibly view the output of different cement plants as close substitutes. We denote the nesting parameter as λ , following Cardell (1997). This parameter characterizes the degree to which valuations of the inside options are correlated across consumers. Valuations are perfectly correlated if $\lambda = 0$ and uncorrelated if $\lambda = 1$; the model collapses to a standard logit in the latter case.

The nested logit demand system conveys two critical advantages in estimation. First, it makes available analytical expressions for the sales of cement (i.e., for $q_{jn}(\mathbf{p}_n; \mathbf{x}_n, \boldsymbol{\beta})$) which ease substantially the computational burden of estimation. As we discuss below, the estimation routine we develop requires that equilibrium be computed numerically for every candidate parameter vector considered. Absent analytical solutions for demand, demand would have to be evaluated numerically for each candidate equilibrium price vector even as the equilibrium price vector is computed numerically for every candidate parameter vector. Second, it provides “smoothness” to the objective function by making demand a continuous function of prices. Were the model to treat all consumers in a given county as homogeneous, small changes in parameter values could lead the entire demand of the county to swing from one plant to another, creating discontinuities in the objective function.

The use of logit demand has economic implications as well. First, it imposes that competition is global, in the sense that each plant in the model ships at least some cement to each county—even if the distance would make transportation costs prohibitive in actuality. We do not find this troubling because our results imply that long-distance shipments are unusual (e.g., 90% of cement is shipped under 200 miles in the baseline specification). Second, and more meaningfully, the use of logit demand helps determine how localized market power and spatial price discrimination manifest in the model.²³ We examine later how model predictions are affected by scaling the variance of the consumer-specific preference shock (i.e., the “logit error”) away from the standard normalization of $\pi^2/6$. The model dictates that when the variance of the logit error is smaller, market power and price discrimination are greater and shipping distances are shorter because consumers place more weight on transportation costs relative to their idiosyncratic preferences. The robustness analysis helps inform how sensitive the results are to distributional assumptions.

The indirect utility equation does not incorporate unobserved plant attributes, such as quality, or unobserved demand shocks. The presence of such factors would create omitted variable bias given the estimation strategy we employ. We view the portland cement industry as a good match for the model in large part because unobservable considerations plausibly are much less relevant for consumer decisions than they are in industries with more standard differentiated products such as automobiles or breakfast cereals.

Finally, we normalize the market size or potential demand of each county based on a set of plausibly exogenous demand factors, following Berry, Levinsohn, and Pakes (1995) and Nevo (2001). The factors we use are the number of construction employees and the number of new residential building permits. The procedure imbeds the assumption that construction spending

²² For the foreign import option, we base $DISTANCE_{j+1}$ on the location of the nearest customs office to the county. The data on driving distance and driving time are obtained from Google maps.

²³ Given any set of parameters and prices, the variance of the logit error determines how plants’ market shares decrease with distance. Given any set of parameters, it determines how prices and margins decrease with distance.

is unaffected by cement prices, which seems reasonable because cement accounts for a small fraction of total construction expenditures (e.g., see Syverson, 2004).²⁴ The results indicate that potential demand is concentrated in a small number of counties. In 2003, the largest 20 counties account for 90% of potential demand, the largest 10% counties account for 65% of potential demand, and the largest two counties—Maricopa County and Los Angeles County—together account for nearly 25% of potential demand.

4. Estimation strategy

□ **The estimator.** We use a minimum distance estimator that compares the available endogenous data against the competitive outcomes implied by the model, aggregated to the level of the endogenous data. We denote the vector of endogenous data for period t as y_t . In our application, this vector contains production, consumption, and average prices for various geographic regions, as well as trade flows between some of those regions. We stack the exogenous data—the distances and the demand and cost shifters—for period t into a single matrix X_t . We stack the aggregated equilibrium predictions of the model into the vector $\tilde{y}_t(\theta; X_t)$, which is a function of the candidate parameter vector and the exogenous data.

The estimator minimizes the weighted sum of squared deviations between the endogenous data and the aggregated equilibrium predictions:

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T [y_t - \tilde{y}_t(\theta; X_t)]' C_T^{-1} [y_t - \tilde{y}_t(\theta; X_t)], \quad (4)$$

where Θ is some compact parameter space. Each element of the vector $y_t - \tilde{y}_t(\theta; X_t)$ defines a single nonlinear equation and C_T is a positive definite matrix that weights the equations.²⁵ The aggregated equilibrium predictions are obtained by solving for equilibrium at each candidate parameter vector and aggregating to the level of the data. We detail this process in Appendix B. The asymptotic properties of the estimator are complicated by the fact that the aggregate equilibrium predictions are functions of the implicit solution to the firms' first-order conditions. Nonetheless, the estimates are consistent and asymptotically normal provided that firms' first-order conditions have certain properties. We describe these properties and provide proofs in a companion article (Miller and Osborne, 2014).²⁶

We follow the two-step procedure of Hansen (1982) when we compute our estimates. In the first stage, we find a consistent estimate of the parameter values by using a diagonal weighting matrix in which each element is the sample variance of the relevant endogenous series.²⁷ In the second stage, we use a consistent estimate of the cross-equation variance matrix obtained in

²⁴ To perform the normalization, we regress regional portland cement consumption on the demand predictors (aggregated to the regional level), impute predicted consumption at the county level based on the estimated relationships, and then scale predicted consumption by a constant of proportionality to obtain potential demand. The regression of regional portland cement consumption on the demand predictors yields an R^2 of 0.9786. Additional predictors, such as land area, population, and percent change in gross domestic product, contribute little additional explanatory power. We use a constant of proportionality of 1.4, which is sufficient to ensure that potential demand exceeds observed consumption in each region-year observation.

²⁵ We use 11 aggregated equilibrium predictions for which empirical analogs are available: average mill prices (production weighted) charged by plants in northern California, in southern California, and in Arizona and Nevada; total production by plants in the same three regions; total consumption by consumers in northern California, in southern California, in Arizona, and in Nevada; and shipments from plants in California to consumers in northern California. The empirical analogs are available annually over 1983–2003 for the first 10 predictions (prices, production, and consumption) and over 1990–2003 for the 11th prediction (cross-region shipments). Thus, estimation exploits variation in 21 time-series observations on 10 nonlinear equations and 14 time-series observations on one nonlinear equation.

²⁶ Theoretical identification can fail if multiple candidate parameters produce equilibrium predictions that are identical once aggregated to the level of the available data, or if multiple equilibria exist for a single candidate parameter vector. We provide evidence that these problems do not arise in our application in a web Appendix.

²⁷ We found that weighting by sample variances, rather than using the identity matrix, results in estimates that better fit the price data in the first stage.

the first stage to weight the nonlinear equations. We use the methods of Hansen (1982), Newey and McFadden (1994), and Newey and West (1987) to calculate standard errors that are robust to heteroscedasticity and arbitrary correlations among the equations of each period, as well as first-order autocorrelation. We provide computational details in Appendix B.

□ **Identification.** We exploit cross-sectional and time-series variation in the data to identify the parameters. The relative importance of each depends on the parameter. For instance, the coal parameter is identified largely based on the correlation between the cement and coal prices, the latter of which vary more substantially over 1983–2003 than within the US Southwest. By contrast, the distance parameter is determined in part based on the magnitude of gaps between production and consumption within regions (e.g., that consumption exceeds production in Arizona-Nevada evidences trade flows), and in part based on how these gaps increase and decrease over time with diesel prices.

We plot selected empirical relationships that are important for identification in Figure 3. On the demand side, the price coefficient is primarily determined by the relationship between consumption and price. In panel A, we plot cement prices and the ratio of consumption to potential demand (“market coverage”) over the sample period. There is weak negative correlation, consistent with downward-sloping but inelastic aggregate demand. In panel B, we plot the gap between production and consumption (“excess production”) for each region. Excess production often is positive in southern California and negative elsewhere; the magnitude of the implied cross-region trade flows helps drive the distance coefficient. The implied trade flows are higher later in the sample when the diesel fuel is less expensive.

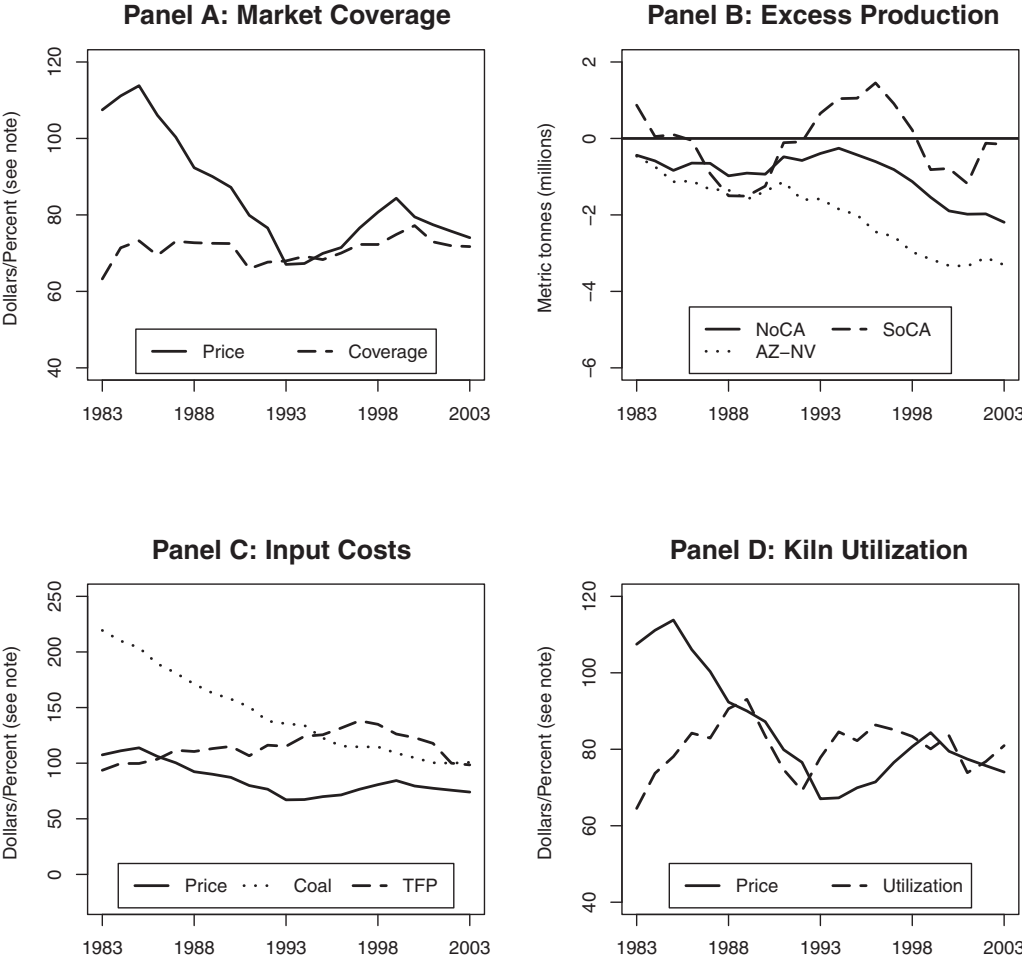
In panel C, we plot the coal price and the NBER’s measure of total factor productivity in the cement industry, together with cement prices. There is a strong positive correlation between coal and cement prices (both slope downward) and this relationship drives the coal coefficient. The relationship between productivity and prices is less clear. Finally, the utilization parameters are primarily determined by (i) the relative procyclicality of production and consumption, and (ii) the relationship between utilization and prices. We explore the second source of identification in panel D, which shows cement prices and industry-wide utilization over the sample period. The two metrics are negatively correlated over 1983–1987 and positively correlated over 1988–2003.

5. Results

□ **Parameter estimates and derived statistics.** Table 3 presents the results of estimation. Column (1), which we refer to as our baseline specification, features a distance measure of driving miles interacted with the diesel price index, and a marginal cost time trend in logs. The remaining columns use either an alternative distance measure (based on driving time or straight-line distance) or an alternative time trend (linear or based on total factor productivity). Shown in the table are parameter estimates and standard errors, derived economics statistics such as the transportation cost and selected demand elasticities, and the root mean squared error (RMSE) of estimation.

The price and distance coefficients are the two primary objects of interest on the demand side; both are negative and precisely estimated in each specification. The other demand parameters are also robust and precisely estimated: the negative coefficients on the import dummy are consistent with observed import prices that do not incorporate grinding costs, and the inclusive value coefficients easily reject (non-nested) logit demand. In the baseline specification, the price and distance coefficients together imply transportation costs of \$0.46 per tonne-mile given diesel prices at the 2000 level. The alternative specifications yield transportation costs ranging from \$0.43 to \$0.51 per tonne-mile. The aggregate elasticity in the baseline specification is -0.02 in the median year, consistent with the conventional wisdom that materials such as steel, asphalt, and lumber are poor substitutes for portland cement in most construction projects. The median

FIGURE 3
EMPIRICAL RELATIONSHIPS IN THE US SOUTHWEST



Notes: Panel A plots average cement prices and market coverage. Prices are in dollars per metric tonne and market coverage is defined as the ratio of consumption to potential demand (times 100). Panel B plots excess production in each region, which we define as the gap between production and consumption. Excess production is in millions of metric tonnes. Panel C plots average coal prices, the NBER total factor productivity index for cement, and the cement price. The coal and productivity time-series are converted to indexes that equal one in 2000. Panel D plots the average cement price and industry-wide utilization (times 100).

firm-level elasticity of -3.22 indicates that most buyers view the output of different cement plants as close substitutes.

Two publications confirm that our transportation cost estimates are reasonable. First, the 1974 edition of the *Minerals Yearbook* indicates transportation costs of \$0.35 per tonne-mile when adjusted to real 2000 dollars. Subsequent editions of the *Mineral Yearbook* do not estimate transportation costs. Second, the 20th edition of *Transportation in America* (2007) reports that revenues per tonne-mile for Class I general freight common carriers (i.e., basic truck transport) ranged from roughly \$0.29–\$0.35 over 1983–2003. Revenues for the transportation of cement, which requires specialized trucks, likely are somewhat higher. That our transportation cost estimates are somewhat higher than these numbers also could reflect that we are capturing consumers’

TABLE 3 Parameter Estimates and Derived Statistics

Specification:		(1)	(2)	(3)	(4)	(5)
Distance Measure		Drive Miles × Diesel	Drive Minutes × Diesel	Straight Miles × Diesel	Drive Miles × Diesel	Drive Miles × Diesel
Time Trend		Log	Log	Log	TFP	Linear
<i>Demand</i>						
Cement price	β^p	-0.049 (0.006)	-0.056 (0.006)	-0.056 (0.007)	-0.042 (0.004)	-0.047 (0.004)
Distance	β^d	-22.92 (0.973)	-24.13 (1.304)	-25.03 (1.655)	-21.58 (0.886)	-22.46 (1.07)
Import dummy	β^i	-1.547 (0.299)	-1.826 (0.348)	-2.506 (0.361)	-1.115 (0.179)	-1.376 (0.198)
Intercept	β^c	1.207 (0.045)	1.307 (0.065)	1.284 (0.070)	1.362 (0.073)	1.398 (0.087)
Inclusive value	λ	0.025 (0.009)	0.029 (0.005)	0.021 (0.007)	0.049 (0.009)	0.049 (0.010)
<i>Marginal costs</i>						
Coal price	α^c	0.924 (0.052)	0.948 (0.046)	0.950 (0.046)	0.954 (0.041)	0.878 (0.045)
Time trend	α^t	-0.874 (0.861)	-1.011 (0.331)	-1.086 (0.777)	-14.86 (4.294)	-0.301 (0.127)
Capacity cost	γ	1904 (2057)	670.8 (543.2)	578.7 (227.3)	844.8 (476.3)	784.0 (199.3)
Cap. threshold	ν	0.958 (0.045)	0.906 (0.043)	0.900 (0.034)	0.896 (0.044)	0.904 (0.025)
<i>Derived statistics</i>						
Transport cost		46.5	43.0	44.3	51.1	47.3
Agg. elasticity		-0.02	-0.03	-0.02	-0.04	-0.04
Firm elasticity		-3.22	-3.47	-3.75	-2.84	-3.09
<i>Model fits</i>						
RMSE (in sample)		867	934	998	905	934
RMSE (out sample)		118	104	146	118	125

Notes: Estimation exploits variation in regional consumption, production, and average prices over 1983–2003, as well as variation in shipments from California to northern California over 1990–2003. The prices of cement and coal are in dollars per metric tonne. Miles and minutes are in thousands. The diesel price is an index that equals one in 2000. The marginal cost parameter ϕ is normalized to 1.5. Standard errors are robust to heteroscedasticity, first-order autocorrelation, and contemporaneous correlations between moments.

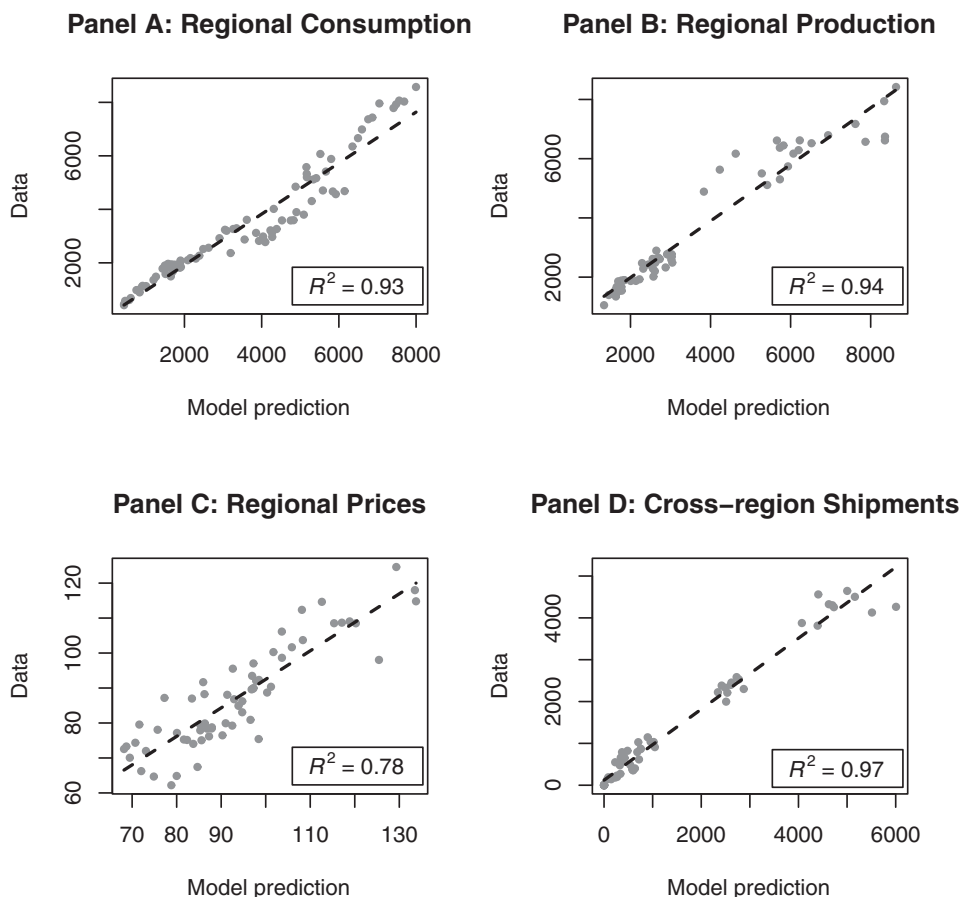
willingness-to-pay for proximity to the plant, which could be affected by other considerations related to distance (e.g., reduced reliability).

On the cost side, the parameter estimates imply marginal costs of \$50.71 per tonne for the average cement plant in 2003. Of this, \$36.28 is attributable to the constant portion of marginal costs and the remaining \$14.43 is attributable to high utilization rates. Integrating the marginal cost function over the levels of production that arise in numerical equilibrium yields an average variable cost of \$55 million. The bulk of these variable costs—71.6%—are due to input costs rather than due to high utilization. Taking the statistics further, we calculate that the average plant has variable revenues of \$87 million in 2003 and that the average gross margin (variable profits over variable revenues) is 36.7%. As argued in Ryan (2012), margins of this magnitude may be needed to rationalize entry given the sunk costs associated with plant construction.

We select column (1) as the baseline specification because it has the smallest in-sample RMSE. It also has among the lowest out-of-sample RMSE, based its ability to fit data on cross-region shipments that were withheld from estimation. Figure 4 explores the fit of the baseline specification in greater detail. We plot observed regional consumption against predicted regional consumption (panel A), observed regional production against predicted regional production (panel B), and observed regional prices against regional predicted prices (panel C). Univariate regressions

FIGURE 4

ESTIMATION FITS FOR REGIONAL METRICS



Notes: Consumption, production, and cross-region shipments are in millions of metric tonnes. Prices are constructed as a weighted average of plants in the region, and are reported as dollars per metric tonne. The lines of best fit and the reported R^2 values are based on univariate OLS regressions.

of the data on the predictions indicate that the model explains 93% of the variation in regional consumption, 94% of the variation in regional production, and 78% of the variation in regional prices. We also plot observations on cross-region shipments against the corresponding model predictions (panel D). We use 14 of these observations in the estimation routine—the shipments from plants in California to consumers in northern California over 1990–2003—but the remaining 82 data points are withheld from the estimation procedure and do not influence the estimated parameters. Even so, the model explains 97% of the variation in these data. That the economic model, evaluated at the parameter estimates, predicts market outcomes similar to those in the data provides confidence that it is a good fit for the portland cement industry.

□ **Market power and price discrimination.** The estimation results imply that transportation costs facilitate the exercise of localized market power in some counties but not others.²⁸ Table 4

²⁸ An interesting implication of the specification—one that we have not fully explored—is that transportation costs and spatial differentiation fluctuate with diesel prices. The extent to which carbon or gasoline taxes would have unintended consequences on the intensity of competition in industries such as portland cement remains an open question.

TABLE 4 Leading Plants in Maricopa County and Los Angeles County in 2003

Plant Owner	Plant Location	Distance	Mill Price	Trans. Cost	Margin	Share
<i>Maricopa County (Phoenix)</i>						
Phoenix Cement	Clarkdale, AZ	128	\$77	\$60	0.52	53%
California Cement	Rillito, AZ	122	\$84	\$57	0.56	43%
Cemex	Victorville, CA	346	\$57	\$162	0.37	1%
<i>Los Angeles County</i>						
Cemex	Victorville, CA	96	\$61	\$45	0.41	18%
National Cement	Encino, CA	22	\$88	\$10	0.59	26%
California Cement	Mojave, CA	95	\$63	\$44	0.43	16%

Notes: Based on estimation results. Distance is the miles between the plant and the county centroid. Mill price and transportation cost are per metric tonne. Mill price is computed based on the estimation results. Margin is based on the mill price and the constant portion of marginal costs (it ignores utilization costs). Share is the proportion of domestic cement consumed in the county that is produced by the plant.

contrasts Maricopa County (Phoenix) and Los Angeles County in 2003, based on the equilibrium computed with the baseline parameter estimates. The results indicate that 96% of the cement consumed in Maricopa County is supplied by two plants. These plants, operated by Phoenix Cement and California Cement, are located about 120 miles north and south of the county, respectively, and charge mill prices of \$77 and \$84. Consumers must spend around \$60 on transportation. Although the mill prices of the southern California plants to Maricopa County are lower, the transportation costs are much higher (e.g., the mill price of the Cemex plant is \$57 but transportation is \$162). This enables the plants in Arizona to support mill prices to Maricopa County well above the cost of production.²⁹ By contrast, the leading suppliers of Los Angeles County tend to be less differentiated spatially and have less localized market power.

The estimation results also imply that the geographic configuration of the US Southwest permits some plants to discriminate among consumers. In Figure 5, we plot the “total cost of purchase” (i.e., the mill price plus the transportation cost) for counties within 400 miles of the Cemex plant in southern California and the Phoenix Cement plant in Arizona. In the absence of price discrimination, one would expect the total cost of purchase to increase linearly in distance. This is precisely what the model implies for the Cemex plant. The line of best fit is produced from a regression of total purchase cost on distance, using *only* counties farther than 200 miles from the plant. Yet it predicts total purchase costs for closer plants equally well. Further, because the slope of the line is 0.46, total purchase costs increase at the same rate as transportation costs (which we estimate at \$0.46 per tonne-mile). By contrast, the model implies that the Phoenix Cement plant discriminates among its consumers. The total costs of purchase for consumers in counties within 200 miles exceed the line of best fit based on counties farther than 200 miles from the plant by \$30.89 on average; this is due to higher mill prices for consumers in nearby counties.³⁰ That the slope of the best fit line is again 0.46 indicates that spatial price discrimination is a local phenomenon—the plant does not discriminate between “distant” and “very distant” consumers.

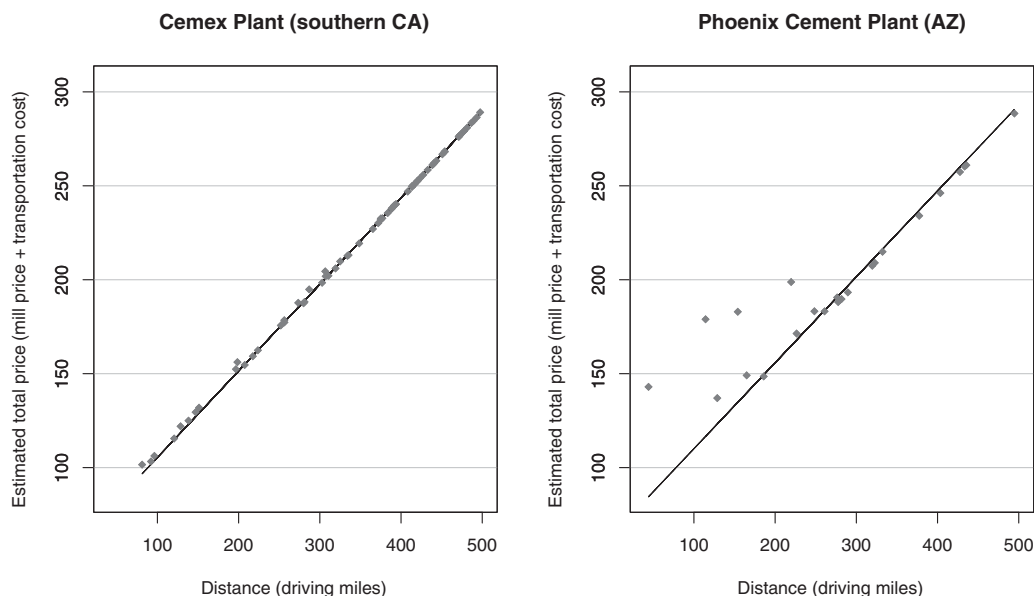
The key difference between the Cemex plant in southern California and the Phoenix Cement plant in Arizona is location: the presence of nearby competitors constrains price discrimination on the part of the Cemex plant, whereas the Phoenix Cement plant is more differentiated spatially. To provide a fuller treatment, we plot the plant-county specific margins implied by the model

²⁹ The margin shown is based on the mill price and the constant portion of marginal costs, and approximates a variable cost margin. In the notation established, $m = (p_{jn} - \mathbf{w}'\hat{\boldsymbol{\alpha}}) / p_{jn}$. Incorporating utilization costs would yield the Lerner index. We find that plants with localized market power typically operate at higher utilization rates, presumably due to the economic profits available.

³⁰ The gap between equilibrium prices and the line of best fit can be interpreted as a back-of-the-envelope calculation of how much localized market power increases prices, albeit one that does not account for competitive interactions. If Phoenix Cement were to change its price schedule then, presumably, so would its competitors. We account for these interactions in a counterfactual policy experiment presented in Section 6.

FIGURE 5

PRICE DISCRIMINATION AT TWO PLANTS IN 2003



Notes: The vertical axis is the total cost of purchase, that is, the mill price plus the transportation cost incurred by the consumer. The mill price is computed based on the estimation results. The horizontal axis is the distance in miles between the plant and the county centroid. Each dot represents the total cost of purchase for a plant-county pair. The line of best fit is from a regression of total cost of purchase on distance, using pairs with distance greater than 200 miles.

in Figure 6. The most pronounced discrimination arises for plants that are relatively insulated from domestic and import competition—for instance, at the Phoenix Cement and California Cement plants in Arizona, the Centex plant in Nevada, and the Lehigh Cement plant in northern California.³¹ Price discrimination is more subdued at the plants in southern California and near San Francisco.

□ **Role of heterogeneous preferences.** The nested logit framework we employ incorporates idiosyncratic consumer tastes for cement from each plant. The degree of this heterogeneity—equivalently, the variance of the “logit error”—is not separably identifiable in estimation and determines how localized market power and spatial price discrimination manifest in the model given the results of estimation.³² Here, we examine how selected model predictions are affected by varying the degree of consumer heterogeneity, taking as given the results of estimation. In particular, we recompute equilibrium alternately scaling the idiosyncratic portion of indirect utility by 0.75, 1.25, and 1.50. This is equivalent to normalizing the variance of the logit error to be $(9/16) * (\pi^2/6)$, $(25/16) * (\pi^2/6)$, and $(36/16) * (\pi^2/6)$, respectively. The analysis helps inform how sensitive the results are to distributional assumptions.

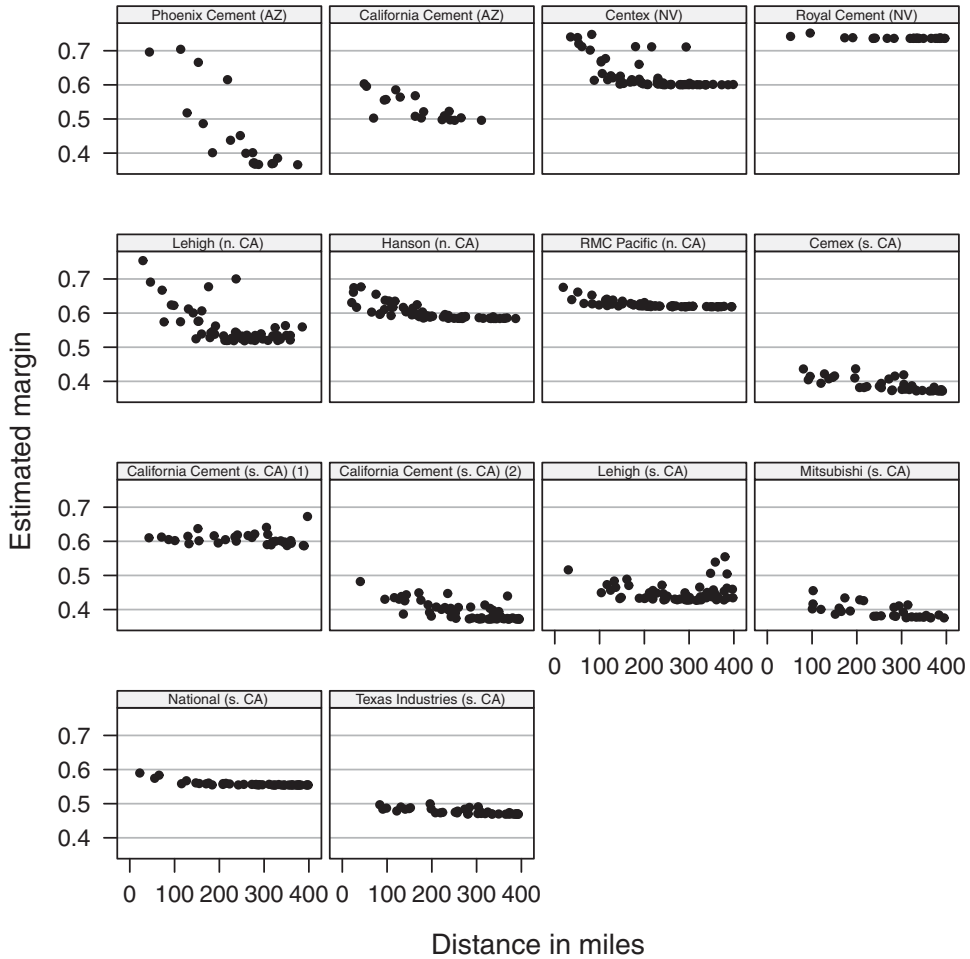
Table 5 shows the implications for the computed distributions of mill prices and shipping distances in 2003. Under the baseline variance assumption of $\pi^2/6$, which is standard in the discrete-choice literature, the mean mill price is \$80 per metric tonne, the mean shipping distance

³¹ The exception is the low-capacity Royal Cement plant in southern Nevada. The plant ships more than 95% of its output to consumers in Clark County (i.e., Las Vegas), and it incurs substantial utilization costs that prevent the plant from profitably lowering its price to consumers in more distant counties.

³² Notably, though, the transportation cost is unaffected by the variance of the logit error because it is constructed as the ratio of two coefficients—the distance parameter and the price parameter. Scaling the variance of logit error affects both of those coefficients equally, so the ratio is unchanged.

FIGURE 6

PRICE DISCRIMINATION AND MARGINS IN 2003, BY PLANT



Notes: The vertical axis is a margin based on mill prices and the constant portion of marginal costs (i.e., it excludes utilization costs). Margins are computed based on the estimation results. The horizontal axis is the distance in miles between the plant and the county centroid. Each dot represents the margin for a plant-county pair.

is 122 miles, and 90% of cement is shipped less than 208 miles. Relative to this baseline, mill prices are higher and shipping distances are shorter with a smaller variance of $(9/16) * (\pi^2/6)$. The reverse happens when the variance of the logit error is scaled up. Directionally, these results are as expected. We find it interesting that shipping distances appear somewhat more responsive than mill prices to the variance of the logit error, at least over the range considered. Overall, the results indicate many predictions of the model are materially similar for a range of assumptions pertaining to importance of consumer heterogeneity.

6. Counterfactual experiments

□ **Spatial discrimination and consumer surplus.** We conduct a counterfactual policy experiment to evaluate the implications of spatial price discrimination in the portland cement industry. In particular, we compute equilibrium under the restriction that each plant sets the same mill price to each county, taking as given the baseline parameter estimates and the topology of the

TABLE 5 Alternative Variances for Preference Shocks

Model Prediction	Mean	Percentiles				
		10th	25th	50th	75th	90th
<i>Baseline—Variance of $\pi^2/6$</i>						
Mill prices	\$80	\$61	\$64	\$77	\$92	\$104
Driving miles	122	40	83	121	150	208
<i>Variance of $(9/16) * (\pi^2/6)$</i>						
Mill prices	\$81	\$61	\$64	\$77	\$92	\$106
Driving miles	116	32	81	116	142	199
<i>Variance of $(25/16) * (\pi^2/6)$</i>						
Mill prices	\$79	\$61	\$63	\$77	\$91	\$104
Driving miles	129	42	84	123	162	220
<i>Variance of $(36/16) * (\pi^2/6)$</i>						
Mill prices	\$79	\$60	\$63	\$76	\$90	\$100
Driving miles	136	43	87	123	174	244

Notes: Model predictions for 2003 generated by computing equilibrium given the parameter estimates and different assumptions about the variance of the consumer-specific preference shock (i.e., the “logit error”). Mill prices are per metric tonne.

industry in the year 2003. The consumer surplus implications of spatial price discrimination have long been recognized as theoretically ambiguous (e.g., Gronberg and Meyer, 1982; Katz, 1984; Hobbs, 1986; Anderson, de Palma, and Thisse, 1989) and, to our knowledge, we provide the first empirical evidence on the matter.

Figure 7 characterizes the consumer surplus implications of disallowing spatial price discrimination. Counties that are shaded in dark gray or black are harmed by the ban whereas counties shaded in light gray or white are benefited. The net effect, aggregating across all counties, is a \$28.8 million gain in consumer surplus. We calculate a 95% confident interval of (\$7.2 million, \$35.2 million) by sampling from the estimated distribution of parameters.³³ This can be compared to the volume of commerce in the US Southwest of \$1.3 billion in 2003.³⁴ The effects of the ban vary widely across counties, with consumers near cement plants benefitting and more distant consumers being harmed. As nearby consumers tend to be inframarginal whereas distant consumers tend to be marginal, this results follows the economics of the model—price discrimination enables plants to extract surplus accruing to inframarginal consumers without sacrificing sales to marginal consumers.

Table 6 shows more specific results for Maricopa County (Phoenix) and two counties immediately to the north and south (Yavapai County and Pima County, respectively). The starkest effects of the ban arise here. As shown, the mill price of the Phoenix Cement plant to consumers in Yavapai falls from \$122 per metric tonne to \$81, and the mill price of the California Cement plant to consumers in Pima falls from \$92 to \$87. Due to these price effects, disallowing price discrimination creates \$5.1 million and \$1.2 million in consumer surplus in these counties, respectively. By contrast, the prices that these plants charge to the consumers in Maricopa increase due to the price discrimination ban, leading to \$6.8 million in lost consumer surplus.

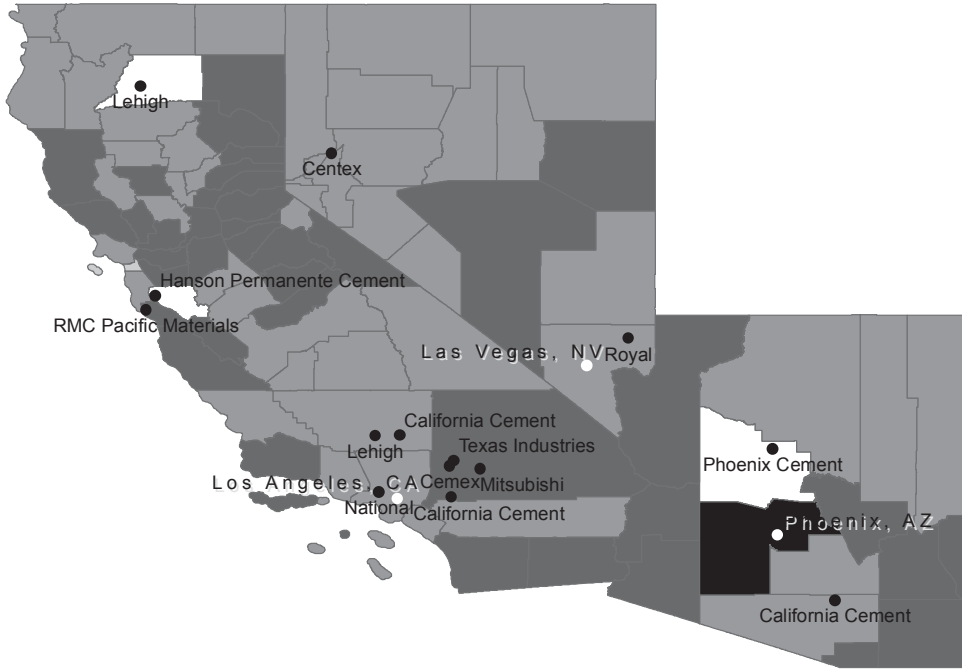
□ **Merger simulation.** Antitrust authorities sometimes have access to incomplete data on market outcomes due to limitations of the Hart-Scott-Rodino Act. The full complement of firm-level data needed to estimate the spatial models of Thomadsen (2005), Davis (2006), McManus (2007), and Houde (2012) rarely is available. The flexible data requirements of our estimator can be valuable in such settings. To illustrate, we evaluate a hypothetical merger between Calmat and

³³ Based on 1000 random draws.

³⁴ Volume of commerce is calculated as price times quantity for all sales by plants in the US Southwest.

FIGURE 7

EFFECTS OF DISALLOWING PRICE DISCRIMINATION ON CONSUMER SURPLUS



Change in consumer surplus (thousands)




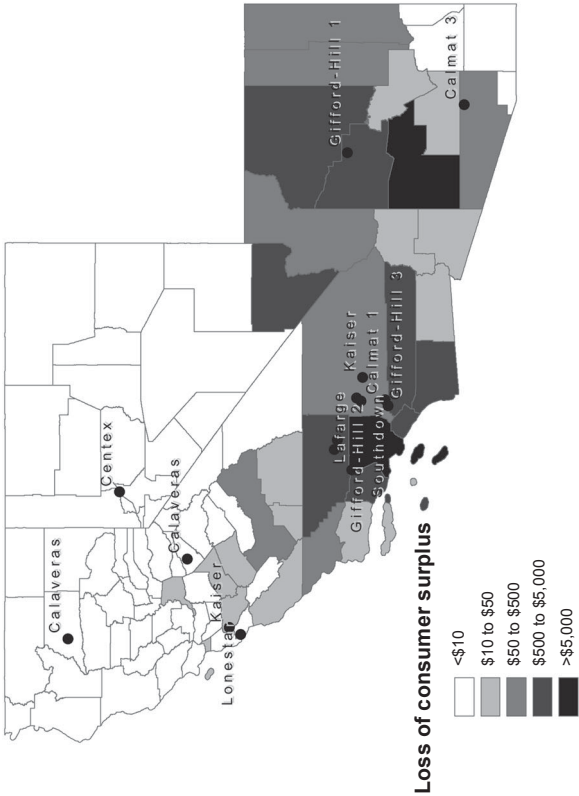
TABLE 6 Effects of Disallowing Price Discrimination on Prices in Selected Counties

Plant Owner	Plant Location	Distance	Trans. Cost	Price PreBan	Price PostBan
<i>Maricopa County (Phoenix)</i>					
Phoenix Cement	Clarkdale, AZ	129	60	77	81
California Cement	Rillito, AZ	123	57	84	87
<i>Yavapai County (Clarkdale)</i>					
Phoenix Cement	Clarkdale, AZ	44	20	122	81
California Cement	Rillito, AZ	205	95	76	87
<i>Pima County (Rillito)</i>					
Phoenix Cement	Clarkdale, AZ	278	129	59	81
California Cement	Rillito, AZ	68	32	92	87

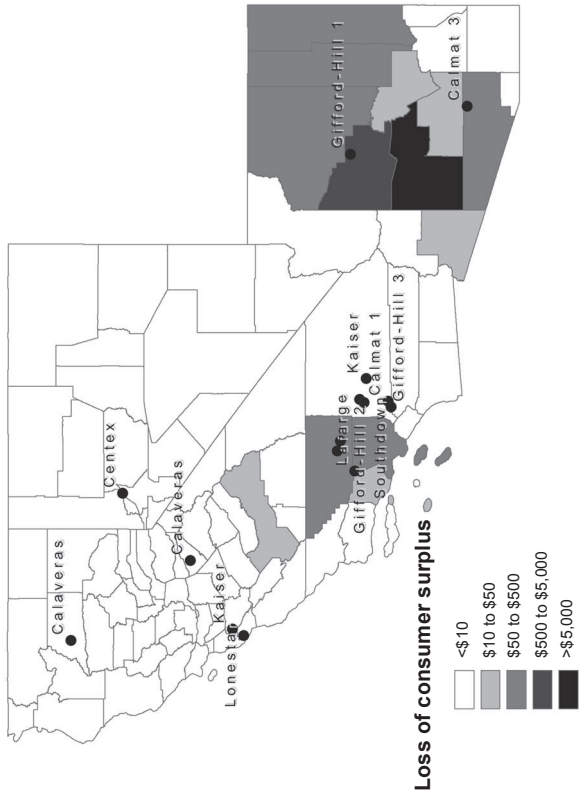
Notes: Results of the counterfactual experiment. Distance is the miles between the plant and the county centroid. Transportation cost is per metric tonne. Preprice is the mill price in the discriminatory regime and postprice is the mill price in the nondiscriminatory regime; both are per metric tonne.

FIGURE 8
LOSS OF CONSUMER SURPLUS DUE TO A HYPOTHETICAL MERGER

Map A: No Divestiture



Map B: Optimal Divestiture



Gifford-Hill in 1986. These two firms operated four of the eight plants in southern California and both of the plants in Arizona.³⁵

Simulation results based on the baseline parameter estimates indicate that the merger leads to prices at the Calmat and Gifford-Hill plants that are 4.9% higher in southern California and 9.8% higher in Arizona, on average. This induces consumer switching; and consumers that do switch split evenly between other domestic plants (46%) and foreign importers (54%). Prices at other domestic plants increase by only 0.8% on average. Total consumer surplus falls by \$31 million, relative to a total volume of commerce in southern California and Arizona of \$801 million.³⁶ Figure 8 maps the distribution of consumer harm that arises from the hypothetical merger of Calmat and Gifford-Hill. Panel A focuses on effects of the merger absent any divestitures. Harm is concentrated in the counties surrounding Los Angeles and Phoenix. Panel B plots harm under the most powerful single-plant divestiture, that of Gifford-Hill's Oro Grande plant ("Gifford-Hill 2" in the figure). This divestiture eliminates 62% of total consumer harm. This relief occurs mainly in southern California; the divestiture does little to reduce harm in Arizona. Additional counterfactual exercises indicate that another divestiture is needed to mitigate this harm as well.

We contrast these simulation results to those that one would obtain by supposing that (i) competition is Cournot in each region, (ii) demand has constant elasticity, (iii) plants share a marginal cost, and (iv) there are no foreign imports. This is a standard framework for analyzing the cement industry—for instance, aside from the marginal cost assumption, it mimics the modelling framework of Ryan (2012). Postmerger prices are

$$p^{post} = \frac{N(N-1)e - (N-1)}{N(N-1)e - N} p^{pre}, \quad (5)$$

where N is the number of firms and e is the aggregate elasticity of demand.³⁷ The merger has the effect of reducing the number of firms from six to five in southern California and from two to one in Arizona. Using the aggregate elasticity estimate of Ryan, 2012 obtains predicted price increases of 1% in southern California and 25% in Arizona. This diverges starkly from the results of our model, which more seriously treats the spatial elements of competition.

7. Conclusion

■ We estimate a structural model of the cement industry that incorporates spatial differentiation and spatial price discrimination, focusing on the US Southwest over 1983–2003. In doing so, we develop an estimation strategy for dealing with data on market outcomes that are substantially aggregated. In the broadest sense, our work extends on literature of the “new empirical industrial organization,” which focuses largely on the structural estimation of economic models. One area of particular interest in the literature has been the estimation of product differentiation models, as in Berry et al. (1995) and Nevo (2001). Geographic considerations, with some exceptions, have received less attention. Our research develops methods for industries in which the primary source of differentiation is spatial.

Our hope is that the estimation strategy we introduce extends the reach of empirical researchers. In a counterfactual policy experiment, we find that disallowing spatial price

³⁵ In the Working Paper version of this article, we show that merger simulation based on market delineation and constant elasticity demand yields substantially different predictions.

³⁶ We follow McFadden (1981) and Small and Rosen (1981) in calculating consumer surplus. Volume of commerce is calculated as price times quantity for all sales by plants in southern California and Arizona.

³⁷ In obtaining this expression, it is useful to keep in mind the relationship between firm elasticities and the aggregate elasticity, that is, that $e_j = Ne$ where e_j is the firm elasticity. Then, manipulation of the Lerner index yields a familiar expression for post-merger prices:

$$p^{post} = \frac{(N-1)e}{(N-1)e-1} c \quad \text{where} \quad c = \frac{Ne-1}{Ne} p^{pre}.$$

discrimination in the portland cement industry would increase consumer surplus by a modest \$30 million, relative to a volume of commerce of \$1.3 billion. Other applications have equal promise. Researchers could study the relationship between transportation costs and the intensity of competition or the proper construction of antitrust markets. In addition, although our application is static, the estimator could be used to define payoffs in strategic dynamic games. Such extensions could examine an array of interesting topics, including entry deterrence, optimal location choice, and the effects of various government policies (e.g., carbon taxes or import duties) on welfare and the long-run location of production. Indeed, the first of these proposed extensions now is the subject of Chicu (2012), which uses a model similar to ours to define the payoffs of dynamic investment games to study entry deterrence in the cement industry.

Appendix A

Appendix A provides details on the data.

We make various adjustments to the data in order to improve consistency over time and across different sources. We discuss some of these adjustments here, in an attempt to build transparency and aid replication. To start, we note that the *California Letter* is based on a monthly survey rather than on the annual USGS census, which creates minor discrepancies. We normalize the *California Letter* data prior to estimation so that total shipments equal total production in each year. The 96 cross-region data points include:

-
- | | |
|---------------------------------|---------------------------------|
| ● CA to N. CA over 1990–2003 | ● S. CA to N. CA over 1990–1999 |
| ● CA to S. CA over 2000–2003 | ● S. CA to S. CA over 1990–1999 |
| ● CA to AZ over 1990–2003 | ● S. CA to AZ over 1990–1999 |
| ● CA to NV over 2000–2003 | ● S. CA to NV over 1990–1999 |
| ● N. CA to N. CA over 1990–1999 | ● N. CA to AZ over 1990–1999. |
-

The (single) Arizona-Nevada region includes Nevada data only over 1983–1991. Starting in 1992, the USGS combined Nevada with Idaho, Montana, and Utah to form a new reporting region. We tailor the estimator accordingly. Additionally, this region also includes information from a small plant located in New Mexico. We scale the USGS production data downward, proportional to plant capacity, to remove for the influence of this plant. As the two plants in Arizona account for 89% of kiln capacity in Arizona and New Mexico in 2003, we scale production by 0.89. We do not adjust prices.

The portland cement plant in Riverside closed its kiln permanently in 1988 but continued operating its grinding mill with purchased clinker. We include the plant in the analysis over 1983–1987, and we adjust the USGS production data to remove the influence of the plant over 1988–2003 by scaling the data downward, proportional to plant grinding capacities. As the Riverside plant accounts for 7% of grinding capacity in southern California in 1988, we scale the production data for that region by 0.93.

We exclude one plant in Riverside that produces white portland cement. White cement takes the color of dyes and is used for decorative structures. Production requires kiln temperatures that are roughly 50°C hotter than would be needed for the production of grey cement. The resulting cost differential makes white cement a poor substitute for grey cement.

The Portland Cement Association reports that the California Cement Company idled one of two kilns at its Colton plant over 1992–1993 and three of four kilns at its Rillito plant over 1992–1995, and that the Calaveras Cement Company idled all kilns at the San Andreas plant following the plant's acquisition from Genstar Cement in 1986. We adjust plant capacity accordingly.

We multiply kiln capacity by 1.05 to approximate cement capacity, consistent with the industry practice of mixing clinker with a small amount of gypsum (typically 3% to 7%) in the grinding mills.

The data on coal prices from the Energy Information Agency are available at the state level starting in 1990. Only national-level data are available in earlier years. We impute state-level data over 1983–1989 by (i) calculating the average discrepancy between each state's price and the national price over 1990–2000, and (ii) adjusting the national-level data upward or downward, in line with the relevant average discrepancy.

Appendix B

Appendix B provides details on the estimation strategy.

□ **Obtaining aggregate equilibrium predictions.** In this Appendix, we detail how aggregate equilibrium conditions can be obtained given some candidate parameter vector. The key ingredient is the equilibrium price vector, which can be computed from the first-order conditions of the firms' profit maximization problem. There are $J \times N$ first-order

conditions, reflecting the modelling assumption that each plant can discriminate between the consumers of different areas. For notational convenience, we define the block-diagonal matrix $\Omega(\mathbf{p}; \mathbf{X}_t, \boldsymbol{\theta})$ as the combination of $n = 1, \dots, N$ submatrices, each of dimension $J \times J$. The elements of the submatrices are defined as follows:

$$\Omega_{jk}^n(\mathbf{p}_n; \mathbf{X}_t, \boldsymbol{\theta}) = \begin{cases} \frac{\partial q_{jn}(\mathbf{p}_n; \mathbf{X}_t, \boldsymbol{\theta})}{\partial p_{kn}} & \text{if } j \text{ and } k \text{ have the same owner} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B1})$$

The elements of each submatrix characterize substitution patterns within area \mathbb{C}_n , and Ω has a block-diagonal structure because $q_{jn}(\mathbf{p}_n; \mathbf{X}_t, \boldsymbol{\theta})$ is free of \mathbf{p}_{-n} . Thus, the construction of Ω builds on the premises that (i) consumers in each area \mathbb{C}_n select among all J plants, and (ii) demand in area \mathbb{C}_n is unaffected by mill prices in area \mathbb{C}_m for $n \neq m$. With this notation in hand, the first-order conditions take the form

$$\mathbf{f}(\mathbf{p}; \mathbf{X}_t, \boldsymbol{\theta}) \equiv \mathbf{p} - \mathbf{c}(\mathbf{Q}(\mathbf{p}; \mathbf{X}_t, \boldsymbol{\theta}); \mathbf{X}_t, \boldsymbol{\theta}) + \Omega^{-1}(\mathbf{p}; \mathbf{X}_t, \boldsymbol{\theta})\mathbf{q}(\mathbf{p}; \mathbf{X}_t, \boldsymbol{\theta}) = \mathbf{0}. \quad (\text{B2})$$

A vector of prices that solves this system of equations is a Bertrand-Nash equilibrium. In most applications, however, analytic solutions are unobtainable. Rather, one must solve equation (B2) numerically using a nonlinear equation solver to produce a vector of *computed equilibrium prices*, which we denote $\tilde{\mathbf{p}}^*(\boldsymbol{\theta}; \mathbf{X}_t)$. Specifically, the nonlinear equation solver selects the vector $\tilde{\mathbf{p}}^*$ to satisfy

$$\frac{1}{JN} \|\mathbf{f}(\tilde{\mathbf{p}}^*; \mathbf{X}_t, \boldsymbol{\theta})\| < \delta, \quad (\text{B3})$$

where δ is a user specified tolerance. A tolerance of 1E-13 performs well in numerical experiments based on our application. Numerical error can propagate into the objective function when the tolerance is substantially looser (e.g., 1E-7), which slows overall estimation time and can produce poor estimates. These thresholds are specific to our application because tolerance is not unit free and must be evaluated relative to the price level.

Once the equilibrium price vector is obtained, it can be manipulated into the aggregated equilibrium predictions. To formalize this process, we define a function $S : \mathbb{R}^{JN} \rightarrow \mathbb{R}^L$ that maps from the equilibrium price vector to the aggregate equilibrium predictions; L is the number of predictions that must be calculated (i.e., the length of \mathbf{y}_t). The aggregate equilibrium predictions that enter the objective function are given by $\tilde{\mathbf{y}}_t(\boldsymbol{\theta}; \mathbf{X}_t) = S(\tilde{\mathbf{p}}^*(\boldsymbol{\theta}; \mathbf{X}_t))$. We assume that $S(\cdot)$ is continuously differentiable, which holds for applications based on averaged or summed endogenous data.

Example: In our application, the estimator makes use of 11 nonlinear equations in most time periods. Three of these relate to the average mill prices (production weighted) charged by plants in specific geographic regions. Thus, denoting the set of plants in region r as A_r , these aggregate equilibrium predictions can be calculated as

$$\tilde{P}_{rt}(\boldsymbol{\theta}, \mathbf{X}_t) = \sum_{j \in A_r} \sum_n \frac{q_{jn}(\tilde{\mathbf{p}}_n^*; \mathbf{X}_t, \boldsymbol{\theta})}{\sum_{j \in A_r} \sum_n q_{jn}(\tilde{\mathbf{p}}_n^*; \mathbf{X}_t, \boldsymbol{\theta})} \tilde{p}_{jn}^*$$

The aggregate equilibrium predictions for production, consumption, and cross-region shipments can be written analogously. These predictions can be stacked into the vector $\tilde{\mathbf{y}}_t(\boldsymbol{\theta}; \mathbf{X}_t)$ and compared to the data.

The estimator has a nested structure in which a numerical optimizer finds the parameter vector that minimizes the objective function and a nonlinear equation solver computes equilibrium prices conditional on the parameter vector. This structure complicates implementation because the dimensionality of the equilibrium price vector that must be computed can be quite large. In our application, there are 90 consumer-areas and 14 plants (in a typical year), resulting in a price vector with 1260 elements. In many standard numerical packages, solving for such a large price vector is computationally intensive. One way to reduce computational complexity is to assume that the firm's marginal cost function is constant. Under this assumption, one can solve for the equilibrium prices in each consumer-area individually, substantially saving computational time.

In many applications, marginal costs are unlikely to be constant and the prices that characterize equilibrium in different consumer areas are not independent. If marginal costs increase with production (e.g., due to capacity constraints), then lowering price in one consumer-area will increase overall quantity sold by a plant, raising its cost, and hence, its equilibrium price, in other areas. In general, one may need to solve for the entire vector of prices jointly. We use a large-scale nonlinear equation solver developed in La Cruz, Martinez, and Raydan (2006) to compute equilibrium in our application. The equation solver employs a quasi-Newton method and exploits simple derivative-free approximations to the Jacobian matrix; it converges more quickly than other algorithms and does not sacrifice precision. This algorithm is available as part of the BB package in the statistical programming language R. Our application uses a Fortran version of the nonlinear equation solver, which significantly increases computational speed.³⁸

³⁸ The function that implements the solver is titled *dfsane*. Our experience is that Fortran reduces the computational time of the inner loop by a factor of 30 or more, relative to the *dfsane* function in R. The numerical computation of equilibrium takes between 2 and 12 seconds for most candidate parameter vectors when run on a 2.40GHz dual core processor with 4.00GB of RAM. We have been able to achieve faster speeds using a computational server.

□ **Estimation.** We minimize the objective function in two steps. First, to get in the vicinity of the minimum of the objective function, we use a simplex algorithm with a loose relative convergence tolerance of 0.01 (Nelder and Mead, 1965). Then, starting at the output of the simplex optimizer, we use the Levenberg-Marquardt algorithm (Levenberg, 1944; Marquardt, 1963) to find the optimum of the objective function, using a tight tolerance of the square root of machine epsilon (roughly 1E-8). The Levenberg-Marquardt algorithm interpolates between the Gauss-Newton algorithm and the method of gradient descent, and we find that it outperforms quasi-Newton methods such as BFGS. We implement the minimization procedure using the `nls.lm` function in R, which is downloadable as part of the `minpack.lm` package.

We use observed prices to form the basis of the initial vector in the inner loop computations, which limits the distance that the nonlinear equation solver must walk to compute numerical equilibrium. The nonlinear equation solver `dfsane` essentially has two loops, an outer loop and an inner loop, with parameters that control how the algorithm operates in each loop.³⁹ In the outer loop, at iteration k , one chooses the next iterate of \mathbf{p}_{k+1} as

$$\mathbf{p}_{k+1} = \mathbf{p}_k + \alpha_k d_k. \quad (\text{B4})$$

The variable α_k is a spectral steplength, whereas the variable d_k is equal to the negative of the first-order conditions, $-\mathbf{f}(\mathbf{p}_k)$.⁴⁰ Varadhan and Gilbert (2009) provide three different methods of choosing the α_k . The inner loop is a line search over \mathbf{p}_{k+1} . Occasionally the candidate \mathbf{p}_{k+1} will appear to be further from the solution of the system than \mathbf{p}_k . If this happens, the line search adjusts the candidate \mathbf{p}_{k+1} so that the following inequality holds:

$$\mathbf{f}(\mathbf{p}_{k+1})' \mathbf{f}(\mathbf{p}_{k+1}) \leq \max_{0 \leq j \leq M} \mathbf{f}(\mathbf{p}_{k-j})' \mathbf{f}(\mathbf{p}_{k-j}) + \eta_k - \gamma \alpha_k^2 \mathbf{f}(\mathbf{p}_k)' \mathbf{f}(\mathbf{p}_k). \quad (\text{B5})$$

The parameter M determines the degree of nonmonotonicity in the line search—if $M = 0$, then the line search must be monotone in the sense that iteration $k + 1$ must get closer to the root of the system than iteration k . When we compute numerical equilibrium, we use a tolerance level of 1E-13, a maximum number of iterations of 600, $M = 10$, and Varadhan and Gilbert's (2009) method number two for choosing α_k . Occasionally, `dfsane` fails to compute a numerical equilibrium at these default values. If this happens, we follow Varadhan and Gilbert (2009) and attempt to find an equilibrium for different methods of computing α_k , and different values of M .⁴¹ It is only in extremely rare cases that the algorithm fails to find an equilibrium in the inner loop. When this happens, we have found that the associated candidate parameter vectors tend to be less economically reasonable and may be consistent with equilibria that are simply too distant from observed prices. When this occurs, we construct regional-level metrics based on the price vector that comes closest to satisfying our definition of numerical equilibrium. We note that the optimizer moves away from such points relatively quickly, and at the estimates parameters the inner loop converges in all time periods.

To further speed the inner loop computations, we reexpress the first-order condition of B2 such that inversion of $\mathbf{\Omega}(\mathbf{p}; \mathbf{X}, \boldsymbol{\theta})$ is avoided. The computation of equilibrium for each time period can be parallelized, which further speeds the inner loop calculations. We also note that were production characterized by constant marginal costs, then one could further ease the computational burden of the inner loop by solving for equilibrium prices in each consumer area separately.

We constrain the signs and/or magnitudes of some parameters based on our understanding of economic theory and the economics of the portland cement industry, because some parameter vectors hinder the computation of numerical equilibrium in the inner loop. For instance, a positive price coefficient would preclude the existence of Bertrand-Nash equilibrium. We use the following constraints: the price and distance coefficients (β_1 and β_2) must be negative; the coefficients on the marginal cost shifters ($\boldsymbol{\alpha}$) and the overutilization cost (γ) must be positive; and the coefficients on the inclusive value (λ) and the utilization threshold (ν) must be between zero and one. We use nonlinear transformations to implement the constraints. We estimate the price coefficient using $\tilde{\beta}_1 = \log(-\beta_1)$ in the GMM procedure, and we estimate the inclusive value coefficient using $\tilde{\lambda} = \log(\frac{\lambda}{1-\lambda})$. We calculate standard errors with the delta method.

References

- ANDERSON, S.P. AND DE PALMA, A. "Spatial Price Discrimination with Heterogeneous Products." *Review of Economic Studies*, Vol. 55 (1988), pp. 573–592.
- ANDERSON, S.P., DE PALMA, A., AND THISSE, J.-F. "Spatial Price Policies Reconsidered." *Journal of Industrial Economics*, Vol. 38 (1989), pp. 1–18.
- BARTLESMAN, E.J., BECKER, R.A., AND GRAY, W.B. NBER-CES manufacturing industry database, 2000. www.nber.org/data/nberces5809.html (Last accessed March 4, 2014).

³⁹ The procedure we outline here generally follows that of Varadhan and Gilbert (2009), who also provide a more in-depth explanation of the algorithm.

⁴⁰ Note that the standard approach to solving nonlinear equations uses Newton's method to compute the iterates:

$$\mathbf{p}_{k+1} = \mathbf{p}_k - J(\mathbf{p}_k)' \mathbf{f}(\mathbf{p}_k),$$

where J is the Jacobian of the first order conditions. Calculating the Jacobian would be extremely slow in a large scale system such as ours.

⁴¹ We iterate over methods 2, 1, and 3 for α_k , and over $M = 10$ and $M = 50$.

- BERRY, S. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, Vol. 25 (2) (1994), pp. 242–262.
- BERRY, S., LEVINSOHN, J., AND PAKES, A. "Automobile Prices in Market Equilibrium." *Econometrica*, Vol. 64 (4) (July 1995), pp. 841–890.
- CARDELL, S.N. "Variance Components Structures for the Extreme Value and Logistic Distributions with Applications to Models of Heterogeneity." *Journal of Economic Theory*, Vol. 13 (1997), pp. 185–213.
- CHICU, M. "Dynamic Investment and Deterrence in the US Cement Industry." Job Market Paper, Northwestern University, 2012.
- DAVIS, P. "Spatial Competition in Retail Markets: Movie Theaters." *RAND Journal of Economics*, Vol. 37 (2006), pp. 964–982.
- DECANIO, S.J. "Delivered Pricing and Multiple Basing Point Equilibria: A Reevaluation." *Quarterly Journal of Economics*, Vol. 99 (1984), pp. 329–349.
- DUNNE, T., KLIMEK, S., AND SCHMITZ, J.A. "Does Foreign Competition Spur Productivity? Evidence from Post-WWII US Cement Manufacturing." Federal Reserve Bank of Minneapolis Staff Report, (2009).
- ENGIN, A., SYVERSON, C., AND HORTAÇSU, A. "Why Do Firms Own Production Chains?" Mimeo, 2012.
- EPA. *Regulatory Impact Analysis: National Emission Standards for Hazardous Air Pollutants from the Portland Cement Manufacturing Industry*. Prepared by RTI International, 2009.
- GRONBERG, T.J. AND MEYER, J. "Spatial Pricing, Spatial Rents, and Spatial Welfare." *Quarterly Journal of Economics*, Vol. 97 (1982), pp. 633–644.
- HADDOCK, D.D. "Basing-Point Pricing: Competitive vs. Collusive Theories." *American Economic Review*, Vol. 72 (1982), pp. 289–306.
- HANSEN, L. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, Vol. 50 (1982), pp. 1029–1054.
- HOBBS, B.F. "Mill Pricing versus Spatial Price Discrimination under Bertrand and Cournot Spatial Competition." *Journal of Industrial Organization*, Vol. 35 (1986), pp. 173–191.
- HOTELLING, H. "Stability in Competition." *Economic Journal*, Vol. 39 (1929), pp. 41–57.
- HOUE, J.-F. "Spatial Differentiation and Vertical Contracts in Retail Markets for Gasoline." *American Economic Review*, Vol. 102 (5) (2012), pp. 2147–2182.
- JANS, I. AND ROSENBAUM, D. "Multimarket Contact and Pricing: Evidence from the US Cement Industry." *International Journal of Industrial Organization*, Vol. 15 (May 1997), pp. 391–412.
- KARLSON, S.H. "Competition and Cement Basing Point: F.O.B. Destination, Delivered from Vol. Where?" *Journal of Regional Science*, Vol. 30 (1990), 75–88.
- KATZ, M.J. "Price Discrimination and Monopolistic Competition." *Econometrica*, Vol. 52 (1984), pp. 1453–1471.
- KAYSER, C. "Basing Point Pricing and Public Policy." *Quarterly Journal of Economics*, Vol. 63 (1949), pp. 189–314.
- LA CRUZ, W., MARTÍNEZ, J., AND RAYDAN, M. "Spectral Residual Method without Gradient Information for Solving Large-Scale Nonlinear Systems of Equations." *Mathematics of Computation*, Vol. 75 (2006), pp. 1429–1448.
- LEVENBERG, K. "A Method for the Solution of Certain Non-linear Problems in Least squares." *Quarterly Journal of Mathematics*, Vol. 2 (1944), pp. 164–168.
- MARQUARDT, D. "An Algorithm for Least-Squares Estimation of Nonlinear Parameters." *SIAM Journal of Applied Mathematics*, Vol. 11 (1963), 431–331.
- McFADDEN, D. "Econometric Models of Probabilistic Choice." In C.F. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data*. Cambridge, Mass.: MIT Press, 1981.
- McMANUS, B. "Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Coffee." *RAND Journal of Economics*, Vol. 38 (2007), 512–532.
- MILLER, N. AND OSBORNE, M. "Consistency and Asymptotic Normality for Equilibrium Models with Partially Observed Outcome Variables." *Economics Letters*, Vol. 123 (2) (2014), pp. 70–75.
- NELDER, J. AND MEAD, R. "A Simplex Method for Function Minimization." *Computer Journal*, Vol. 7 (1965), pp. 308–313.
- NEVO, A. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, Vol. 69 (2001), pp. 307–342.
- NEWBY, W. AND WEST, K. "A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, Vol. 55 (1987), pp. 703–708.
- NEWBY, W.K. AND McFADDEN, D. "Large Sample Estimation and Hypothesis Testing." *Handbook of Econometrics*, Vol. 4 (1994), 2112–2245.
- NEWMARK, C. "Price and Seller Concentration in Cement: Effective Oligopoly or Misspecified Transportation Cost?" *Economics Letters*, Vol. 60 (2) (August 1998), pp. 243–250.
- PINSKE, J., SLADE, M.E. AND BRETT, C. "Spatial Price Competition: A Semiparametric Approach." *Econometrica*, Vol. 70 (3)(2002), pp. 1111–1153.
- PORTER, R.H. AND ZONA, J.D. "Ohio School Milk Markets: An Analysis of Bidding." *RAND Journal of Economics*, Vol. 30 (1999), pp. 263–288.
- ROSENBAUM, D. AND READING, S. "Market Structure and Import Share: A Regional Market Analysis." *Southern Economic Journal*, Vol. 54 (3)(1988), pp. 694–700.
- ROSENBAUM, D.I. "Efficiency versus Collusion: Evidence Cast in Cement." *Review of Industrial Organization*, Vol. 9 (1994), pp. 379–392.

- RYAN, S. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica*, Vol. 80, No. 3 (2012), pp. 1019–1062.
- SALOP, S. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics*, Vol. 10 (1979), pp. 141–156.
- SALVO, A. "Inferring Market Power under the Threat of Entry: The Case of the Brazilian Cement Industry." *RAND Journal of Economics*, Vol. 41 (2)(2010a), pp. 326.
- . "Trade Flows in a Spatial Oligopoly: Gravity Fits Well, But What Does It Explain?" *Canadian Journal of Economics*, Vol. 43 (2010b), pp. 63–96.
- SCHERER, F.M., WITH BECKENSTEIN, AL, KAUFER, E., AND MURPHY, R.D. *The Economics of Multi-Plant Operation: An International Comparisons Study*. Cambridge, Mass.: Harvard University Press, 1975.
- SMALL, K. AND ROSEN, H. "Applied Welfare Economics with Discrete Choice Methods." *Econometrica*, Vol. 49 (1981), pp. 105–130.
- SYVERSON, C. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, Vol. 112 (2004), pp. 1181–1222.
- SYVERSON, C. AND HORTAÇSU, A. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." *Journal of Political Economy*, Vol. 115 (2007), pp. 250–301.
- THISSE, J.-F. and VIVES, X. "On the Strategic Choice of Spatial Price Policy." *American Economic Review*, Vol. 78 (1988), pp. 122–137.
- THOMADSEN, R. "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries." *RAND Journal of Economics*, Vol. 36 (2005), pp. 908–929.
- VARADHAN, R. AND GILBERT, P. "BB: An R Package for Solving a Large System of Nonlinear Equations and for Optimizing a High-Dimensional Nonlinear Objective Function." *Journal of Statistical Software*, Vol. 32 (2009), pp. 1–26.
- VOGEL, J. "Spatial Competition with Heterogeneous Firms." *Journal of Political Economy*, Vol. 116 (3)(June 2008), pp. 423–466.