ECONOMIC ANALYSIS GROUP
DISCUSSION PAPER

Competition among Spatially Differentiated Firms: An Empirical Model with an Application to Cement

By

Nathan H. Miller* and Matthew Osborne**

EAG 10-2 March 2010

EAG Discussion Papers are the primary vehicle used to disseminate research from economists in the Economic Analysis Group (EAG) of the Antitrust Division. These papers are intended to inform interested individuals and institutions of EAG's research program and to stimulate comment and criticism on economic issues related to antitrust policy and regulation. The Antitrust Division encourages independent research by its economists. The views expressed herein are entirely those of the authors and are not purported to reflect those of the United States Department of Justice or the Bureau of Economic Analysis.

Information on the EAG research program and discussion paper series may be obtained from Russell Pittman, Director of Economic Research, Economic Analysis Group, Antitrust Division, U.S. Department of Justice, LSB 9446, Washington DC 20530, or by e-mail at russell.pittman@usdoj.gov. Comments on specific papers may be addressed directly to the authors at the same mailing address or at their email addresses.

Recent EAG Discussion Paper and EAG Competition Advocacy Paper titles are listed at the end of this paper. To obtain a complete list of titles or to request single copies of individual papers, please write to Janet Ficco at the above mailing address or at janet.ficco@usdoj.gov or call (202) 307-3779. In addition, recent papers are now available on the Department of Justice website at http://www.usdoj.gov/atr/public/eag/discussion_papers.htm. Beginning with papers issued in 1999, copies of individual papers are also available from the Social Science Research Network at www.ssrn.com.

----

* Economic Analysis Group, Antitrust Division, U.S. Department of Justice. Email: nathan.miller@usdoj.gov.
** Bureau of Economic Analysis. Email: matthew.osborne@bea.gov. We thank Robert Johnson, Ashley Langer, Russell Pittman, Chuck Romeo, Jim Schmitz, Charles Taragin, Raphael Thomadsen, and seminar participants at the Bureau of Economic Analysis, Georgetown University, and the U.S. Department of Justice for valuable comments. Sarolta Lee, Parker Sheppard, and Vera Zavin provided research assistance.

Abstract

The theoretical literature of industrial organization shows that the distances between consumers and firms have first-order implications for competitive outcomes whenever transportation costs are large. To assess these effects empirically, we develop a structural model of competition among spatially differentiated firms and introduce a GMM estimator that recovers the structural parameters with only regional-level data. We apply the model and estimator to the portland cement industry. The estimation fits, both in-sample and out-of-sample, demonstrate that the framework explains well the salient features of competition. We estimate transportation costs to be \$0.30 per tonne-mile, given diesel prices at the 2000 level, and show that these costs constrain shipping distances and provide firms with localized market power. To demonstrate policy-relevance, we conduct counter-factual simulations that quantify competitive harm from a hypothetical merger. We are able to map the distribution of harm over geographic space and identify the divestiture that best mitigates harm.

1 Introduction

Geography is understudied in the empirical literature of industrial organization. Although the theoretical literature has established that the physical distances between firms and consumers have first-order implications for competitive outcomes whenever transportation costs are large (e.g., Hotelling (1929), d'Aspremont, Gabszewicz, and Thisse (1979), Salop (1979), Thisse and Vives (1988), Economides (1989), Vogel (2008)), the complexities associated with modeling spatial differentiation have made it difficult to translate theoretical insights into workable empirical models.[1]

Standard empirical methodologies simply sidestep spatial differentiation through the delineation of distinct geographic markets. This simplifies estimation but requires the dual assumptions that (1) transportation costs are sufficiently large to preclude substantive competition across market boundaries, and (2) transportation costs are sufficiently small that spatial differentiation is negligible within markets. It can be difficult to meet both conditions.[2] In practice, markets are often based on political borders of questionable economic significance (e.g., state or county lines). Nonetheless, market delineation is employed routinely in studies of industries characterized by high transportation costs, including ready-mix concrete (e.g., Syverson (2004), Syverson and Hortaçsu (2007), Collard-Wexler (2009)), portland cement (e.g., Salvo (2008), Ryan (2009)), and paper (e.g., Pesendorfer (2003)).[3]

Our purpose is to introduce an alternative empirical framework. To that end, we develop and estimate a structural model of competition among spatially differentiated firms that accounts for transportation costs in a realistic and tractable manner. We focus on production and consumption within a two dimensional Euclidean space, which we refer to informally as a geographic space. Competition involves a discrete number of plants, each endowed with a physical location, and a continuum of consumers that spans the space. Each plant sets a distinct price to each consumer, taking into consideration its proximity to the consumer and the proximities of its competitors. Thus, plants discriminate between elastic and inelastic consumers based on the pre-determined plant and consumer locations, and the model resembles the theoretical work of Thisse and Vives (1988). Competitive outcomes depend on the magnitude of transportation costs and the degree of spatial differentiation

[1] Surely it would be too strong to claim that, in empirical industrial organization, space is the final frontier.

[2] Syverson (2004) discusses how this tension can compel researchers to seek compromise between markets that are "too small" and markets that are "too large". It is sometimes argued that markets that are too large may overstate the intensity of competition while markets that are too small may understate competition.

[3] These valuable contributions focus on wide range of topics, including the competitive impacts of horizontal and vertical mergers, heterogeneity in plant productivity and its implications for competition, the inference of market power, and dynamic investment decisions.

within the geographic space.

We discretize the geographic space into small "consumer areas" to operationalize the model. Since each plant may ship to each consumer area, the model diverges starkly from more conventional approaches in which plants do not compete across market boundaries. We employ standard differentiated-product methods to model competition within consumer areas. On the supply side, domestic plants compete in prices given capacity constraints and the existence of a competitive fringe of foreign importers. On the demand side, consumers select the plant that maximizes utility, taking into consideration their proximity to the plants, the plant prices, and a nested logit error term. To be clear, the plants are differentiated primarily by price and location – we assume that the product itself is homogeneous. We derive an equation that characterizes the equilibrium price and market share for each plant-area pair as a function of data and parameters.[4]

The central challenge for estimation is that prices and market shares are unobserved in the data, at least at the plant-area level. We develop a generalized method of moments (GMM) estimator that exploits variation in data that are more often observed: regional level consumption, production, and prices. The key insight is that each candidate parameter vector corresponds to an equilibrium set of plant-area prices and market shares. We compute numerical equilibrium for each candidate parameter vector using a large-scale nonlinear equation solver developed in La Cruz, Martínez, and Raydan (2006). We then aggregate the predictions of the model to the regional level and evaluate the distance between the data and the predictions. The estimator can be interpreted as having inner and outer loops: the outer loop minimizes an objective function over the parameter space while the inner loop computes numerical equilibrium for each candidate parameter vector. We show that the estimator consistently recovers the structural parameters of the data generating process in an artificial data experiment.

We apply the model and the estimator to the portland cement industry in the U.S. Southwest over the period 1983-2003. The choice of industry conveys at least three substantive advantages to the analysis. First, transportation costs contribute substantially to overall consumer costs because portland cement is inexpensive relative to its weight. Second, it may be reasonable to treat portland cement as a homogenous product because strict industry standards govern the production process.[5] This conformity matches the simplicity of

[4]We refer to the fraction of potential demand in a consumer area that is captured by a given plant as the "market share" of the plant. We select the term purely for expositional convenience. We do not argue that consumer areas reflect antitrust markets in any sense; indeed, a defining characteristic of the model is that it avoids market delineation entirely.

[5]Many plants produce a number of different types of portland cement, each with slightly different spec-

the demand system, in which spatial considerations (e.g., plant and consumer locations) are the main source of plant heterogeneity. Third, high quality data on the industry are publicly available. We obtain information on regional consumption, production, and average prices, as well as limited information on cross-region shipments, from annual publications of the United States Geological Survey. We pair these regional-level metrics with information on the location and characteristics of portland cement plants from publications of the Portland Cement Association. We exploit variation in these data to estimate the model.

The results of estimation suggest that consumers pay roughly $0.30 per tonne-mile, given diesel prices at the 2000 level.[6] Given the shipping distances that arise in numerical equilibrium, this translates into an average transportation cost of $24.61 per metric tonne over the sample period – sufficient to account for 22 percent of total consumer expenditure. Costs of this magnitude have real effects on competition in the industry. We focus on two such effects: First, transportation costs constrain the distance that cement can be shipped economically. The results indicate that cement is shipped only 92 miles on average between the plant and the consumer; by contrast, a counter-factual simulation suggests that the average shipment would be 276 miles absent transportation costs. Second, transportation costs insulate firms from competition and provide localized market power. For instance, the prices that characterize numerical equilibrium decrease systematically in the distance between the plant and the consumer, as do the corresponding market shares.[7]

The estimation procedure produces impressive in-sample and out-of-sample fits despite parsimonious demand and marginal cost specifications. The model predictions explain 93 percent of the variation in regional consumption, 94 percent of the variation in regional production, and 82 percent of the variation in regional prices. The model predictions also explain 98 percent of the variation in cross-region shipments, even though we withhold the bulk of these data from estimation. As we detail in an appendix, the quality of these fits is underscored by the rich time-series variation in these data. Together, the regression fits suggest that a small quantity of exogenous data, properly utilized, may be sufficient to explain some of the most salient features of competition in the portland cement industry. We interpret this as substantial support for the power of the analytical framework.

---

ifications and characteristics (e.g., superior early strength or higher sulfate resistance). These products are close substitutes for most construction projects.

[6]Strictly speaking, the model identifies consumer willingness-to-pay for proximity to the plant, which incorporates transportation costs as well as any other distance-related costs (e.g., reduced reliability). We refer to this willingness-to-pay as the transportation cost, although the concepts may not be precisely equivalent.

[7]These patterns are precisely what economic theory would predict given that consumers pay the costs of transportation in the portland cement industry.

We suspect that our method may prove useful for future research and policy endeavors relating to international trade, environmental economics, and industrial organization. One such application is merger simulation, an important tool for competition policy. We use counter-factual simulations to evaluate the effects of a hypothetical merger between two multi-plant cement firms in 1986. We find that the merger reduces consumer surplus by $1.4 million if no divestitures are made, and we are able to map the distribution of this harm over the U.S. Southwest. The overall magnitude of the effect is modest relative to the amount of commerce – by way of comparison, we calculate total pre-merger consumer surplus to be more than $239 million. We then consider the six possible single-plant divestitures, and find that the most powerful reduces consumer harm by 56 percent.[8]

At least two caveats are important. First, the estimation procedure rests on the uniqueness of equilibrium at each candidate parameter vector, but there is no theoretical reason to expect this condition to hold generally. To assess the issue, we conduct a Monte Carlo experiment that computes numerical equilibrium using several different starting points for each of 6,300 randomly drawn candidate parameter vectors. The results are strongly supportive of uniqueness, at least in our application (see the appendices for details). Second, the promise of spatial differentiation creates incentives for firms to locate optimally in order to secure a base of profitable customers, provide separation from an efficient competitor, and/or deter nearby entry. We abstract from these considerations entirely and instead assume that firm location is pre-determined and exogenous. Nonetheless, our framework could help define stage-game payoffs in more dynamic models that endogenize firm location choices (e.g., as in Seim (2006) and Aguirregabiria and Vicentini (2006)).

Our work builds on recent contributions to the industrial organization literature that model competition among spatially differentiated retail firms (e.g., Thomadsen (2005), Davis (2006), McManus (2009)). These papers employ a framework in which each firm sets a single price to all consumers, consistent with practice in most retail settings, and recover the structural parameters with firm-level data and more standard estimation techniques (e.g., Berry, Levinsohn, and Pakes (1995)).[9] We make two distinct contributions to this literature. First, our model extends the existing framework to incorporate firms that price discriminate

[8]We refer to the divesture plan that offsets the greatest amount of consumer harm among the set of single-plant divestitures as the most powerful. We do not attempt to characterize, in any way, the appropriate course of action for an antitrust authority.

[9]Thomadsen (2005) may be the closest antecedent to our work. Thomadsen shows that a supply-side equilibrium condition can be substituted for firm-level market shares in estimation. We develop the potency of equilibrium conditions more fully: given an equilibrium condition and aggregate data, estimation is feasible with neither firm-level prices nor firm-level market shares. Aguirregabiria and Vicentini (2006) develop a sophisticated model of spatial differentiation but do not take the model to data.

among consumers. In such a setting, firm-level data are no longer sufficient to support standard estimation techniques. Thus, our second contribution is the demonstration that estimation is feasible with relatively aggregated data. Of course, the regional-level data we use in our application would also support estimation in simpler retail settings.

More generally, our results indicate that firm-level heterogeneity can affect competitive outcomes even when the product is relatively homogenous. Refinements to the standard toolbox available for structural research in these settings have been outstripped in empirical industrial organization by models of observed and (especially) unobserved product heterogeneity. We suspect that models of competition for homogenous-product industries are currently of heightened value, not only because of the imbalance in the literature, but also because these industries provide fertile testing grounds for methodological innovations regarding industry dynamics (e.g., as in Ryan (2009)). We hope that the framework we introduce helps, incrementally, to redirect the attention of researchers to these settings.

The paper proceeds as follows. In Section 2, we sketch the relevant institutional details of the portland cement industry, focusing on transportation costs, production technology, and trends in production and consumption. We develop the empirical model of Bertrand-Nash competition in Section 3 and discuss our data sources in Section 4. Then, in Section 5, we develop the estimator and provide identification arguments. We present the estimation results in Section 6, discuss the merger simulations in Section 7, and then conclude.

2 The Portland Cement Industry

Portland cement is a finely ground dust that forms concrete when mixed with water and coarse aggregates such as sand and stone. Concrete is an essential input to many construction and transportation projects, either as pourable fill material or as pre-formed concrete blocks, because its local availability and lower maintenance costs make it more economical than substitutes such as steel, asphalt, and lumber (e.g., Van Oss and Padovani (2002)).[10]

Most portland cement is shipped by truck to ready-mix concrete plants or construction sites, in accordance with contracts negotiated between individual purchasers and plants.[11] Transportation costs contribute substantially to overall consumer expenditures because port-

[10] 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. We defer detailed discussion of these sources for expositional convenience.

[11] A smaller portion is shipped by train or barge to terminals, and only then distributed to consumers by truck. Shipment via terminals reduces transportation costs for more distant consumers. Roughly 23 percent of portland cement produced in the United States was shipped through terminals in 2003.

land cement is inexpensive relative to its weight, a fact that is well understood in the academic literature. For example, Scherer et al (1975) estimates that transportation would have accounted for roughly one-third of total consumer expenditures on a hypothetical 350-mile route between Chicago and Cleveland, and a 1977 Census Bureau study determines that most portland cement is consumed locally – for example, more than 80 percent is transported within 200 miles.[12] More recently, Salvo (2010) presents evidence consistent with the importance of transportation costs in the Brazilian portland cement industry.

A recent report prepared for the Environmental Protection Agency identifies five main variable input costs of production: raw materials, coal, electricity, labor, and kiln maintenance (EPA (2009)). In the production process itself, a feed mix composed of limestone and supplementary materials is fed into large rotary kilns that reach temperatures of 1400-1450° Celsius. The combustion of coal is the most efficient way to generate this extreme heat. Kilns generally operate at peak capacity with the exception of an annual shutdown period for maintenance. It is possible to adjust output by extending or shortening the maintenance period – for example, plants may simply forego maintenance at the risk of kiln damage and/or breakdowns. The feed mix exits the kiln as semi-fused clinker modules. Once cooled, the clinker is mixed with a small amount of gypsum, placed into a grinding mill, and ground into tiny particles averaging ten micrometers in diameter. This product – portland cement – is shipped to purchasers either in bulk or packaged in smaller bags.

We focus on the production and consumption of portland cement in the U.S. Southwest – by which we mean California, Arizona and Nevada – over the period 1983-2003. This region accounts for roughly 15 percent of domestic portland cement production and consumption during the sample period. Figure 1 provides a map of the region based on the plant locations in the final year of the sample. As shown, most plants are located along an interstate highway and nearby one or more population centers. Although some firms own more than one plant, production capacity in the area is not particularly concentrated – the capacity-based Herfindahl-Hirschman Index (HHI) of 1260 is well below the threshold level that defines highly concentrated markets in the 1992 Merger Guidelines. The plants also face competition from foreign importers that ship portland cement through the customs offices of San Francisco, Los Angeles, San Diego, and Nogales. Still, transportation costs may insulate some plants from both foreign and domestic competition.[13]

[12]Scherer et al (1975) examined more than 100 commodities and determined that the transportation costs of portland cement were second only to those of industrial gases. Other commodities identified as having particularly high transportation costs include concrete, petroleum refining, alkalies/chlorine, and gypsum.

[13]We observe little entry and exit over the sample period. The sole entrant (Royal Cement) began operations in 1994 and the only two exits occurred in 1988. This stability is consistent with substantial sunk costs

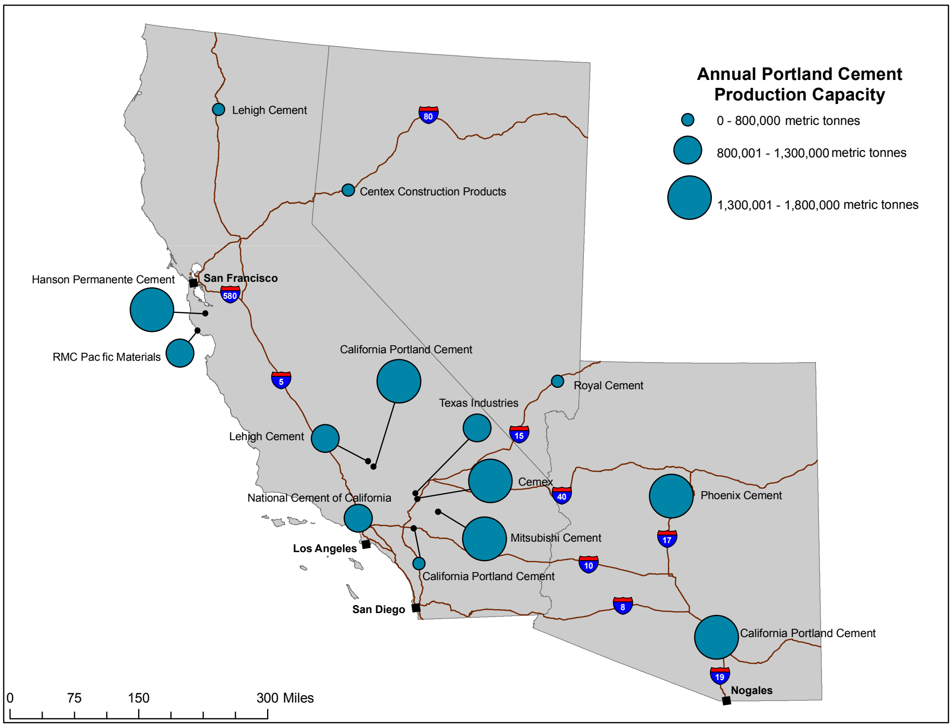

Figure 1: Portland Cement Production Capacity in the U.S. Southwest circa 2003.

In Figure 2, we plot total consumption and production in the U.S. Southwest over the sample period, together with two measures of foreign imports. Several patterns are apparent. Both consumption and production increase over 1983-2003, and both metrics are highly cyclical. However, consumption is more cyclical than production, so that the gap between consumption and production increases in overall activity; foreign importers provide additional supply whenever domestic demand outstrips domestic capacity. Strikingly, observed foreign imports are nearly identical to "apparent imports," which we define as consumption minus production, consistent with negligible net trade between the U.S. Southwest and other domestic areas. Finally, we note that the average free-on-board price charged by domestic plants in the region (not shown) falls over the sample period from \$107 per metric tonne to \$74 per metric tonne, primarily due to lower coal and electricity costs.[14]

of plant construction, as documented in Ryan (2009). Plant ownership is somewhat more fluid; we observe fourteen changes in plant ownership, spread among nine of the sixteen plants.

[14]Prices are in real 2000 dollars.

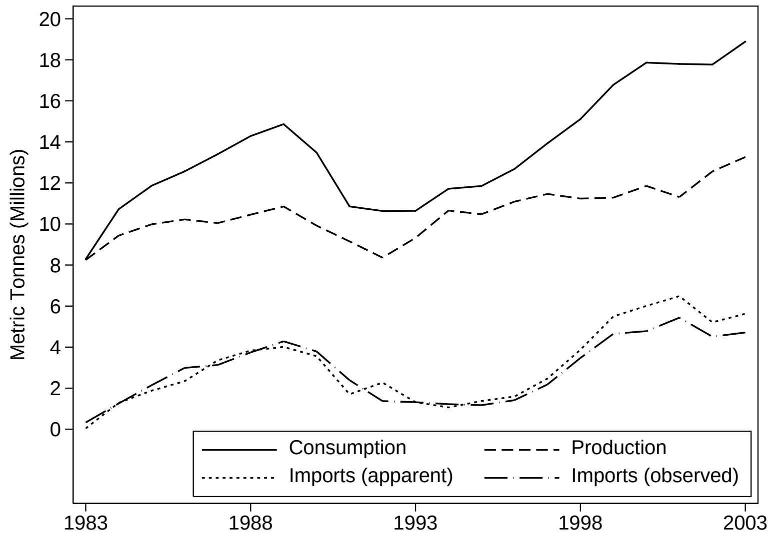

Figure 2: Consumption, Production, and Imports of Portland Cement. Apparent imports are defined as consumption minus production. Observed imports are total foreign imports shipped into San Francisco, Los Angeles, San Diego, and Nogales.

3 Empirical Model

3.1 Overview

We develop a structural model of competition that accounts for transportation costs in a realistic and tractable manner. We focus on production and consumption within a single geographic space. Competition involves a discrete number of plants, each endowed with a physical location, and a continuum of consumers that spans the space. Each plant sets a distinct price to each consumer, taking into consideration its proximity to the consumer and the proximities of its competitors. To operationalize the model, we discretize the geographic space into small consumer areas. We employ standard techniques to model competition within consumer areas: plants charge area-specific prices subject to exogenous capacity constraints, and consumer demand is nested logit. We derive an equation that characterizes the equilibrium price and market share for each plant-area pair.

3.2 Supply

We take as given that there are F firms, each of which operates some subset $\Im_f$ of the J plants. Each plant is endowed with a set of attributes, including a physical location. Consumers exist in N different areas that span the geographic space. Every plant can ship into every area. Firms set prices at the plant-area level to maximize variable profits:

$$\pi_f = \underbrace{\sum_{j \in \Im_f} \sum_n P_{jn} Q_{jn}(\boldsymbol{\theta_d})}_{\text{variable revenues}} - \underbrace{\sum_{j \in \Im_f} \int_0^{Q_j(\boldsymbol{\theta_d})} MC_j(Q; \boldsymbol{\theta_c}) dQ,}_{\text{variable costs}} \tag{1}$$

where Q_{jn} and P_{jn} are the quantity and price, respectively, of plant j in area n, MC_j is the marginal cost of production, and Q_j is total plant production (i.e., $Q_j = \sum_n Q_{jn}$). The vectors $\boldsymbol{\theta_c}$ and $\boldsymbol{\theta_d}$ include the supply and demand parameters, respectively, and together form the joint parameter vector $\boldsymbol{\theta} = (\boldsymbol{\theta_c}, \boldsymbol{\theta_d})$.

The marginal cost function can accommodate most forms of scale economies or diseconomies. We choose a functional form that captures the plant-level capacity constraints that are important for portland cement production. In particular, we let marginal costs increase flexibly in production whenever production exceeds some threshold level. The marginal cost function is:

$$MC(\boldsymbol{w}_j, Q_j(\boldsymbol{\theta_d}); \boldsymbol{\theta_c}) = \boldsymbol{w}_j'\boldsymbol{\alpha} + \gamma\, 1\left\{\frac{Q_j(\boldsymbol{\theta_d})}{CAP_j} > \nu\right\}\left(\frac{Q_j(\boldsymbol{\theta_d})}{CAP_j} - \nu\right)^{\phi}, \tag{2}$$

where $\boldsymbol{w_j}$ is a vector that includes the relevant marginal cost shifters and CAP_j is the maximum quantity that plant j can produce. The parameter ν is the utilization threshold above which marginal costs increase in production, the parameter ϕ determines the curvature of the marginal cost function when utilization exceeds ν, and the combination $\gamma(1-\nu)^{\phi}$ represents the marginal cost penalty associated with production at capacity. The marginal cost function is continuously differentiable in production for $\phi > 1$. The vector of cost parameters is $\boldsymbol{\theta_c} = (\boldsymbol{\alpha}, \gamma, \nu, \phi)$.

We let the domestic firms compete against a competitive fringe of foreign importers, which we denote as plant $J+1$. We assume that the fringe is a non-strategic actor and that import prices are exogenously set based on some marginal cost common to all importers. The fringe is endowed with one or more geographic locations and ships into every consumer area. It sets a single price across all consumer areas (i.e., the fringe does not price discriminate), consistent with perfect competition among importers.

3.3 Demand

We specify a nested logit demand system that captures the two most important characteristics that differentiate portland cement plants: price and location. To that end, we assume that each area features many potential consumers. Each consumer either purchases cement from one of the domestic plants or the importer, or foregoes a purchase of portland cement altogether. We refer to the domestic plants and the importer as the inside goods, and refer to the option to forego a purchase as the outside good. We place the inside goods in a separate nest from the outside good.

We express the indirect utility that consumer i receives from plant j as a function of the relevant plant and location observables:

$$u_{ij} = \beta_0 + \boldsymbol{x}'_{jn}\boldsymbol{\beta} + \zeta_i + \lambda\epsilon_{ij}, \tag{3}$$

where the vector $\boldsymbol{x}_{jn}$ includes the price of cement and the distance between the plant and the area, as well as other plant-specific demand shifters. The idiosyncratic portion of the indirect utility function is composed of consumer-specific shocks to the desirability of the inside good (ζ_i) and the desirability of each plant (ϵ_{it}). We assume that the combination $\zeta_i + \lambda\epsilon_{ij}$ has an extreme value distribution in which the parameter λ characterizes the extent to which valuations of the inside good are correlated across consumers.[15] We normalize the mean utility of the outside good to zero, so that the indirect utility associated with foregoing cement purchase is $u_{i0} = \epsilon_{i0}$, with ϵ_{i0} also having the extreme value distribution. The vector of demand parameters is $\boldsymbol{\theta_d} = (\beta_0, \boldsymbol{\beta}, \lambda)$.[16]

The nested logit structure yields an analytical expression for the market shares captured by each plant. The market shares are specific to each plant-area pair because the relative desirability of each plant varies across areas:

$$S_{jn}(\boldsymbol{P}_n; \boldsymbol{\theta_d}) = \frac{\exp(\beta_0 + \lambda I_n)}{1 + \exp(\beta_0 + \lambda I_n)} * \frac{\exp(\boldsymbol{x}'_{jn}\boldsymbol{\beta})}{\sum_{k=1}^{J+1} \exp(\boldsymbol{x}'_{kn}\boldsymbol{\beta})}, \tag{4}$$

where $\boldsymbol{P}_n$ is a vector of area-specific prices and $I_n = \ln\left(\sum_{k=1}^{J+1} \exp(\boldsymbol{x}'_{kn}\boldsymbol{\beta})\right)$ is the inclusive value of the inside goods. The first factor in this expression is the marginal probability that a consumer in area n selects an inside good, and the second factor is the conditional probability that the consumer purchases from plant j given selection of an inside good. Both

[15] Cardell (1997) derives the conditions under which the specified taste shocks produce the extreme value distribution. Tastes are perfectly correlated if $\lambda = 0$ and tastes are uncorrelated if $\lambda = 1$. In the latter case the model collapses to a standard logit.

[16] We exclude product characteristics from the specification because portland cement is a homogenous good, at least to a first approximation. It may be desirable to control for observed product characteristics in other applications. (Though the presence of unobserved characteristics would pose a challenge to the estimation procedure.) Also, we assume that consumers pay the cost of transportation, consistent with practice in the cement industry. One could alternatively incorporate distance into the marginal cost function.

take familiar logit forms due to the distributional assumptions on idiosyncratic consumer tastes. The demand system maps cleanly into supply: the quantity sold by plant j to area n is $Q_{jn} = S_{jn}M_n$, where M_n is the potential demand of area n.[17]

3.4 Equilibrium

The profit function yields the first-order conditions that characterize equilibrium prices for each plant-area pair:

$$Q_{jn} + \sum_{n} \sum_{k \in \Im_f} (P_{kn} - MC_{kn}) \frac{\partial Q_{kn}}{\partial P_{jn}} = 0. \tag{5}$$

Since each plant competes in every consumer area and may price differently across areas, there are $J \times N$ first-order conditions. For notational convenience, we define the block-diagonal matrix $\boldsymbol{\Omega(P)}$ as the combination of $n = 1, \ldots, N$ sub-matrices, each of dimension $J \times J$. The elements of the sub-matrices are defined as follows:

$$\Omega^n_{jk}(\boldsymbol{P}) = \begin{cases} \frac{\partial Q_{jn}}{\partial P_{kn}} & \text{if } j \text{ and } k \text{ have the same owner} \\ 0 & \text{otherwise,} \end{cases} \tag{6}$$

where the demand derivatives take the nested logit forms. Thus, the elements of each sub-matrix $\Omega^n_{jk}(\boldsymbol{P})$ characterize substitution patterns within area n, and the matrix $\boldsymbol{\Omega(P)}$ has a block diagonal structure because prices in one area do not affect demand in other areas. We can now stack the first-order conditions:

$$\boldsymbol{P} = \boldsymbol{MC(P)} - \boldsymbol{\Omega(P)}^{-1}\boldsymbol{Q(P)}, \tag{7}$$

where $\boldsymbol{P}$, $\boldsymbol{MC(P)}$, and $\boldsymbol{Q(P)}$ are vectors of prices, marginal costs and quantities, respectively. Provided that marginal cost parameter ϕ exceeds one, the mappings $\boldsymbol{MC(P)}$, $\boldsymbol{\Omega(P)}$, and $\boldsymbol{Q(P)}$ are each continuously differentiable. Further, the price vector belongs to a compact set in which each price is (1) greater or equal to the corresponding marginal cost, and (2) smaller than or equal to the price a monopolist would charge to the relevant

[17]The substitution patterns between cement plants are characterized by the independence of irrelevant alternatives (IIA) within the inside good nest. We argue that IIA is a reasonable approximation for our application. Portland cement is purchased nearly exclusively by ready-mix concrete plants and other construction companies. These firms employ similar production technologies and compete under comparable demand conditions. We are therefore skeptical that meaningful heterogeneity exists in consumer preferences for plant observables (e.g., price and distance). Without such heterogeneity, the IIA property arises quite naturally – for example, the random coefficient logit demand model collapses to standard logit when the distribution of consumer preferences is degenerate.

area. Therefore, by Brower's fixed-point theorem, Bertrand-Nash equilibrium exists and is characterized by the price vector $\boldsymbol{P^*}$ that solves Equation 7.[18]

Spatial price discrimination is at the core of the firm's pricing problem: firms maximize profits by charging higher prices to nearby consumers and consumers without a close alternative. Aside from price discrimination, the firm's pricing problem follows standard intuition. For example, a firm that contemplates a higher price for cement from plant j to area n must evaluate a number of effects: (1) the tradeoff between lost sales to marginal consumers and greater revenue from inframarginal consumers; (2) whether the firm would recapture lost sales with its other plants; and (3) whether the lost sales would ease capacity constraints and make the plant more competitive in other consumer areas.

4 Data Sources and Summary Statistics

4.1 Data sources

We cull the bulk of our data from the Minerals Yearbook, an annual publication of the U.S. Geological Survey (USGS). The Minerals Yearbook is based on an annual census of portland cement plant and contains regional-level information on portland cement consumption, production, and free-on-board prices.[19] The four relevant regions include Northern California, Southern California, Arizona, and Nevada. We observe annual consumption in each region over the period 1983-2003. The USGS combines the Arizona and Nevada regions when reporting production and prices over 1983-1991, and no usable production or price information is available for Nevada over 1992-2003.[20] The Minerals Yearbook also includes information on the price and quantity of portland cement that is imported into the U.S. Southwest via the customs offices in San Francisco, Los Angeles, San Diego, and Nogales.

We also make use of more limited data on cross-region shipments from the California

[18]The existence proof follows Aguirregabiria and Vicentini (2006). Multiple equilibria may exist.

[19]The census response rate is typically well over 90 percent (e.g., 95 percent in 2003), and USGS staff imputes missing values for the few non-respondents based on historical and cross-sectional information. Other academic studies that feature these data include McBride (1983), Rosenbaum and Reading (1988), Rosenbaum (1994), Jans and Rosenbaum (1997), Ryan (2009), and Syverson and Hortaçsu (2007). The Minerals Yearbook provides the average free-on-board price charged the plants located in each region, rather than the price paid by the consumers in each region.

[20]The USGS combines Nevada with Idaho, Montana and Utah starting in 1992. We adjust the Arizona data to remove the influence of a single plant located in New Mexico whose production is aggregated into the region. We detail this adjustment in an appendix. Also, it is worth noting that the USGS does not intend for its regions to approximate geographic markets. Rather, the regions are delineated such that plant-level information cannot be backward-engineered from the Minerals Yearbook.

Letter, a second annual publication of the USGS. The California Letter provides information on the quantity of portland cement that is shipped from plants in California to consumers in Northern California, Southern California, Arizona, and Nevada. However, the level of aggregation varies over the sample period, some data are redacted to protect sensitive information, and no information is available before 1990. In total, we observe 96 data points:

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

We withhold the bulk of these data from the estimation procedure and instead use the data to conduct out-of-sample checks on the model predictions.[21]

We supplement the USGS data with basic plant-level information from the Plant Information Survey (PIS), an annual publication of the Portland Cement Association. The PIS provides the location of each portland cement plant in the United States, together with its owner, the annual kiln capacity, and various other kiln characteristics.[22] We approximate consumer areas using counties, which meet the criterion of being small relative to the overall geographic space – there are 90 counties in the U.S. Southwest. We collect county-level data from the Census Bureau on construction employment and residential construction permits to account for county-level heterogeneity in potential demand. Finally, we collect data on diesel, coal, and electricity prices from the Energy Information Agency, data on average wages of durable good manufacturing employees from the BEA, and data on crushed stone prices from the USGS; we exploit state-level variation for all but the diesel data.

[21] The California Letter is based on a monthly survey rather than on the annual USGS census. As a result, the data are not always consistent with the Minerals Yearbook. We normalize the data prior to estimation so that total shipments equal total production in each year.

[22] 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 percent) in the grinding mills.

4.2 Summary statistics

We provide summary statistics on consumption, production, and prices for each of the regions in Table 1. Some patterns stand out: First, substantial variation in each metric is available, both inter-temporally and across regions, to support estimation. Second, Southern California is larger than the other regions, whether measured by consumption or production. Third, consumption exceeds production in Northern California, Arizona, and Nevada; these shortfalls must be countered by cross-region shipments and/or imports. The observation that plants in these regions charge higher prices is consistent with transportation costs providing some degree of local market power.[23] Finally, imports are less expensive than domestically produced portland cement. This discrepancy may exist in part because the reported prices exclude duties; more speculatively, domestic producers may be more reliable or may maintain relationships with consumers that support higher prices.

In Table 2, we explore the spatial characteristics of the regions in more depth, based on the plant locations of 2003. First, the average county in Northern California is 65 miles from the nearest domestic cement plant. Since the comparable statistics for Southern California, Arizona, and Nevada are 72 miles, 92 miles, and 100 miles, respectively, one may infer that average transportation costs may differ substantively across the four regions. Second, the average *additional* distance to the second closest domestic plant is 44 miles in Northern California, 11 miles in Southern California, 82 miles in Arizona, and 77 miles in Nevada, suggestive that plants in Arizona and Nevada may hold more local market power than plants elsewhere. Finally, the average distance to the nearest customs office is 123 miles in Northern California, 110 miles in Southern California, 181 miles in Arizona, and 281 miles in Nevada, suggestive that imports may constrain domestic prices less severely in Arizona and Nevada. The latter two empirical patterns are consistent with the higher cement prices observed in Arizona and Nevada.

5 Estimation

5.1 Overview

The challenge for estimation is that prices and market shares are unobserved in the data, at least at the plant-area level. We develop a GMM estimator that exploits variation in data

[23]The data on cross-region shipments are also suggestive of large transportation costs. For example, more than 90 percent of portland cement produced in Northern California was shipped to consumers in Northern California over 1990-1999.

Table 1: Consumption, Production, and Prices

Description	Mean	Std	Min	Max
Consumption				
Northern California	3,513	718	2,366	4,706
Southern California	6,464	1,324	4,016	8,574
Arizona	2,353	650	1,492	3,608
Nevada	1,289	563	416	2,206
Production				
Northern California	2,548	230	1,927	2,894
Southern California	6,316	860	4,886	8,437
Arizona-Nevada	1,669	287	1050	2,337
Domestic Prices				
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
Import Prices				
U.S. Southwest	50.78	9.30	39.39	79.32

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.

Table 2: Distances between Counties and Plants

Description	Mean	Std	Min	Max
Miles to the closest plant				
Northern California	64.65	30.00	7.36	115.39
Southern California	72.28	39.74	18.58	127.46
Arizona	91.62	40.01	29.13	163.99
Nevada	100.04	61.97	17.38	232.03
Additional miles to the second closest plant				
Northern California	43.95	49.84	0.49	176.07
Southern California	11.22	8.83	0.49	31.08
Arizona	81.96	55.65	0.69	172.38
Nevada	77.38	43.28	12.82	177.20
Miles to the closest import point				
Northern California	122.65	67.42	4.33	283.00
Southern California	110.17	67.43	30.29	221.28
Arizona	180.80	90.26	14.07	314.91
Nevada	281.45	81.93	170.09	442.02

Distances are calculated based on plant locations in 2003. There are 46 counties in Northern California, 12 counties in Southern California, 15 counties in Arizona, and 17 counties in Nevada.

that are more often observed: regional level consumption, production, and prices. The key insight is that each candidate parameter vector corresponds to an equilibrium set of plant-area prices and market shares. We compute equilibrium numerically for each candidate parameter vector, aggregate predictions of the model to the regional-level, and then evaluate the "distance" between the data and the aggregate predictions. The estimation routine is an iterative procedure defined by the following steps:

1. Select a candidate parameter vector $\widetilde{\boldsymbol{\theta}}$.
2. Compute the equilibrium price and market share vectors.
3. Calculate regional-level metrics based on the equilibrium vectors.
4. Evaluate the regional-level metrics against the data.
5. Update $\widetilde{\boldsymbol{\theta}}$ and repeat steps 1-5 to convergence.

The estimation procedure can be interpreted as having both an inner loop and an outer loop: the inner loop computes equilibrium for each candidate parameter vector and the outer loop minimizes an objective function over the parameter space. We discuss the inner loop and the outer loop in turn, and then address some additional details.

5.2 Computation of numerical equilibrium

We use a large-scale nonlinear equation solver developed in La Cruz, Martínez, and Raydan (2006) to compute equilibrium. 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. We define a numerical Bertrand-Nash equilibrium as a price vector for which $\parallel \boldsymbol{\iota}(\boldsymbol{P}) \parallel / \dim(\boldsymbol{\iota}(\boldsymbol{P})) < \delta$, where $\parallel \cdot \parallel$ denotes the Euclidean norm operator and

$$\boldsymbol{\iota}(\boldsymbol{P}) = \boldsymbol{\Omega}(\boldsymbol{P})(\boldsymbol{P} - \boldsymbol{MC}(\boldsymbol{P})) - \boldsymbol{Q}(\boldsymbol{P}). \tag{8}$$

We denote the equilibrium prices that correspond to the candidate parameter vector $\widetilde{\boldsymbol{\theta}}$ as $\widetilde{P}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$, where $\boldsymbol{\chi}$ includes the exogenous data. We denote the corresponding equilibrium market shares as $\widetilde{S}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$. From a computational standpoint, our construction of $\boldsymbol{\iota}(\boldsymbol{P})$ avoids the burden of inverting $\boldsymbol{\Omega}(\boldsymbol{P})$ that would be required by the straight application of Equation 7. Further, the structure of the problem permits us to compute equilibrium separately for each period. The price vector that characterizes the equilibrium of a given period has length $J_t \times N$ so that, for example, the equilibrium price vector for 2003 has $14 \times 90 = 1,260$ elements.[24]

The empirical analogs to the computed plant-area prices and market shares are not observed in the data, so we aggregate the prices and market shares to construct more useful regional metrics. For notational precision, we define the sets $\aleph_r$ and $\jmath_r$ as the counties and plants, respectively, located in region r. The aggregated region-period metrics take the form:

$$\begin{aligned}
\widetilde{C}_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) &= \sum_{n \in \aleph_r} \left(1 - \widetilde{S}_{0nt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})\right) M_{nt} \\
\widetilde{Q}_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) &= \sum_{j \in \jmath_r} \widetilde{S}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) M_{nt} \\
\widetilde{P}_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) &= \sum_{j \in \jmath_r} \sum_{n} \frac{\widetilde{S}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) M_{nt}}{\sum_{j \in \jmath_r} \sum_n \widetilde{S}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) M_{nt}} \widetilde{P}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}),
\end{aligned} \tag{9}$$

where $\widetilde{C}_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$, $\widetilde{Q}_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$, and $\widetilde{P}_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$ are total consumption, total production, and weighted-average price, respectively. We calculate regional consumption for each of the four regions in the data – Northern California, Southern California, Arizona, and Nevada. We calculate production and prices for Northern California, Southern California, and a combined Arizona-Nevada region. We denote the empirical analogs C_{rt}, Q_{rt}, and P_{rt}.

We find that, in practice, we can better identify some of the model parameters by exploiting information on aggregated cross-region shipments. We denote the quantity of shipments from region r to region s as $\widetilde{Q}^s_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$. The shipments take the form:

$$\widetilde{Q}^s_{rt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) = \sum_{j \in \jmath_r} \sum_{n \in \aleph_s} \widetilde{S}_{jnt}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}) M_{nt}, \tag{10}$$

We calculate the quantity of portland cement produced by plants in California (both Northern and Southern) that is shipped to consumers in Northern California. The empirical analog, which we denote Q^s_{rt}, is available over the period 1990-2003. We withhold data on other cross-region shipments from estimation because the number of parameters that we estimate exceeds the lengths of the available time-series. The withheld data, however, provide

[24] We set δ =1e-13. Numerical error can propagate into the outer loop when the inner loop tolerance is substantially looser (e.g., 1e-7), which slows overall estimation time and can produce poor estimates. The inner loop tolerance is not unit free and must be evaluated relative to the price level. We also note that, in settings characterized by constant marginal costs, one could ease the computational burden of the inner loop by solving for equilibrium prices in each consumer area separately.

natural out-of-sample tests on the model predictions.

5.3 Objective function

We estimate the parameters using the standard GMM framework for systems of nonlinear regression equations (e.g., Greene (2003), page 369). The equations we use compare the metrics computed in the inner loop to their empirical analogs:

$$
\begin{aligned}
\boldsymbol{C}_r &= \widetilde{\boldsymbol{C}}_r(\boldsymbol{\theta}, \boldsymbol{\chi}) + \boldsymbol{e}_{\boldsymbol{r}}^{\mathbf{1}} \\
\boldsymbol{Q}_r &= \widetilde{\boldsymbol{Q}}_r(\boldsymbol{\theta}, \boldsymbol{\chi}) + \boldsymbol{e}_{\boldsymbol{r}}^{\mathbf{2}} \\
\boldsymbol{P}_r &= \widetilde{\boldsymbol{P}}_r(\boldsymbol{\theta}, \boldsymbol{\chi}) + \boldsymbol{e}_{\boldsymbol{r}}^{\mathbf{3}} \\
\boldsymbol{Q}_r^s &= \widetilde{\boldsymbol{Q}}_r^s(\boldsymbol{\theta}, \boldsymbol{\chi}) + \boldsymbol{e}_{\boldsymbol{rs}}^{\mathbf{4}}.
\end{aligned} \tag{11}
$$

We write the equations in vector form with one element per period. There are four consumption moments, three production moments, three price moments, and one cross-region shipments moment. We have 21 observations on each of the consumption, production and price moments, and 14 observations on the cross-region shipments moment. We interpret the disturbances as measurement error, and assume the disturbances to have expectation zero and contemporaneous covariance matrix $\boldsymbol{\Sigma}$.

The GMM estimator is:

$$
\widehat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\Theta}} \boldsymbol{e}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})' \boldsymbol{A}^{-1} \boldsymbol{e}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi}), \tag{12}
$$

where $\boldsymbol{e}(\widetilde{\boldsymbol{\theta}}, \boldsymbol{\chi})$ is a vector of empirical disturbances obtained by stacking the nonlinear regression equations and $\boldsymbol{A}$ is a positive definite weighting matrix. We employ the usual two-step procedure to obtain consistent and efficient estimates (Hansen (1982)). We first minimize the objective function using $\boldsymbol{A} = \boldsymbol{I}$. We then estimate the contemporaneous covariance matrix and minimize the objective function a second time using the weighting matrix $\boldsymbol{A} = \widehat{\boldsymbol{\Sigma}} \otimes \boldsymbol{I}$. We compute standard errors that are robust to both heteroscedasticity and arbitrary correlations among the error terms of each period, using the methods of Hansen (1982) and Newey and McFadden (1994).[25]

[25] Measurement error that is zero in expectation is sufficient for consistency. Estimation of the contemporaneous covariance matrix $\boldsymbol{\Sigma}$ is complicated by the fact that we observe consumption, production, and prices over 1983-2003 but cross-region shipments over 1990-2003. We use methods developed in Srivastava and Zaatar (1973) and Hwang (1990) to account for the unequal numbers of observations.

5.4 Potential demand

We normalize the potential demand of each county using two exogenous demand predictors that we observe at the county level: the number of construction employees and the number of new residential building permits. 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.[26] The results indicate that potential demand is concentrated in a small number of counties. In 2003, the largest 20 counties account for 90 percent of potential demand, the largest 10 counties account for 65 percent of potential demand, and the largest two counties – Maricopa County and Los Angeles County – together account for nearly 25 percent of potential demand.[27] In the time-series, potential demand more than doubles over 1983-2003, due to greater activity in the construction sector and the onset of the housing bubble.

5.5 The geographic space

Our restricted geographic focus eases the computational burden of the estimation routine. For instance, a national sample would require the computation of more than 300 thousand equilibrium plant-county prices in each time period, for every outer loop iteration. The geographical restriction is valid provided that gross domestic inflows/outflows are insubstantial. The data provide some support. Most directly, the California Letter indicates that more than 98 percent of cement produced in Southern California was shipped within the U.S. Southwest over the period 1990-1999, and more than 99 percent of cement produced in California (both Northern and Southern) was shipped within the U.S. Southwest over the period 2000-2003.[28] We consider outflows from Arizona unlikely because the Minerals Yearbook indicates that consumption routinely exceeds production in that state, and we consider outflows from Nevada unlikely because production capacity is low relative to potential demand. Since *net* domestic inflows/outflows are insubstantial (see Figure 2), these data patterns

[26] The regression of regional portland cement consumption on the demand predictors yields an R^2 of 0.9786, which foreshadows an inelastic estimate of aggregate demand. 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.

[27] The largest five counties are Maricopa County (3,259 thousand metric tonnes), Los Angeles County (3,128 thousand metric tonnes), Clark County (1,962 thousand metric tonnes), Riverside County (1,803 thousand metric tonnes) and San Diego County (1,733 thousand metric tonnes).

[28] Analogous statistics for Northern California over 1990-1999 are unavailable due to data redaction.

suggest that gross inflows are also insubstantial.

5.6 Identification

We use an artificial data experiment to test identification. We draw 40 data sets, each with 21 time periods, using our model and a vector of "true" parameters as the data generating process. We then seek to recover the parameter values with the GMM estimation procedure. The exogenous data includes the plant capacities, the potential demand of counties, the diesel price, the import price, and two cost shifters. We randomly draw capacity and potential demand from the data (with replacement), and we draw the remaining data from normal distributions.[29] We hold plant and county locations fixed to maintain tractability, and rely on the random draws on capacity, potential demand, and diesel prices to create variation in the distances between production capacity and consumers.

Table 3 shows the results of the experiment. Interpretation is complicated somewhat because we use non-linear transformations to constrain the price and distance coefficients below zero, constrain the coefficients on the cost shifters and the over-utilization cost above zero, and constrain the inclusive value and utilization threshold coefficients between zero and one. We defer details on these transformations to Appendix C. As shown, the means of the estimated coefficients are close to (transformed) true parameters. The means of the price and distance coefficients, which are of particular interest, are within 6 percent and 11 percent of the truth, respectively. The root mean-squared errors tend to be between 0.45 and 0.66 – the two exceptions that generate higher mean-squared errors are the import dummy and the over-utilization cost, which appear to be less cleanly identified.

To further build intuition, we explore some of the empirical relationships (graphed in Figure 3) that drive parameter estimates in our application. On the demand side, the price coefficient is primarily identified by the relationship between the consumption and price moments. In panel A, we plot cement prices and the ratio of consumption to potential demand ("market coverage") over the sample period. The two metrics have a weak negative correlation, consistent with downward-sloping but inelastic aggregate demand. The distance coefficient is primarily identified by (1) the cross-region shipments moment, and (2) the relationship between the consumption and production moments. To explore the second source of identification, we plot the gap between production and consumption ("excess production") for each region over the sample period. In many years, excess production is positive in

[29]Specifically, we use the following distributions: diesel price $\sim N(1, 0.28)$, import price $\sim N(50, 9)$, cost shifter 1 $\sim N(60, 15)$, and cost shifter 2 $\sim N(9, 2)$. We redraw data that are below zero. We also redraw data that lead the estimator to nonsensical areas of parameter space.

Table 3: Artificial Data Test for Identification

Variable	Parameter	Truth (θ)	Transformed ($\widetilde{\theta}$)	Mean Est	RMSE
Demand					
Cement Price	β_1	-0.07	-2.66	-2.51	0.66
Miles×Diesel Price	β_2	-25.00	3.22	2.86	0.59
Import Dummy	β_3	-4.00	-4.00	-6.07	1.23
Intercept	β_0	2.00	2.00	1.11	0.51
Inclusive Value	λ	0.09	-2.31	-1.73	0.54
Marginal Costs					
Cost Shifter 1	α_1	0.70	-0.36	-0.88	0.51
Cost Shifter 2	α_2	3.00	1.10	0.54	0.45
Utilization Threshold	ν	0.90	2.19	1.71	0.59
Over-Utilization Cost	γ	300.00	5.70	6.14	1.05

Results of GMM estimation on 40 data sets that are randomly drawn based on the "true" parameters listed. The parameters are transformed prior to estimation to place constraints on the parameter signs/magnitudes (see Appendix C). Mean Est and RMSE are the mean of the estimated (transformed) parameters and the root mean-squared error, respectively.

Southern California and negative elsewhere, consistent with inter-regional trade flows. The magnitude of these implied trade flows helps drive the distance coefficient. Interestingly, the implied trade flows are higher later in the sample, when the diesel fuel is less expensive.

On the supply side, the parameters on the marginal cost shifters are primarily identified by the price moments. In panel C, we plot the coal price, the electricity price, the durable-goods manufacturing wage, and the crushed stone price for California. Coal and electricity prices are highly correlated with the cement price (e.g., see panel A), consistent with a strong influence on marginal costs; inter-regional variation in input prices helps disentangle the two effects. It is less clear that wages and crushed stone prices are positively correlated with cement prices. The utilization parameters are primarily identified by the relationship between the production moments (which determine utilization) and the price moments. In panel D, we plot 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.

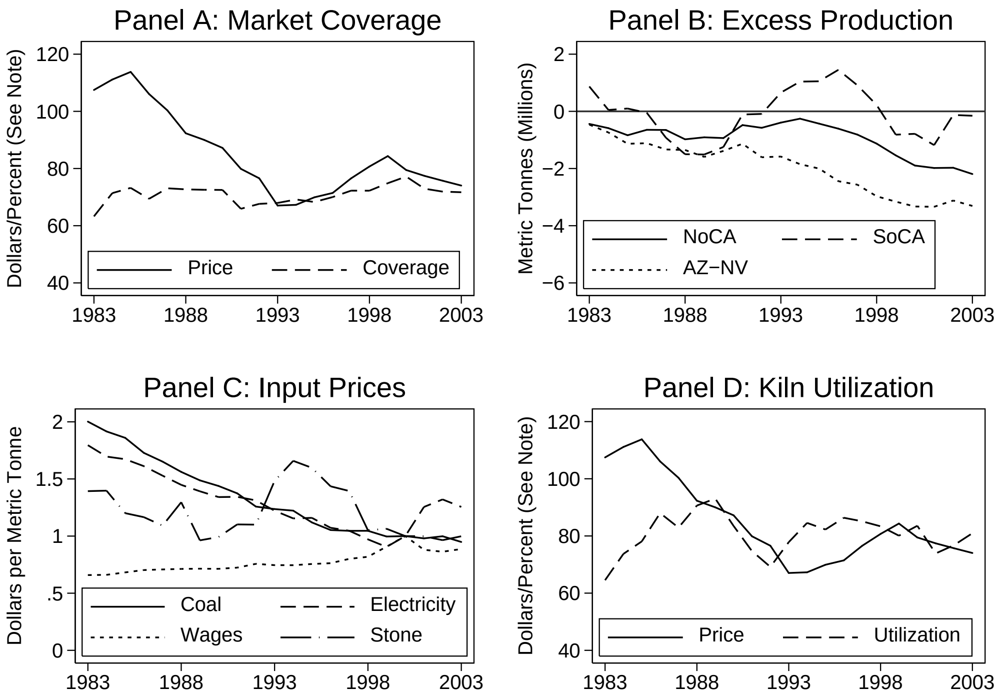

Figure 3: Empirical Relationships in the U.S. Southwest. Panel A plots average cement prices and market coverage. Prices are in dollars per metric tonne and market coverage is defined as the ratio of consumption to potential demand (times 100). Panel B plots excess production in each region, which we define as the gap between between production and consumption. Excess production is in millions of metric tonnes. Panel C plots average coal prices, electricity prices, durable-goods manufacturing wages, and crushed stone prices in California. For comparability, each time-series is converted to an index that equals one in 2000. Panel D plots the average cement price and industry-wide utilization (times 100).

6 Estimation Results

6.1 Specification and fits

We estimate the model with parsimonious specifications of the utility and marginal cost functions. Specifically, the utility specification includes the plant-county price, the "distance" between the plant and county, a dummy for the import option, and an intercept. We proxy distance using a diesel price index interacted with the miles between the plant and the center of the county (in thousands). The marginal cost specification incorporates the five variable inputs identified by EPA (2009). The constant portion of marginal costs includes shifters for the price of coal, the price of electricity, the average wages of durable-goods manufacturing employees, and the price of crushed stone. We let marginal costs increase in production once utilization exceeds some (estimated) threshold value, as written in Equation 2, and we

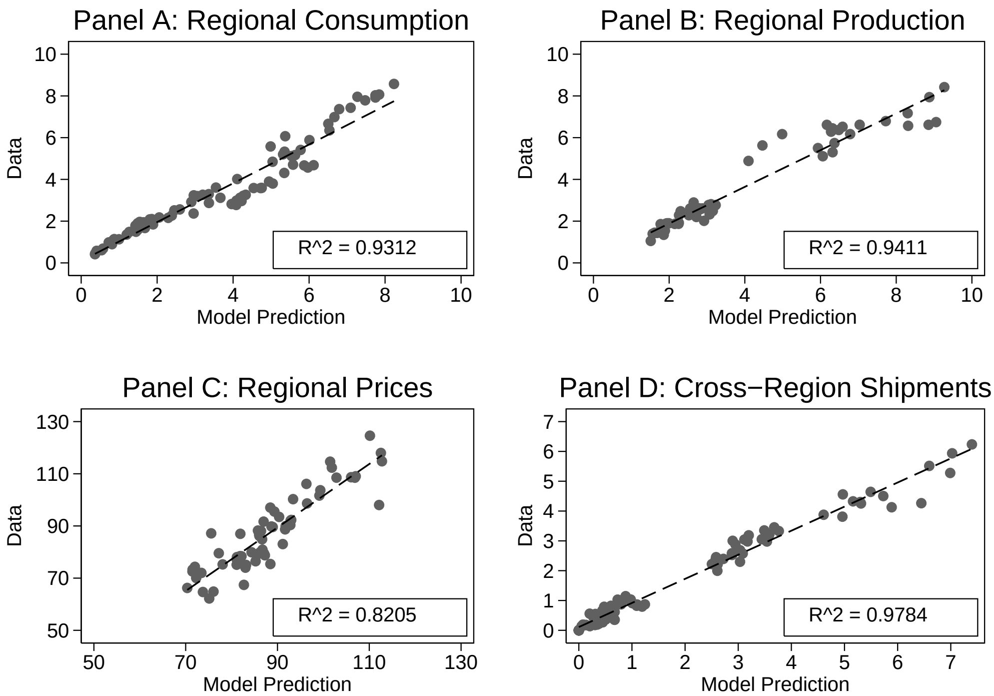

Figure 4: GMM Estimation Fits for Regional Metrics. Consumption, production, and cross-region shipments are in millions of metric tonnes. Prices are constructed as a weighted-average of plants in the region, and are reported as dollars per metric tonne. The lines of best fit and the reported R^2 values are based on univariate OLS regressions.

normalize $\phi = 1.5$ to ensure the theoretical existence of equilibrium.

Before turning to the parameter estimates, we note that this specification produces impressive in-sample and out-of-sample fits. In Figure 4, we plot observed consumption against predicted consumption (panel A), observed production against predicted production (panel B), and observed prices against predicted prices (panel C). The model explains 93 percent of the variation in regional consumption, 94 percent of the variation in regional production, and 82 percent of the variation in regional prices. The model also generates accurate out-of-sample predictions. In panel D, we plot observations on cross-region shipments against the corresponding model predictions. We use only 14 of these observations in the estimation routine – the remaining 82 data points are withheld from the estimation procedure and do not directly influence the estimated parameters. Even so, the model explains 98 percent of the variation in these data.[30]

[30] We provide additional information on the estimation fits in Appendix A.

Table 4: Estimation Results

Variable	Parameter	Estimate	St. Error
Demand			
Cement Price	β_1	-0.087	0.002
Miles×Diesel Price	β_2	-26.42	1.78
Import Dummy	β_3	-3.80	0.06
Intercept	β_0	1.88	0.08
Inclusive Value	λ	0.10	0.004
Marginal Costs			
Coal Price	α_1	0.64	0.05
Electricity Price	α_2	2.28	0.47
Hourly Wages	α_3	0.01	0.04
Crushed Stone Price	α_4	0.29	0.31
Utilization Threshold	ν	0.86	0.01
Over-Utilization Cost	γ	233.91	38.16

GMM estimation results. Estimation exploits variation in regional consumption, production, and average prices over the period 1983-2003, as well as variation in shipments from California to Northern California over the period 1990-2003. The prices of cement, coal, and crushed stone are in dollars per metric tonne. Miles are in thousands. The diesel price is an index that equals one in 2000. The price of electricity is in cents per kilowatt-hour, and hourly wages are in dollars per hour. The marginal cost parameter ϕ is normalized to 1.5, which ensures the theoretical existence of equilibrium. Standard errors are robust to heteroscedasticity and contemporaneous correlations between moments.

6.2 Demand estimates and transportation costs

Table 4 presents the parameter estimates of the GMM procedure. The price and distance coefficients are the two primary objects of interest on the demand side; both are negative and precisely estimated.[31] The ratio of these coefficients identifies consumers' willingness-to-pay for proximity, incorporating transportation costs and any other distance-related costs (e.g., reduced reliability). In the following discussion, we refer to the willingness-to-pay as the transportation cost, although the two concepts may not be strictly equivalent.

First, however, we briefly summarize the implied price elasticities. We estimate that

[31] The other demand parameters take reasonable values and are precisely identified. The negative coefficient on the import dummy may be due to consumer preferences for domestic plants or to the fact that observed import prices do not reflect the full price of imported cement (e.g., the data exclude duties). The inclusive value coefficient suggests that consumer tastes for the different cement providers are highly correlated, inconsistent with the standard (non-nested) logit model.

the aggregate elasticity is -0.16 in the median year. This inelasticity is precisely what one should expect based on economic theory and the fact that portland cement composes only a small fraction of total construction expenses. Indeed, Syverson (2004) makes a similar argument for ready-mix concrete, which accounts for only two percent of total construction expenses according to the 1987 Benchmark Input-Output Tables. The cost share of portland cement (an input to ready-mix concrete) is surely even lower. By contrast, we estimate that the median firm-level elasticity is -5.70, consistent with substantive competition between firms. Finally, we estimate that the domestic elasticity – which captures the responsiveness of domestic demand to domestic prices, holding import prices constant – is -1.11 in the median year. The discrepancy between the aggregate and domestic elasticities is suggestive that imports have a disciplining effect on domestic prices.

Returning to transportation costs, we estimate that consumers pay roughly \$0.30 per tonne mile, given diesel prices at the 2000 level.[32] Given the shipping distances that arise in numerical equilibrium, this translates into an average transportation cost of \$24.61 per metric tonne over the sample period – sufficient to account for 22 percent of total consumer expenditure. Transportation costs of this magnitude have real effects on the industry. We develop two such effects here: (1) transportation costs constrain the distance that cement can be shipped economically; and (2) transportation costs insulate firms from competition and provide some degree of localized market power.

In Figure 5, we plot the estimated distribution of the shipping distances over 1983-2003. We calculate that portland cement is shipped an average of 92 miles, that 75 percent of portland cement is shipped under 110 miles, and that 90 percent is shipped under 175 miles.[33] To better place these numbers in context, we ask the question: "How far would portland cement be shipped if transportation costs were negligible?" We perform a counterfactual simulation in which we normalize the distance coefficient to zero, keeping the other coefficients at their estimated values. The results suggest that, in an average year, portland cement would have been shipped on average 276 miles absent transportation costs. Intriguingly, the ratio of actual miles shipped to this simulated measure provides a unit-free measure that could enable cross-industry comparisons. The ratio in our application is 0.33.

We now develop the empirical evidence regarding localized market power. We start with an illustrative example. Figure 6 shows the prices (Map A) and market shares (Map B) that characterize numerical equilibrium for the Clarksdale plant in 2003, evaluated the

[32]The calculation is simply $\frac{26.42}{0.087}\frac{index}{1000} = 0.3037$, where $index = 1$ in 2000.

[33]The average shipping distance fluctuates between a minimum of 72 miles in 1983 and a maximum of 114 miles in 1998, and is highly correlated with the diesel price index.

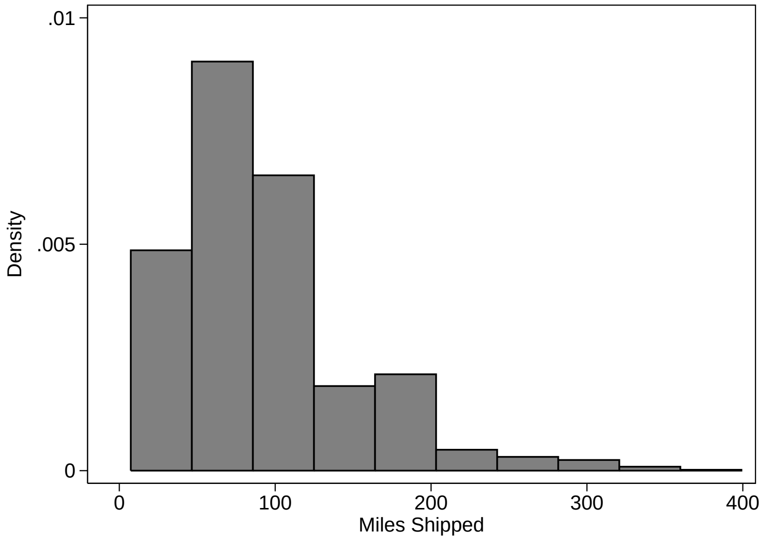

Figure 5: The Estimated Distribution of Miles Shipped over 1983-2003.

optimized coefficient vector. We mark the location of the Clarksdale plant with a star, and mark other plants with circles. The Clarksdale plant captures more than 40 percent of the market in the central and northeastern counties of Arizona. It charges consumers in these counties its highest prices, typically \$80 per metric tonne or more. Both market shares and prices are lower in more distant counties, and in many counties the plant captures less than one percent of demand despite substantial discounts. The locations of competitors may also influence market share and prices, though these effects are difficult to discern on the map.

We explore these relationships more rigorously with regression analysis, based on the prices and market shares that characterize equilibrium at the optimized coefficient vector. We regress price and market share on three independent variables: the distance between the plant and the county, the distance between the county and the nearest other domestic plant, and the estimated marginal cost of the plant. We define distance as miles times the diesel index, and use a log-log specification to ease interpretation. We focus on three data samples, composed of the plant-county pairs with distances of 0-100, 100-200, and 200-300, respectively. Our objective is purely descriptive and the regression coefficients should not be interpreted as consistent estimates of any underlying structural parameters.

Table 5 presents the results, which are consistent with the illustrative example and demonstrate that (1) plants have higher prices and market shares in counties that are closer; and (2) plants have lower prices and market shares in counties that have nearby alternatives. For the closest plants and counties, a 10 percent reduction in distance is associated with prices

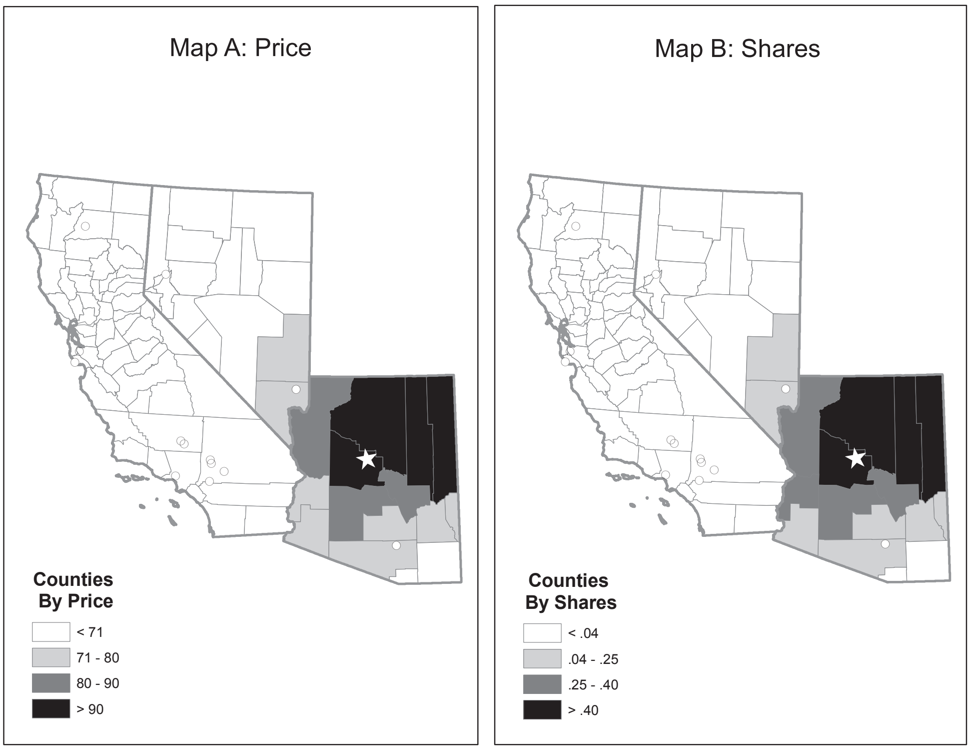

Figure 6: Equilibrium Prices and Market Shares for the Clarksdale Plant in 2003. The Clarksdale plant is marked with a star, and other plants are marked with circles.

and market shares that are 0.9 percent and 14 percent higher, respectively. For the same sample, a 10 percent reduction in the distance separating the county from its the closest alternative is associated with prices and market shares that are 0.7 percent and 11 percent lower, respectively. Interestingly, these price effects attenuate for plants and counties that are somewhat more distant, whereas the market share effects amplify.

6.3 Marginal cost estimates

The marginal costs estimates shown in Table 4 correspond to a marginal cost of $69.40 in the mean plant-year (weighted by production). Of these marginal costs, $60.50 is attributable to costs related to coal, electricity, labor and raw materials, and the remaining $8.90 is attributable to high utilization rates. In Figure 7, we plot the three metrics over 1983-2003, together with the average prices that arise in numerical equilibrium at the optimized parameter vector. The constant portion of marginal costs declines through the sample period, due primarily to cheaper coal and electricity, whereas the utilization portion increases. We estimate that utilization-related expenses account for roughly 25 percent of overall marginal

Table 5: Plant Prices and Market Shares

Dependent Variable:	ln(Price)	ln(Price)	ln(Price)	ln(Share)	ln(Share)	ln(Share)
Distance from Plant:	0-100	100-200	200-300	0-100	100-200	200-300
ln(Distance from Plant)	-0.098*	-0.038*	-0.003	-1.369*	-2.655*	-5.902*
	(0.026)	(0.009)	(0.009)	(0.102)	(0.186)	(0.195)
ln(Distance to Nearest Alternative)	0.071*	0.018*	0.007*	1.073*	0.813*	1.279*
	(0.019)	(0.004)	(0.002)	(0.117)	(0.081)	(0.098)
ln(Marginal Cost)	0.723*	0.835*	0.841*	-1.126*	-2.057*	-3.212*
	(0.054)	(0.014)	(0.013)	(0.304)	(0.645)	(0.874)
R^2	0.7186	0.9209	0.9499	0.5278	0.4735	0.5407
# of Obs.	2,840	5,460	6,088	2,840	5,460	6,088

Results of OLS regression. The units of observation are at the plant-county-year level. The dependent variables are the natural logs of the plant-county specific prices and market shares that characterize numerical equilibrium at the GMM estimates. Distance from Plant is the miles between the plant and county, times a diesel price index that equals one in 2000. Distance to Nearest Alternative is the miles between the county and the nearest other domestic plant, times the diesel price index. Marginal Cost is the marginal cost of the plant implied by the GMM estimates. All regressions include an intercept. Standard errors are robust to heteroscedasticity and correlations among observations from the same plant. Statistical significance at the one percent level is denoted by *.

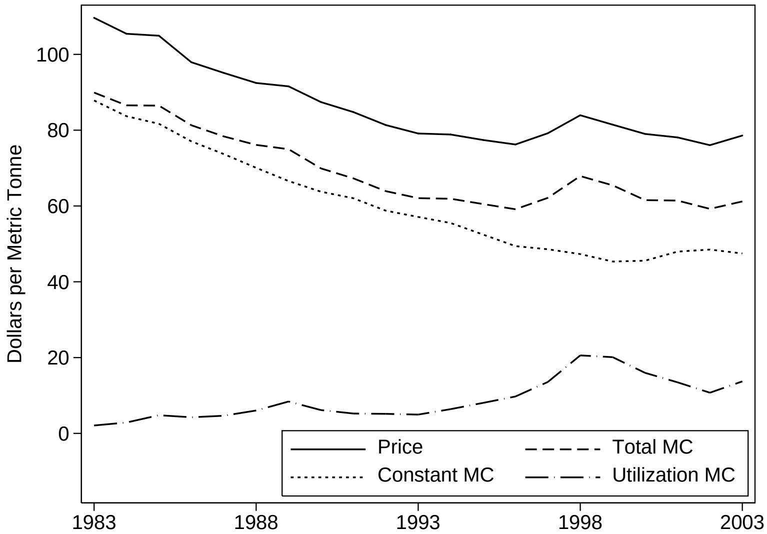

Figure 7: Estimated Marginal Costs and Average Prices.

costs over 1997-2003, during the onset of the housing bubble. Finally, we note that the average markup (i.e., price minus marginal cost) is quite stable through the sample period around its mean of $17.20.

We calculate that the average plant-year observation has variable costs of $51 million by integrating the marginal cost function over the production levels that arise in numerical equilibrium. Virtually all of these variable costs – 98.5 percent – are due to coal, electricity, labor and raw materials, rather than due to high utilization. Thus, although capacity constraints may have substantial affects on marginal costs, the results suggest that their cumulative contribution to plant costs can be minimal. Taking the accounting statistics further, we calculate that the average plant-year has variable revenues of $73 million and that the average gross margin (variable profits over variable revenues) is 0.32. As argued in Ryan (2009), margins of this magnitude may be needed to rationalize entry given the sunk costs associated with plant construction.[34,35]

[34]These gross margins are consistent with publicly-available accounting data. For instance, Lafarge North America – one of the largest domestic producers – reports an average gross margin of 0.33 over 2002-2004.

[35]Fixed costs are well understood to be important for production, as well. The trade journal *Rock Products* reports that high capacity portland cement plants incurred averaged $6.96 in maintenance costs per production tonne in 1993 (Rock-Products (1994)). Evaluated at the production levels that correspond to numerical equilibrium in 1993, this number implies that the average plant would have incurred $5.7 million in maintenance costs relative to variable profits of $17.7 million. The GMM estimation results suggest that the bulk of these maintenance costs are best considered fixed rather than due to high utilization rates. Of course, the static nature of the model precludes more direct inferences about fixed costs.

Finally, we discuss the individual parameter estimates shown in Table 4, each of which deviates somewhat from production data available from the Minerals Yearbooks and EPA (2009). To start, the coal parameter implies that plants burn 0.64 tonnes of coal to produce one tonne of cement, whereas in fact plants burn roughly 0.09 tonnes of coal to produce each tonne of cement. The electricity parameter implies that plants use 228 kilowatt-hours per tonne of cement, whereas the true number is closer to 150. Each tonne of cement requires approximately 0.34 employee-hours yet the parameter on wages is essentially zero. Lastly, the crushed stone coefficient of 0.29 is too small, given that roughly 1.67 tonnes of raw materials (mostly limestone) are used per tonne of cement. We suspect that these discrepancies are due to measurement error in the data.[36]

6.4 A comparison to standard methods

The standard method of structural analysis for homogenous product industries assumes independent markets and Cournot competition. In this section, we contrast some of our results to those generated by the standard method in Ryan (2009), a recent paper that estimates a structural model of the portland cement industry based on data from the Minerals Yearbook and the Plant Information Summary. We focus on two economic concepts – the aggregate elasticity of demand and the consequences of high utilization – for which our model generates distinctly different estimates than the standard method. These discrepancies do not diminish the substantial contribution of Ryan (2009), which embeds the standard method within an innovative dynamic discrete choice game and focuses primarily on the dynamic parameters. Rather, the discrepancies suggest two reasons that our model may sometimes provide more reasonable results than conventional approaches.

First, we estimate the aggregate elasticity of demand to be -0.16 in the median sample year whereas Ryan works with an aggregate elasticity of -2.96, obtained from a constant elasticity demand system. The difference is due to specific specification choices – the constant elasticity demand system produces an aggregate elasticity of -0.15 once housing permits are included as a control.[37] Ryan cannot use the inelastic estimate because, within the context of Cournot competition, it implies firm elasticities that are small and inconsistent with profit

[36] In particular, the coal prices in the data are free-on-board and do not reflect any transportation costs paid by cement plants; cement plants may negotiate individual contracts with electrical utilities that are not reflected in the data; the wages of cement workers need not track the average wages of durable-goods manufacturing employees; and cement plants typically use limestone from a quarry adjacent to the plant, so the crushed stone price may not proxy the cost of limestone acquisition (i.e., the quarry production costs).

[37] See Table 3 in Ryan (2009). We consider the inelastic estimate more plausible because portland cement is a minor cost for most construction projects (see Section 6.2).

maximization. This occurs because the Cournot model restricts each firm elasticity to be linearly related to the aggregate elasticity according the relationship $e_j = e/s_j$, where e_j, e, and s_j denote the firm elasticity, the aggregate elasticity, and the firm market shares, respectively. This critique is fundamental: the standard method can be inappropriate for intermediate goods, such as portland cement, that account for only a fraction for the total production costs of the final good. By contrast, the nested logit demand system divorces the firm elasticities from the aggregate elasticity and, in our case, produces inelastic aggregate demand and elastic firm demand.

Second, the two methods produce vastly different estimates of the marginal cost curve once utilization reaches the threshold level (which both methods place just above 0.85). We estimate that marginal costs increase gradually so that full utilization increases marginal costs by a total of \$12.25 relative to utilization below the threshold. By contrast, Ryan estimates that the slope of the marginal cost curve past the threshold is nearly infinite.[38] We suspect that the difference is data driven. The standard method requires data on firm-level utilization. However, firm production is not available from the publicly-available data, and Ryan imputes utilization as annual capacity divided by annualized daily capacity.

In Figure 8, we plot total production, total annual capacity, and total annualized daily capacity in the U.S. Southwest over 1983-2003, together with total consumption. Both production and consumption are pro-cyclical, and actual utilization (i.e., production over annual capacity) varies substantively and predictably with demand conditions. By contrast, annual capacity simply tracks annualized daily capacity so that Ryan's utilization proxy is uncorrelated with demand conditions. The strength of the relationship between utilization and demand is precisely what identifies the magnitude of utilization costs. Thus, we suspect that the lower data requirements of our model – estimation is feasible when some variables of interest (e.g., firm-level production) are unobserved – may improve economic estimates.

7 An application to competition policy

The model and estimator may prove useful for a variety of policy endeavors. One potential application is merger simulation, an important tool for competition policy. In this subsection, we use counter-factual simulations to evaluate a hypothetical merger between Calmat and Gifford-Hill in 1986. During that year, Calmat and Gifford-Hill operated six plants and accounted for 43 percent of industry capacity in the U.S. Southwest. We calculate the loss

[38]Our coefficient γ is roughly analogous to Ryan's δ_2 coefficient. We estimate γ to be 233.91, while Ryan estimates δ_2 to be 1157×10^7. See Table 4 in Ryan (2009).

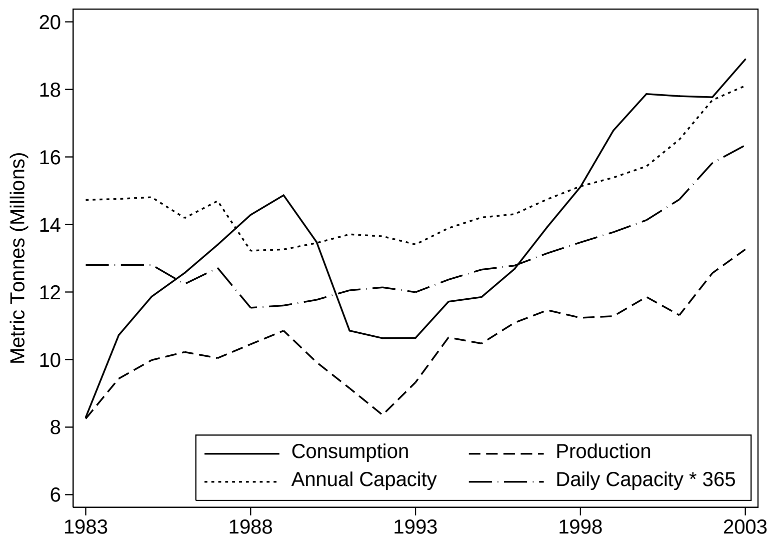

Figure 8: Consumption, Production, and Two Capacity Measures.

of consumer surplus due to the unilateral effects of the merger, map the distribution of harm over the U.S. Southwest, and evaluate six alternative divestiture plans.[39]

Table 6 shows that the merger reduces consumer surplus by $1.40 million in 1986, absent a divestiture. The magnitude of the effect is modest relative to the amount of commerce; by way of comparison, we calculate total pre-merger consumer surplus to be more than $239 million. We refer to the six plants available for divestiture as Calmat1, Calmat2, Calmat3, Gifford-Hill1, Gifford-Hill2, and Gifford-Hill3, respectively. The single-plant divestitures mitigate between 31 percent and 56 percent of the harm. The "optimal" divestiture – that of Gifford-Hill2 – results in consumer harm of only $614 thousand.

We map the distribution of consumer harm over the U.S. Southwest in Figure 9, both for the merger without divestiture (panel A) and under the optimal divestiture plan (panel B). As shown in panel A, the unilateral effects of the merger are concentrated in Southern California and Arizona. Together, Maricopa County and Los Angeles County account for

[39] We follow standard practice to perform the counterfactuals. For each merger simulation, we define a matrix $\boldsymbol{\Omega^{post}(P)}$ using Equation 6 and the post-merger structure of the industry. We compute the equilibrium post-merger price vector as the solution to Equation 7, substituting $\boldsymbol{\Omega^{post}(P)}$ for $\boldsymbol{\Omega(P)}$. Following McFadden (1981) and Small and Rosen (1981), the change in consumer surplus due to the merger is:

$$\Delta CS = \sum_{n=1}^{N} \frac{\ln(1+\exp(\beta_0+\lambda I_{nt}^{pre})) - \ln(1+\exp(\beta_0+\lambda I_{nt}^{post}))}{\beta_1} M_n,$$

where I_n^{pre} is the inclusive value of the inside goods calculated using equilibrium pre-merger prices, I_n^{post} is the inclusive value calculated using equilibrium post-merger prices, and β_1 is the price coefficient.

Table 6: Divestitures and Consumer Surplus

		Required Divestiture					
	None	Calmat1	Calmat2	Calmat3	Gifford-Hill1	Gifford-Hill2	Gifford-Hill3
Δ Surplus	-1,397	-964	-618	-797	-827	-614	-891
% Mitigated	·	31%	56%	43%	41%	56%	36%

Results of counterfactual simulations. Δ Surplus is the change in consumer surplus due to a hypothetical merger between Calmat and Gifford-Hill in 1986, and is reported in thousands of 2000 dollars. % Mitigated is calculated relative to the change in consumer surplus that occurs when no plant is divested. We consider six single-plant divestiture plans, and refer to the different plants as Calmat1, Calmat2, Calmat3, Gifford-Hill1, Gifford-Hill2, and Gifford-Hill3.

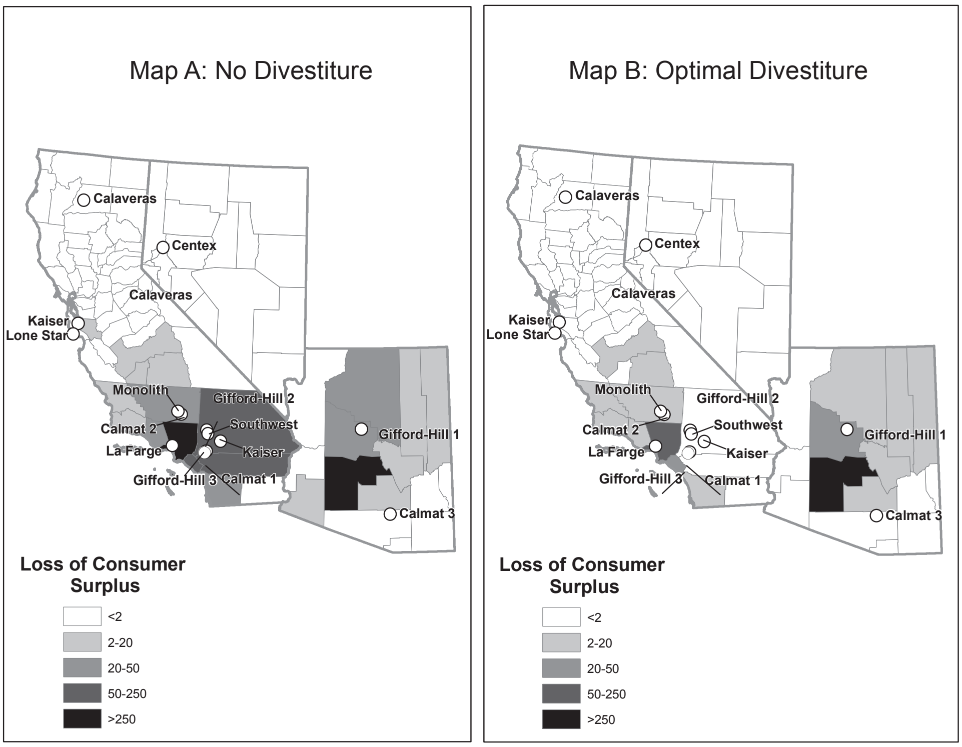

Figure 9: Loss of Consumer Surplus Due to a Hypothetical Merger between Calmat and Gifford-Hill

60 percent of consumer harm, and more than 90 percent of consumer harm occurs only 10 counties. The best single-plant divestiture mitigates consumer harm in Southern California but does little to reduce harm in Maricopa County (see panel B). The results of an additional counterfactual exercises, in which we also divest one of the Arizona plants, suggest that a two-plant divestiture can mitigate this harm as well (results not shown).

8 Conclusion

We develop a structural model of competition among spatially differentiated firms. The model accounts for transportation costs in a realistic and tractable manner. We estimate the model with relatively disaggregated data and recover the underlying structural parameters. We argue that the model and estimator together provide an appealing framework with which to evaluate competition in industries characterized by transportation costs and relatively homogenous products. We apply the model and estimator to the portland cement industry and demonstrate that (1) the framework explains the salient features of competition, (2) the framework provides novel insights regarding transportation costs and spatial differentiation, and (3) the framework could inform merger analysis and other competition policy endeavors. Although the model is static, it could be utilized to define payoffs in more dynamic settings. Such extensions could examine a number of research topics – such as entry deterrence and product differentiation – that have been emphasized in the theoretical literature of industrial organization since at least Hotelling (1929).

References

Aguirregabiria, V. and G. Vicentini (2006). Dynamic spatial competition between multi-store firms. *Mimeo*.

Berry, S., J. Levinsohn, and A. Pakes (1995, July). Automobile prices in market equilibrium. *Econometrica*, 841–890.

Cardell, S. N. (1997). Variance components structures for the extreme value and logistic distributions with applications to models of heterogeneity. *Journal of Economic Theory 13*, 185–213.

Collard-Wexler, A. (2009). Productivity dispersion and plant selection in the ready-mix concrete industry. *Mimeo*.

d'Aspremont, C., J. Gabszewicz, and J.-F. Thisse (1979). On Hotelling's "Stability in Competion". *Econometrica 47*, 1145–1150.

Davis, P. (2006). Spatial competition in retail markets: Movie theaters. *The RAND Journal of Economics 37*, 964–982.

Dunne, T., S. Klimek, and J. A. Schmitz (2009). Does foreign competition spur productivity? Evidence from post-WWII U.S. cement manufacturing. *Federal Reserve Bank of Minneapolis Staff Report*.

Economides, N. (1989). Stability in competition. *Journal of Economic Theory 47*, 178–194.

EPA (2009). *Regulatory Impact Analysis: National Emission Standards for Hazardous Air Pollutants from the Portland Cement Manufacturing Industry.* Prepared by RTI International.

Greene, W. H. (2003). *Econometric Analysis 5th Ed.* Prentice-Hall Upper Saddle River NJ.

Hansen, L. (1982). Large sample properties of generalized method of moments estimators. *Econometrica 50*, 1029–1054.

Hotelling, H. (1929). Stability in competition. *Economic Journal 39*, 41–57.

Hwang, H.-s. (1990). Estimation of a linear SUR model with unequal numbers of observations. *Review of Economics and Statistics 72*.

Jans, I. and D. Rosenbaum (1997, May). Multimarket contact and pricing: Evidence from the U.S. cement industry. *International Journal of Industrial Organization 15*, 391–412.

La Cruz, W., J. Martínez, and M. Raydan (2006). Spectral residual method without gradient information for solving large-scale nonlinear systems of equations. *Mathematics of Computation 75*, 1429–1448.

Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly Journal of Mathematics 2*, 164–168.

Marquardt, D. (1963). An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal of Applied Mathematics 11*, 431–331.

McBride, M. (1983, December). Spatial competition and vertical integration: Cement - concrete revisited. *American Economic Review 73*, 1011–1022.

McFadden, D. (1981). *Econometric Models of Probabilistic Choice*. MIT Press. Structural Analysis of Discrete Data, C.F. Manski and D. McFadden (eds.).

McManus, B. (Forthcoming 2009). Nonlinear pricing in an oligopoly market: the case of specialty coffee. *RAND Journal of Economics*.

Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. *Handbook of Econometrics 4*.

Pesendorfer, M. (2003). Horizontal mergers in the paper industry. *RAND Journal of Economics 34*, 495–515.

Rock-Products (1994). Cement plant operating cost study. Overland Park, KS: Intertic Publishing.

Rosenbaum, D. and S. Reading (1988). Market structure and import share: A regional market analysis. *Southern Economic Journal 54*(3), 694–700.

Rosenbaum, D. I. (1994). Efficiency v. collusion: evidence cast in cement. *Review of Industrial Organization 9*, 379–392.

Ryan, S. (2009). The costs of environmental regulation in a concentrated industry. Mimeo.

Salop, S. (1979). Monopolistic competition with outside goods. *Bell Journal of Economics 10*, 141–156.

Salvo, A. (2008). Inferring market power under the threat of entry: The case of the Brazilian cement industry. *Mimeo*.

Salvo, A. (2010). Trade flows in a spatial oligolopy: Gravity fits well, but what does it explain? *Canadian Journal of Economics 43*, 63–96.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics 37*, 619–640.

Small, K. and H. Rosen (1981). Applied welfare economics with discrete choice methods. *Econometrica 49*, 105–130.

Srivastava, J. and M. K. Zaatar (1973). A Monte Carlo comparison of four estimators of the dispersion matrix of a bivariate normal population, using incomplete data. *Journal of the American Statistical Association 68*, 180–183.

Syverson, C. (2004, December). Market structure and productivity: A concrete example. *Journal of Political Economy 112*, 1181–1222.

Syverson, C. and A. Hortaçsu (2007). Cementing relationships: Vertical integration, foreclosure, productivity, and prices. *Journal of Political Economy 115*, 250–301.

Thisse, J.-F. and X. Vives (1988). On the strategic choice of spatial price policy. *American Economic Review 78*, 122–137.

Thomadsen, R. (2005). The effect of ownership structure on prices in geographically differentiated industries. *RAND Journal of Economics 36*, 908–929.

Van Oss, H. G. and A. Padovani (2002). Cement manufacture and the environment, Part I: Chemistry and technology. *Journal of Industrial Ecology 6*, 89–105.

Vogel, J. (2008, June). Spatial competition with heterogeneous firms. *Journal of Political Economy 116*(3), 423–466.

A Regression fits

In this appendix, we develop further the quality of the estimation fits. We focus on the ability of the model to predict the inter-temporal variation that exist in the data.

Figure 10 aggregates the data and the model predictions across regions, and plots the resulting time-series. Panel A shows consumption, panel B shows production, panel C shows imports, and panel D shows average prices (imports are defined as production minus consumption). In each case, the model predictions mimic the inter-temporal patterns observed in the data. Univariate regressions of the data on the prediction explain 96 percent of the variation total consumption, 75 percent of the variation in total production, 76 percent of the variation in imports, and 91 percent of the variation in average prices.[40]

Figure 11 provides analogous time-series fits for eight of the time-series for cross-region shipments. The data in panel A pertain to shipments from California (both Northern and Southern) to Northern California, and are used in estimation. The data in the remaining panels are excluded from the estimation procedure, so the corresponding fits are out-of-sample. As shown, the model predictions are close to the data in each panel and tend to track the variation well (when variation exists).

B The uniqueness of equilibrium

The estimation procedure rests on uniqueness of equilibrium at each candidate parameter vector. The results of a Monte Carlo experiment suggest the assumption holds, at least in our application. We consider 300 parameter vectors for each of the 21 years in the sample, for a total of 6,300 candidate parameter vectors. For each $\theta_i \in \boldsymbol{\theta}$, we draw from the distribution $N(\widehat{\mu}_i, \widehat{\sigma}_i^2)$, where $\widehat{\mu}_i$ and $\widehat{\sigma}_i$ are the coefficient and standard error, respectively, reported in Table 4. We then compute the numerical equilibrium for each parameter vector, using eleven different starting vectors. We define the elements of the starting vectors to be $P_{jnt} = \phi \overline{P_t}$, where $\overline{P_t}$ is the average price of portland cement and $\phi = 0.5, 0.6, \ldots, 1.4, 1.5$. Thus, we start the equation solver at initial prices that sometimes understate and sometimes overstate the average prices in the data. The experiment produces eleven equilibrium price vectors for each parameter vector. We calculate the standard deviation of each price element across the eleven observations. Thus, we would calculate 1,260 standard deviations for a typical

[40]The model does not fully capture the fall in average prices over the 1980s and early 1990s. One possible explanation is that the model, as specified, does not incorporate potential changes to total factor productivity. Dunne, Klimek, and Schmitz (2009) review the evidence regarding productivity and argue that the gradual elimination of onerous clauses from labor contracts improved productivity in the 1980s.

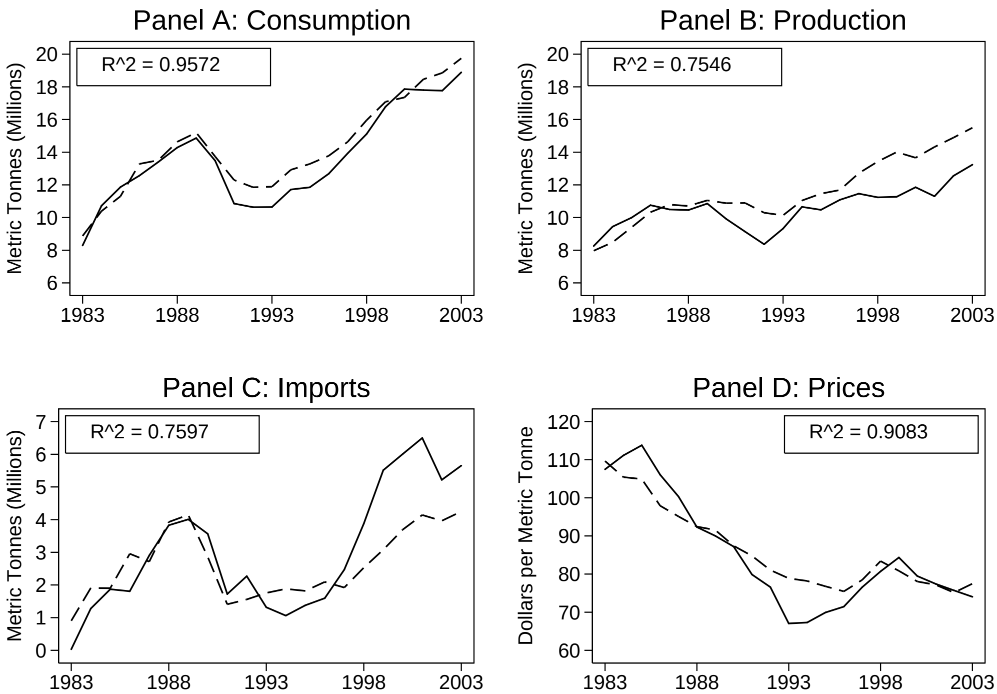

Figure 10: GMM Estimation Fits for Aggregate Metrics. The solid lines plot data and the dashed lines plot predictions. Consumption, production, and imports are in millions of metric tonnes. Imports are defined as production minus consumption. Prices are constructed as a weighted-average of the plant-county prices and are reported in dollars per metric tonne. The R^2 values are calculated from univariate regressions of the observed metric on the predicted metric.

equilibrium price vector with 1,260 plant-county elements. The experiment provides support for the uniqueness condition if these standard deviations are small.[41] This proves to be the case. In fact, the maximum maximum standard deviation is *zero*, considering all prices and draws, so the experiment finds no evidence of multiple equilibria.

C Estimation details

We minimize the objective function using the Levenberg-Marquardt algorithm (Levenberg (1944), Marquardt (1963)), which interpolates between the Gauss-Newton algorithm and the method of gradient descent. We find that the Levenberg-Marquardt algorithm outperforms simplex methods such as simulated annealing and the Nelder-Mead algorithm, as well as

[41]The equal-solver computes numerical equilibria for 90.3 percent of the candidate vectors. See Appendix C for a discussion of non-convergence in the inner-loop.

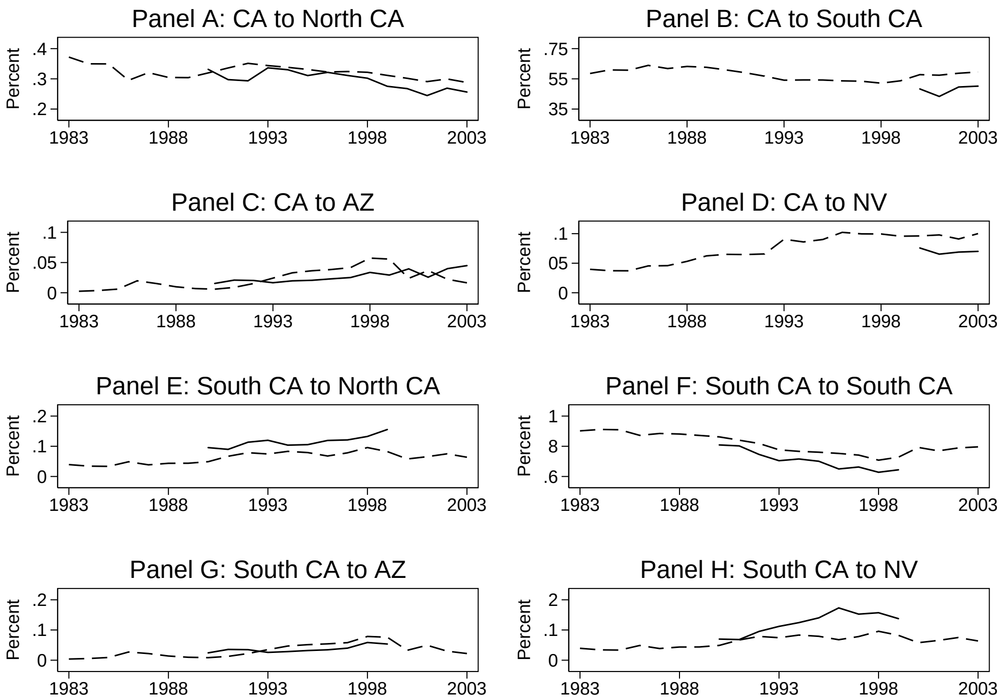

Figure 11: GMM Estimation Fits for Cross-Region Shipments. The solid lines plot data and the dashed lines plot predictions. Shipments are expressed as a percentage of production in California (panels A-D) or Southern California (panels E-H).

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 compute numerical equilibrium using Fortran code that builds on the source code of the dfsane function in R. The dfsane function implements the nonlinear equation solver developed in La Cruz, Martínez, and Raydan (2006) and is downloadable as part of the BB package. We find 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 computation of equilibrium for each time period can be parallelized, which further speeds the inner loop calculations. 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 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. In practice, the equation solver occasionally fails to compute a numerical equilibrium at the specified tolerance level (1e-13) within the specified maximum number of iterations (600). The candidate parameter vectors that generate non-convergence in the inner loop 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 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 over-utilization 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. As examples, we estimate the price coefficient using $\widetilde{\beta_1} = \log(-\beta_1)$ in the GMM procedure, and we estimate the inclusive value coefficient using $\widetilde{\lambda} = \log\left(\frac{\lambda}{1-\lambda}\right)$. We calculate standard errors with the delta method.

D Data details

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.

The Minerals Yearbook reports the total production and average price of plants in the "Nevada-Arizona-New Mexico" region over 1983-1991, and in the "Arizona-New Mexico" region over 1992-2003. We scale the USGS production data downward, proportional to plant capacity, to remove for the influence of the single New Mexico plant. Since the two plants in Arizona account for 89 percent of kiln capacity in Arizona and New Mexico in 2003, we scale production by 0.89.

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. Since the Riverside plant accounts for 7 percent of grinding capacity in Southern California in 1988, so 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 PCA 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.

The data on coal and electricity 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 (1) calculating the average discrepancy between each state's price and the national price over 1990-2000, and (2) adjusting the national-level data upward or downward, in line with the relevant average discrepancy.