

C a S a a D ff a d F : A
E ca M d a A ca C

Nathan H. Miller
Department of Justice

Matthew Osborne
Department of Commerce

September 2010

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 an estimator for models of spatial differentiation and spacial price discrimination that recovers the underlying structural parameters using only aggregate data. We provide conditions under which the estimates are consistent and asymptotically normal. We apply the estimator to the portland cement industry. The estimation tests, both

1 Introduction

In many industries, firms are geographically differentiated and transportation is costly. Yet few empirical studies estimate structural models of spatial differentiation. We attribute this dearth of research to a simple *data availability problem*: the most straight-forward way to identify the degree of spatial differentiation { or, equivalently, the magnitude of transportation costs { is to measure how firms' market shares differ between nearby and distant consumers. But this requires data on the geographic distributions of the market shares. These data are difficult to attain and, indeed, we are unaware of any study that exploits variation in market shares over geographic space.

The data availability problem is only exacerbated for industries characterized by spatial price discrimination because it becomes necessary to account for the geographic distributions of the prices, as well. While three recent studies apply econometric techniques to sidestep the data availability problem in non-discriminatory settings (Thomadsen (2005), Davis (2006),

normal. We also conduct an empirical application and demonstrate that (1) estimation is

markets.⁴ These assumptions preclude inference regarding spatial differentiation because the transportation cost cannot be estimated structurally. Further, markets tend to be delineated based on political borders of questionable economic significance such as state or county lines. Yet this approach has been employed routinely to study of industries characterized by high transportation costs, including ready-mix concrete (e.g., Syverson (2004), Syverson and Hortacsu (2007), Collard-Wexler (2009)), portland cement (e.g., Salvo (2008), Ryan (2009)), and paper (e.g., Pesendorfer (2003)).⁵

In the empirical application, we examine the portland cement industry in the U.S. Southwest over the period 1983-2003. The available data include average prices, production,

the geographic distribution of market shares at each candidate parameter vector. (These market shares have convenient analytical solutions given the assumed logit demand function.) The estimation procedure then selects the parameters that bring the implied equilibrium firm-level prices close to the data. By contrast, Davis (2006) and McManus (2009) exploit variation in firm-level prices and sales. They derive predicted sales in a number of

3 The Model of Price Competition

3.1 The geographic space

We define the relevant geographic *space* to be a compact, connected set \mathbb{C} in the Euclidean space \mathbb{R}^2 . We take as given that J plants compete in the space, and assume that each plant is endowed with a fixed location defined by the geographic coordinates $\{z_1, z_2, \dots, z_J\}$, where $z \in \mathbb{C}$. We further take as given that a continuum of consumers spans the space, and assume that each consumer has unit demand and a fixed location $w \in \mathbb{C}$. The absolute measure $\mu(w)$ characterizes the geographic distribution of consumers and we define $M = \int_{\mathbb{C}} \mu(w) dw$

quantity produced by plant j , and $c(Q; \mathbf{w}; 0$

Figure 1: A Geographic Space.

of \mathbb{C} builds on the premises that (1) consumers in each area \mathbb{C}_n select among all J plants, and (2) demand in area \mathbb{C}_n is unaffected by mill prices in area \mathbb{C}_m for $n \neq m$.

We now rearrange and stack the first-order conditions:

$$\mathbf{f}(\mathbf{p}; \boldsymbol{\theta}; \mathbf{y}_0) = \mathbf{p} - \mathbf{c}(\mathbf{Q}(\mathbf{p}; \boldsymbol{\theta}; \mathbf{y}_0); \boldsymbol{\theta}; \mathbf{y}_0) + \mathbf{1}(\mathbf{p}; \boldsymbol{\theta}; \mathbf{y}_0) \mathbf{q}(\mathbf{p}; \boldsymbol{\theta}; \mathbf{y}_0) = \mathbf{0}. \quad (6)$$

A vector of prices that solves this system of equations is a spatial Bertrand-Nash equilibrium. We define a mapping $\mathbf{H}(\boldsymbol{\theta}; \mathbf{y}_0) : \mathbb{R}^K \rightarrow \mathbb{R}^{JN}$ that matches the parameters of the model to spatial Bertrand-Nash equilibrium given the exogenous data. Formally, the mapping is defined by the equivalence $\mathbf{f}(\mathbf{H}(\boldsymbol{\theta}; \mathbf{y}_0); \boldsymbol{\theta}; \mathbf{y}_0) = \mathbf{0}$.

3.4 Discussion

We offer three comments to help build intuition on the economics of the model. First, spatial price discrimination is at the core of the firm's pricing problem: firms charge higher mill prices to nearby consumers and to consumers for whom the firm's competitors are more distant. However, aside from price discrimination, the firm's pricing problem follows standard intuition. A firm that contemplates a higher mill price from one of its plants to a given area must evaluate (1) the tradeoff between lost sales to marginal consumers and greater revenue from inframarginal consumers; and (2) whether the firm would recapture lost sales with its other plants. If marginal costs are not constant, then the firm must also evaluate how the lost sales would affect the plant's competitiveness in other areas.

Second, the areas $\mathbb{C}_1; \mathbb{C}_2; \dots; \mathbb{C}_N$ are best interpreted as determining the extent which firms engage in spatial price discrimination. Finer partitions of the geographic space correlate with more sophisticated discrimination, and if only a single area exists (i.e., $N = 1$) then firms do not discriminate. *The areas have no economic significance aside from these implications for spatial price discrimination.* Since every plant competes in every area, the partition of the geographic space into distinct areas does not artificially limit competition and is not analogous to a "market delineation" assumption under which plants compete only within prescribed geographic boundaries.

Finally, the indirect utility specification of equation (2) implies that plants are differentiated by both location and idiosyncratic preferences shocks. In the special case of degenerate preference shock distribution, the model collapses to a "pure characteristics model" along the lines of the original Hotelling (1929) formulation. Although the estimation strategy we outline below accommodates the pure characteristics model on a theoretical level, we suspect

already placed on the model. As long as the distributions of \mathbf{z} are known, or reasonable approximations can be made, compute demand can be computed given the relevant prices and the mean distances between plants and areas. We formalize this in Assumption A2.

Assumption A2: *The econometrician knows the distributions of \mathbf{z} .*

Integrating over this distribution yields an equilibrium mapping $\mathbf{H}(\mathbf{z}_0; \mathbf{z})$ that depends on mean distances and the demand and cost shifters. We assume that the price data are generated by the following process:

$$\mathbf{p}^e = \mathbf{H}(\mathbf{z}_0; \mathbf{z}) + \mathbf{v}; \quad (8)$$

where \mathbf{p}^e is a vector of length JN , and \mathbf{v} is a vector of length JN with $\mathbf{v} \sim N(\mathbf{0}, \Sigma)$.

reasonable to further assume that the sampling error is independent of the "right-hand-side" data \mathbf{p} . This simplifies the construction of the estimator, and we impose the additional assumption here:

Assumption A4⁰: *The sampling error is mean zero conditional on \mathbf{p} :*

$$E[\mathbf{p} - \mathbf{S}(\mathbf{H}(\theta_0; \mathbf{p})) | \mathbf{p}] = \mathbf{0}:$$

A4⁰ enables estimation with multiple equation nonlinear least squares, which is equivalent to GMM with the optimal instruments

$$\mathbf{Z} = \frac{\partial \mathbf{S}(\mathbf{H}(\theta_0; \mathbf{p}))}{\partial \theta_0} \mathbf{V}_0^{-1/2}(\theta_0); \quad (10)$$

where $\mathbf{V}_0(\theta_0) = E[\mathbf{S}(\theta_0; \mathbf{p}) \mathbf{S}(\theta_0; \mathbf{p})']$ is the variance matrix of the aggregated error terms. Thus, the sample moment equations that correspond to A4⁰ are

$$\frac{1}{T} \sum_{j=1}^T \frac{\partial \mathbf{S}(\mathbf{H}(\theta; \mathbf{p}))}{\partial \theta} \mathbf{C}^{-1}(\mathbf{p} - \mathbf{S}(\mathbf{H}(\theta; \mathbf{p}))); \quad (11)$$

where \mathbf{C} is some consistent estimate of $\mathbf{V}_0(\theta_0)$ and θ is a candidate parameter vector defined within the compact subspace Θ .

We come now to the central methodological contribution of the paper. Estimation based on the sample moments of equation (11) requires knowledge of equilibrium prices at the plant-area level (i.e., $\mathbf{H}(\theta; \mathbf{p})$). Yet the data generating process provides only prices that are aggregated and measured with error. The solution to this dilemma lies in numerical approximations to equilibrium. Conceptually, it is possible to *compute* the equilibrium price vector for any number of candidate parameter vectors, and then identify the candidate parameter vector that minimizes the "distance" between the aggregated equilibrium price vectors and the data. The power of modern computers makes this procedure feasible given a convenient distribution of the composite error term (see equation (7)). In our application, we are typically able to numerically compute a vector, call it $\tilde{\mathbf{H}}(\theta; \mathbf{p})$, that satisfies the first-order conditions of equation (6) to computer precision in a matter of seconds.

The GMM estimate that utilizes these numerical approximations is:

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T [\mathbf{p}_t - \mathbf{S}(\tilde{\mathbf{H}}(\theta; \mathbf{y}_t))]^0 \mathbf{C}^{-1} [\mathbf{p}_t - \mathbf{S}(\tilde{\mathbf{H}}(\theta; \mathbf{y}_t))]: \quad (12)$$

We think it is intuitive to think of the estimation procedure as combining an outer loop and an inner loop. In the outer loop, the objective function is minimized over the parameter space, whereas in the inner loop equilibrium is computed numerically for each candidate parameter vector considered. This structure makes our estimator broadly analogous to other estimators developed for discrete static games (e.g., Bajari, Hong, and Ryan (2008)), non-strategic dynamic games (e.g., Rust (1987)) and certain strategic dynamic games (e.g., Goettler and Gordon (2009), and Gallant, Hong, and Khwaja (2010)), in the sense that each requires the repeated computation of equilibrium.

4.2 Asymptotic properties

The asymptotic properties of the GMM estimator are unclear without further assumptions, which we develop now:

Assumption A5: *A unique Bertrand-Nash equilibrium exists, and the prices that support it are strictly positive. Formally, for any $\epsilon > 0$ there exists a vector $\mathbf{p}_1 \in \mathbb{R}_+^{JN}$ such that $\mathbf{f}(\mathbf{p}_1; \theta; \mathbf{y}) = \mathbf{0}$. Further, $\mathbf{f}(\mathbf{p}_1; \theta; \mathbf{y}) = \mathbf{f}(\mathbf{p}_2; \theta; \mathbf{y}) = \mathbf{0} \iff \mathbf{p}_1 = \mathbf{p}_2$.*

A5 ensures that the GMM objective function is well-behaved.¹¹ We suspect that uniqueness alone may suffice if, for instance, the econometrician can compute multiple equilibria and select the equilibrium closest to the data (e.g., as in Bisin, Moro, and Topa (2010)). We defer the evaluation of such possibilities to further research. The following lemma clarifies that, given the assumptions of the model, small changes to the parameter vector do not produce large jumps in the objective function:

Lemma 1: *The function $\mathbf{S}(\mathbf{H}(\theta; \mathbf{y}))$ is continuously differentiable in θ and \mathbf{y} for $\theta \in \Theta$, where \mathbf{y} is the vector representation of \mathbf{y} .*

Proof. See appendix A. □

¹¹Recent theoretical contributions demonstrate that A5 holds for two special cases of our model: nested logit demand, convex marginal costs, and single-plant firms (Mizuno 2003), and logit demand, sufficiently increasing marginal costs, and multi-plant firms (Konovalov and Sandor 2010). The assumption is not satisfied generally (e.g., Caplin and Nalebu (1991)).

Assumption A6: The parameter vector θ_0 is globally identified in \mathcal{S} . Formally, $E[\mathbf{p}' \mathbf{S}(\mathbf{H}(\theta; \mathcal{S}))] = \mathbf{0}$ $\mathcal{S} = \theta_0$.

A6 could be violated even if parameters of the model would be globally identified given disaggregate data (i.e., even if $E[\mathbf{p}' \mathbf{H}(\theta; \mathcal{S})] = \mathbf{0}$ $\mathcal{S} = \theta_0$). Such a scenario may be more likely when aggregation is particularly coarse. Empirically, it may be possible to evaluate (imperfectly) the potential for this sort of aggregation problem using artificial data experiments, and we develop one such test in our application.

The asymptotic properties of the GMM estimator follow directly from A1-A6 and the other assumptions placed on the data generating process:

Theorem 1: Under A1-A6 and certain regularity conditions enumerated in the appendix,

- i) $\hat{\theta} \xrightarrow{P} \theta_0$ and
- ii) $\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{D} N(\mathbf{0}; \mathbf{V})$;

where $\mathbf{V} = (\mathbf{G}_0' \mathbf{C}_0 \mathbf{G}_0)^{-1} \mathbf{G}_0' \mathbf{C}_0 \mathbf{C}_0 \mathbf{G}_0 (\mathbf{G}_0' \mathbf{C}_0 \mathbf{G}_0)^{-1}$ and $\mathbf{G}_0 = E[\partial \mathbf{S}(\mathbf{H}(\theta; \mathcal{S})) / \partial \theta]$.

Proof. See appendix A. □

4.3 Incorporating non-price data

The estimation strategy can be extended to exploit variation in other endogenous data, such as observations on production or consumption, that are often available to the econometrician. We focus on production data for expositional brevity; the other extensions are analogous. We assume the data are generated by:

$$\mathbf{q}^\epsilon = \mathbf{q}(\mathbf{H}(\theta_0; \mathcal{S}); \mathcal{S}; \epsilon) + \epsilon; \quad (13)$$

where \mathbf{q}^ϵ is a vector of length JN , and ϵ is a vector of unobserved sampling errors. We define a linear function $\mathbf{R}: \mathbb{R}^{JN} \rightarrow \mathbb{R}^{L^*}$ that maps the plant-area quantities to the aggregate production vector, which we denote as \mathbf{q} . We assume that the aggregate sampling error is mean zero conditional on the exogenous data:

$$E[\mathbf{q} - \mathbf{R}(\mathbf{q}(\mathbf{H}(\theta_0; \mathcal{S}); \mathcal{S}; \epsilon))] = \mathbf{0}; \quad (14)$$

The GMM estimate that incorporates these data is:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{T} \sum$$

Some details of the production process motivate the marginal cost specification we introduce below. Cement plants are typically adjacent to a limestone quarry. The limestone is fed into coal-fired rotary kilns that reach peak temperatures of 1400-1450 Celsius. The output of the kilns (clinker) is cooled, mixed with a small amount of gypsum, and ground in electricity-powered mills to form portland cement. Kilns operate at peak capacity with the exception of an annual maintenance period. When demand is particularly strong, managers sometimes forego maintenance at the risk of breakdowns and kiln damage. Consistent with these stylized facts, a recent report prepared for the Environmental Protection Agency identifies the main variable input costs of production: raw materials, coal, electricity, labor, and kiln maintenance (EPA (2009)).

5.2 The geographic space

We focus on California, Arizona, and Nevada over the period 1983-2003. We refer to these three states as the "U.S. Southwest" for expositional convenience. Figure 2 maps the geographic configuration of the industry in the U.S. Southwest circa 2003. Most plants are located along an interstate highway, nearby one or more population centers. Some firms own multiple plants but ownership 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 figure also plots the four customs offices through which foreign imports enter the region (San Francisco, Los Angeles, San Diego, and Nogales). Most cement imported into the region is produced by large,

Figure 2: Portland Cement Production Capacity in the U.S. Southwest circa 2003.

imports. The similarity of the two imports measures we plot in Figure 3 { actual foreign imports and consumption minus production ("apparent imports") } reveals that net trade flows between the U.S. Southwest and other domestic regions are negligible. Other statistics published by the USGS are strongly suggestive that gross trade flows are also negligible.




Figure 3: 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.

Census response rates are typically well over 90 percent, and the USGS staff imputes missing values for non-respondents based on historical and cross-sectional information.¹⁶ The USGS aggregates the census data to the "regional" level before their publication in the Minerals Yearbook in order to protect the confidentiality of survey respondents. We observe the following endogenous data:

Average mill prices (weighted by production) charged by plants in each of three regions: Northern California, Southern California, and a single Arizona-Nevada region.

Total production by plants in the same three regions.

Consumption in each of four regions: Northern California, Southern California, Arizona, and Nevada.

We also rely on the Minerals Yearbook for information on the price and quantity of portland cement that is imported into the U.S. Southwest.

We make use of more limited data on cross-region shipments from the California Letter, a second annual publication of the USGS. The level of aggregation varies over the

¹⁶The quality of the census has long generated interest among researchers. Other academic studies that feature USGS data include McBride (1983), Pniant(t)-1.361SBJT/F126.9i(1983).

sample period, some data are redacted to protect sensitive information, and no information is available before 1990. For instance, we observe shipments from producers in California (Northern and Southern) to consumers in Northern California over 1990-2003, but shipments from California to Nevada only over 2000-2003. There are 96 data points in total.

The Plant Information Survey (PIS), an annual publication of the Portland Cement Association, provides the geographic location of each portland cement plant as well as the

We augment the theoretical model by letting domestic plants compete against a competitive fringe of foreign importers, which we denote as "plant" $J + 1$. We place the fringe in geographic space at the four customs offices of the U.S. Southwest. Consumers pay the door-to-door cost of transportation from these customs offices. We rule out spatial price discrimination on the part of the fringe, consistent with perfect competition among importers, and assume that the import price is set exogenously (e.g., based on the marginal costs of the importers or other considerations). Thus, the supply specification is capable of generating

collapses to a standard logit in the latter case. The demand parameters to be estimated are $(\beta; \gamma; \delta; \epsilon; \eta) \geq 0$.¹⁷

The nested logit structure yields well-known analytical expressions for the quantity of cement that each plant sells to each area (i.e., $q_{nt}(\mathbf{p}_{nt}; \beta; \gamma)$) and helps make estimation feasible from a computational standpoint. Nonetheless, the structure introduces some tension between the theoretical model and the empirical specification. Recall that the composite error term incorporates both an idiosyncratic preference shock and the consumer-specific deviation from mean distance (e.g., equation (7)). Since the deviation from mean distance is not independently distributed neither is the composite error.¹⁸ The relevance of this discrepancy depends on how much of the variation in the composite error term is due to variation in deviations from mean distance. There should be less tension between the theoretical model and the empirical specification when areas are small, and more tension when areas are large or preference shocks are degenerate (e.g., as in the "pure characteristics model").

6.1.3 Areas and potential demand

We define 90 consumer areas based on the counties of the U.S. Southwest. The choice implies relatively fine spatial price discrimination and enables us to model the geographic distribution of demand using commonly-available data at the county level. We normalize potential demand using exogenous demand factors, following standard practice for discrete-choice systems (e.g., Berry, Levinsohn, and Pakes (1995), Nevo (2001)). The two factors we select are the number of construction employees and the number of new residential building permits. Thus, we implicitly assume that construction spending is unaffected by cement prices, consistent with the fact that that portland cement composes only a small fraction of total construction expenditures.¹⁹

¹⁷The substitution patterns between cement plants are characterized by the independence of irrelevant alternatives (IIA) within the inside good nest. IIA may be 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. Thus, we are 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.

¹ For a given consumer, the deviations from distance can be positively or negatively correlated. For instance, consider two plants located on either side of an area: a consumer that is closer to the first plant is farther from the second plant. But if the two plants are on the same side then a consumer that is closer to the first plant is also closer to the second plant.

¹ 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 concrete) must be even lower.

To perform the normalization, we regress regional portland cement consumption on the demand predictors (aggregated to the regional level), impute predicted consumption at the county level based on the estimated relationships, and then scale predicted consumption by a constant of proportionality to obtain potential demand.²⁰ 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. 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.

6.2 Estimation

in region r . Then the aggregated regional-level metrics take the form:

$$\begin{aligned}
 \tilde{p}(\cdot; \cdot) &= \sum_{z_r} \sum \frac{\tilde{q}(\cdot; \cdot)}{\sum_{z_r} \sum \tilde{q}(\cdot; \cdot)} \tilde{p}(\cdot; \cdot) \\
 \tilde{q}(\cdot; \cdot) &= \sum_{z_r} \tilde{q}(\cdot; \cdot); \\
 \tilde{c}(\cdot; \cdot) &= \sum \sum_{z@r} \tilde{q}(\cdot; \cdot);
 \end{aligned} \tag{19}$$

where $\tilde{p}(\cdot; \cdot)$ is the production-weighted average mill price, $\tilde{q}(\cdot; \cdot)$ is total production, and $\tilde{c}(\cdot; \cdot)$ is total consumption. We calculate regional prices and quantities for Northern California, Southern California, and the combined Arizona-Nevada region, and calculate regional consumption for Northern California, Southern California, Arizona, and Nevada.

We also exploit information on aggregated cross-region shipments to help identify the model.²² We denote the Sout S]em1.79Td[n95(Sou(ts)-331(to)-339Td-308(regional)-307(prices)-308(ar

9(p)[p](6

using the weighting $ma43769.89cmBTiQ$

Table 1: Artificial Data Test for Identification

Variable	Parameter	Truth ()	Transformed ($\tilde{}$)	Mean Est	RMSE
----------	-----------	-----------	---------------------------------------	----------	------

the standard deviation of each equilibrium price across the eleven starting points. (So there are 1,260 standard deviations for a typical equilibrium price vector of 1,260 plant-area elements.) The results indicate that the *maximum* standard deviation, over all candidate parameter vectors and all plant-area prices, is zero to computer precision. Thus, the Monte Carlo experiment finds no evidence of multiple equilibria. This may be unsurprising because, theoretically, uniqueness is ensured for two close cousins of our model: nested logit demand, convex marginal costs, and single-plant firms (Mizuno 2003), and logit demand, sufficiently increasing marginal costs, and multi-plant firms (Konovalov and Sandor 2010).

6.5 Key empirical relationships

Although the estimation routine relies on strong functional form assumptions on demand and marginal costs, it is nonetheless possible to visualize the key empirical relationships that drive the parameter estimates. We explore these relationships in Figure 4.

On the demand side, the price coefficient is primarily determined 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. There is weak negative correlation, consistent with downward-sloping but inelastic aggregate demand. Next, the distance coefficient is primarily determined by (1) the cross-region shipments moment, and (2) the relationship between the consumption and production moments. We plot the gap between production and consumption ("excess production") for each region in panel B. In many years, excess production is positive in Southern California and negative elsewhere, consistent with inter-regional trade flows. The magnitude of these implied trade flows drives 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 determined 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. Finally, the utilization parameters are primarily determined by (1) the relationship between the production moments (which determine utilization) and the consumption moments, and (2) the relationship between the production moments and the price moments. We explore the second source of identification in panel D, which shows




Figure 4: 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 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).

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.

7 Empirical Results

7.1 Demand estimates and transportation costs

Table 2 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.²⁶ The aggregate elasticity implied by the price coefficient is 0.16 in the

²⁶The other demand parameters take reasonable values and are precisely identified. The negative coefficient on the import dummy is likely due to the fact that observed import prices do not reflect the full price of imported cement (see Appendix D). The inclusive value coefficient suggests that consumer tastes for the

Table 2: Estimation Results

Variable	Parameter	Estimate	St. Error
<i>Demand</i>			
Cement Price		-0.087	0.002
Miles Diesel Price		-26.42	1.78
Import Dummy		-3.80	0.06
Intercept		1.88	0.08
Inclusive Value		0.10	0.004
<i>Marginal</i>			
<i>5Tfi-.10</i>			

Figure 5: 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 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.²⁸

Firms appear to exercise some degree of localized market power. To illustrate, we map the prices and market shares of the Clarksdale plant that correspond to numerical equilibrium in Figure 5. We mark the location of the Clarksdale plant with a star, and mark other plants with circles. As shown, 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 steep discounts. The locations of competitors also influence market share and prices, though these effects are more difficult to discern.

We explore these relationships more rigorously with regression analysis. We regress prices and market shares on three independent variables: (1) the distance between the plant and the county, (2) the distance between the county and the nearest other domestic plant,

² 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.

and (3) the estimated marginal cost of the plant. Among plant-county pairs within 100 miles, a 10 percent reduction in distance is associated with prices and market shares that are 0.9 percent and 14 percent higher, respectively; and 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. Each of these patterns is statistically significant at the one percent level.²⁹

7.2 Marginal cost estimates

We estimate marginal costs to be \$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. Integrating the marginal cost function over the levels of production that arise in numerical equilibrium yields an average variable cost of \$51 million. 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 effects on marginal costs, the results suggest that their cumulative contribution to variable 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.^{30,31}

Finally, we discuss the individual parameter estimates shown in Table 2, 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.

² We refer the readers to the working paper for more details on this regression.

³⁰Lafarge North America, one of the largest domestic producers, reports an average gross margin of 0.33 over 2002-2004 in its public accounting records.

³¹Fixed 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. Our 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.

Lastly, the crushed stone coefficient of 0.29 is too small, given that roughly 1.67 tonnes of raw materials are used per tonne of cement. We suspect that these discrepancies are due to measurement error in the data.³² Alternatively, they may be due to a failure of identification (e.g., see Section 6.3) or due to the implicit assumption that plant productivity is fixed over the sample period { it seems clear that the renegotiation of onerous labor contracts improved productivity in the 1980s (e.g., Northrup (1989), Dunne, Klimek, and Schmitz (2009)).

7.3 Regression tests




Figure 6: 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.

percent of the variation in average prices.³⁴

7.4 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. We use counterfactual simulations to evaluate a hypothetical merger between Calmat and Giord-Hill in 1986, when the firms together operated six plants and accounted for 43 percent of industry capacity in the U.S. Southwest.³⁵

We map the distribution of consumer harm over the U.S. Southwest in Figure 8. In

³⁴The model does not fully capture the fall in average prices over the 1980s and early 1990s. One plausible explanation is that the model does not account for the productivity improvements that occurred during the sample period (e.g., Northrup (1989), Dunne, Klimek, and Schmitz (2009)).

³⁵We follow standard practice to perform the counterfactuals. We define an ownership matrix \mathbf{P}^{post} that reflects the post-merger structure of the industry. We then compute the equilibrium post-merger price vector as the solution to Equation 6, substituting \mathbf{P}^{post} for \mathbf{P} : Following McFadden (1981) and

Figure 7: 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.

panel A we focus on the effects of the merger, absent any divestitures. The total loss of consumer surplus is \$1.4 million which is small relative to pre-merger consumer surplus of \$239 million. Consumer harm is concentrated in the counties surrounding Los Angeles and

Figure 8: Loss of Consumer Surplus Due to a Hypothetical Merger between Calmat and Gifford-Hill

mitigates consumer harm in Southern California but do little to reduce harm in Maricopa County. Additional counterfactual exercises indicate that a two-plant divestiture is needed if this harm is to be mitigated as well.

7.5 Comparison to market delineation

In the introduction, we argue that the market delineation model imposes awkward theoretical assumptions. We now contrast some of our results to those of Ryan (2009), a recent paper that uses market delineation in a study of the portland cement industry. In particular, we point out that our approach generates distinctly different estimates of aggregate elasticity than does the market delineation approach. The discrepancy is consistent with the notion that our estimation strategy may sometimes provide more reasonable results than conventional approaches, and that these differences can be sizeable.³⁶

³⁶The discrepancy does not diminish the substantial contribution of Ryan (2009), which estimates an in-

Ryan makes the common assumptions that demand has constant elasticity and supply is Cournot within each market. He estimates the aggregate elasticity to be -2.96 , which is quite different than our estimate of -0.16 . The difference is entirely due to specification choices { the constant elasticity demand system produces an aggregate elasticity of -0.15 once housing permits are included as a control.³⁷ However, Ryan cannot use the inelastic estimate because, within the context of Cournot competition, it would imply that the firm elasticities are small to be consistent with profit maximization. This occurs because the Cournot model restricts each firm elasticity to be linearly related to the aggregate elasticity according to the relationship $e_i = e \cdot s_i$, where e_i , e , and s_i denote the firm elasticity, the aggregate elasticity, and the firm market shares, respectively. Further, Ryan cannot use the nested logit system to divorce the firm elasticities from the aggregate elasticity (as we do) because logit models assume differentiated products whereas Cournot supply models assume homogenous products. Our takeaway is that our econometric strategy can lead to improved estimates by connecting the data to more realistic economic models.

8 Conclusion

production. We are enthused by the breadth of opportunity.

References

- Amemiya, T. (1985). *Advanced Econometrics*. Harvard University Press.
- Bajari, P., H. Hong, and S. Ryan (2008). Identification and estimation of discrete games of complete information. Working paper, Duke University.
- Berry, S., J. Levinsohn, and A. Pakes (1995, July). Automobile prices in market equilibrium. *Econometrica*, 841{890.
- Bisin, A., A. Moro, and G. Topa (2010). The empirical content of models with multiple equilibria in economies with social interactions. Mimeo.
- 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 Competition". *Econometrica* 47, 1145{1150.
- Davis, P. (2006). Spatial competition in retail markets: Movie theaters. *The RAND Journal of Economics* 37, 964{982.

- 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. Martinez, 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*.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69, 307{342.
- Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. *Handbook of Econometrics* 4.
- Northrup, H. R. (1989). From union hegemony to union disintegration: Collective bargaining in cement and related industries. *Journal of Labor Research* 10, 337{376.
- Pesendorfer, M. (2003). Horizontal mergers in the paper industry. *RAND Journal of Economics* 34, 495{515.
- Pinske, J., M. E. Slade, and C. Brett (2002). Spatial price competition: A semiparametric approach. *Econometrica* 70(3), 1111{1153.
- 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.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurchner. *Econometrica* 55, 993{1033.
- Ruud, P. A. (2000). *An Introduction to Classical Econometric Theory*. Oxford University Press.
- 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 oligopoly: 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.

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

A Proofs

Proof of Lemma 1: We first show that $\mathbf{S}(\mathbf{H}(\cdot; \cdot))$ is continuously differentiable in \mathbf{z} . The proof is by contradiction. Suppose that $\mathbf{S}(\mathbf{H}(\cdot; \cdot))$ is not continuously differentiable at some parameter vector $\mathbf{z}_1 \in \mathcal{Z}$, i.e., that $\mathbf{S}(\mathbf{H}(\cdot; \cdot))|_{\mathbf{z}=\mathbf{z}_1}$ is discontinuous at \mathbf{z}_1 . Then, by the linearity of \mathbf{S} and the definition of discontinuity,

$$\lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^-} \left. \frac{\partial \mathbf{H}(\cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} \neq \lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^+} \left. \frac{\partial \mathbf{H}(\cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} :$$

However, the function $\mathbf{f}(\mathbf{p}; \cdot; \cdot)$ is continuously differentiable in \mathbf{p} and \mathbf{z} by the assumptions placed on $\mathbf{q}(\mathbf{p}; \cdot; \cdot)$ and $\mathbf{c}(\mathbf{p}; \cdot; \cdot)$. It follows that $\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)|_{\mathbf{z}=\mathbf{z}_1}$ is continuous, i.e. that

$$\lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^-} \left. \frac{\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} = \lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^+} \left. \frac{\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} :$$

Totally differentiating both sides, using the arbitrary price vector $\mathbf{H}(\cdot; \cdot)$, yields

$$\begin{aligned} \lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^-} \left(\left. \frac{\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)}{\partial \mathbf{H}(\cdot; \cdot)} \frac{\partial \mathbf{H}(\cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} + \left. \frac{\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} \right) \\ = \lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^+} \left(\left. \frac{\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)}{\partial \mathbf{H}(\cdot; \cdot)} \frac{\partial \mathbf{H}(\cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} + \left. \frac{\partial \mathbf{f}(\mathbf{p}; \cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} \right) : \end{aligned}$$

Since $\mathbf{f}(\mathbf{p}; \cdot; \cdot)$ is continuously differentiable in \mathbf{p} and \mathbf{z} , it follows that

$$\lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^-} \left. \frac{\partial \mathbf{H}(\cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} = \lim_{\mathbf{z}' \rightarrow \mathbf{z}_1^+} \left. \frac{\partial \mathbf{H}(\cdot; \cdot)}{\partial \theta} \right|_{\mathbf{z}=\mathbf{z}'} ;$$

which creates the contradiction. It remains to show that $\mathbf{S}(\mathbf{H}(\cdot; \mathbf{y}))$ is continuously differentiable in \mathbf{y} for $\mathbf{z} \in \mathcal{Z}$, where \mathbf{y} is the vector representation of \mathcal{Y} . The proof is obtainable by contradiction, using the same steps employed above, and we omit the explicit derivation for expositional brevity. □

Proof of Theorem 1: We first place regularity conditions on the data generating process. Let \mathbf{y} be the vector representation of the set \mathcal{Y} . We assume that $\{y_j\}$ is a sequence of i.i.d. random vectors. We further assume that $\sup_{\theta \in \Theta} \int \mathbf{p} \cdot \mathbf{S}(\tilde{\mathbf{H}}(\cdot; \cdot)) y_j < 1$, that $\sup_{\theta \in \Theta} \int \mathbf{p} \cdot \mathbf{H}(\cdot; \cdot) = \theta < 1$, and that $\sup_{\theta \in \Theta} \int \mathbf{p} \cdot \mathbf{S}(\tilde{\mathbf{H}}(\cdot; \cdot)) [\mathbf{p} \cdot \mathbf{S}(\tilde{\mathbf{H}}(\cdot; \cdot))]^{\theta} < 1$. Amemiya (1985) proves that these conditions, along with the assumptions already introduced

in the body of the text, imply the following properties:

- (i) $\frac{1}{2} \sum_{=1} [\mathbf{p} \quad \mathbf{S}(\widetilde{\mathbf{H}}(\mathbf{z};))] ! \dot{\mathbf{f}} E[\mathbf{p} \quad \mathbf{S}(\widetilde{\mathbf{H}}(\mathbf{z};))]$ uniformly in \mathcal{Z} ,
- (ii) $\frac{1}{2} \sum_{=1} [\textcircled{\mathbf{H}}(\mathbf{z};)=\textcircled{\mathbf{H}}] ! \dot{\mathbf{f}} E[\textcircled{\mathbf{H}}(\mathbf{z};)=\textcircled{\mathbf{H}}]$ uniformly in \mathcal{Z} ,
- (iii) $\frac{1}{2} \sum_{=1} [\mathbf{p} \quad \mathbf{S}(\widetilde{\mathbf{H}}(\mathbf{z};))][\mathbf{p} \quad \mathbf{S}(\widetilde{\mathbf{H}}(\mathbf{z};))]$

JF11J/F16 1

Table 3: Consumption, Production, and Prices

Description	Mean	Std	Min	Max
<i>Consumption</i>				
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
<i>Production</i>				
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
<i>Domestic Prices</i>				
Northern California	85.81	11.71	67.43	108.68
Southern California	82.81	16.39	62.21	114.64
Arizona-Nevada	92.92	14.24	75.06	124.60
<i>Import Prices [excludes duties and grinding costs]</i>				
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.

simplex methods such as simulated annealing and the Nelder-Mead algorithm, as well as 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

GMM procedure, and we estimate the inclusive value coefficient using $\tilde{\lambda} = \log\left(\frac{\lambda}{1-\lambda}\right)$. 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. To start, we note that the California Letter is based on a monthly survey rather than on the annual USGS census, which creates minor discrepancies. We normalize the California Letter data prior to estimation so that total shipments equal total production in each year. The 96 cross-region data points include:

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

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

The portland cement plant in Riverside closed its kiln permanently in 1988 but continued operating its grinding mill with purchased clinker. We include the plant in the analysis over 1983-1987, and we adjust the USGS production data to remove the influence of the plant over 1988-2003 by scaling the data downward, proportional to plant grinding capacities. 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.

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.

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.