Competition among Spatially Di erentiated Firms: An Empirical Model with an Application to Cement*

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 e ects empirically, we develop an estimator for m71(75)-4ts of spatial di erentiation 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 fits, both in-sample and out-of-sample, that the framework explains well the salient features of competition. We estimate transportation costs to be $0.30 per tonne-mile and show that these costs constrain shipping distances and create localized market power. To policy-relevance, we conduct counter-factual simulations that quantify competitive harm from a hypothetical merger. We map the distribution of harm over geographic space and identify the divestiture that best mitigates harm.

Keywords: transportation costs; spatial di erentiation; price discrimination; cement
JEL classification: C51; L11; L40; L61

*Miller: nathan.miller@usdoj.gov. Osborne: matthew.osborne@bea.gov. We thank Masak Ishihara, Robert Johnson, Scott Kominers, Ashley Langer, Russell Pittman, Chuck Romeo, Jim Schmitz, Adam Shapiro, Charles Taragin, Raphael Thomadsen, Glen Weyl and seminar participants at the Bureau of Economic Analysis, Georgetown University, and the U.S. Department of Justice for v comments. Sarolta Lee, Parker Sheppard, Gian Wrobel, and Vera Zavin provided research assistance. The views expressed herein are entirely those of the authors and should not be purported to reflect those of the U.S. Department of Justice or attributed to the Bureau of Economic Analysis.

1 Introduction

In many industries, firms are geographically di erentiated and transportation is costly. Yet few empirical studies estimate structural models of spatial di erentiation. We attribute this dearth of research to a simple *data availability problem*: the most straight-forward way to identify the degree of spatial di erentiation { or, equivalently, the magnitude of transportation costs { is to measure how firms' market shares di er between nearby and distant consumers. But this requires data on the geographic distributions of the market shares. These data are di cult 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), McManus (2009)),[1] to our knowledge no previous work estimates structural parameters in the face of spatial price discrimination { despite the oft-cited result of Greenhut, Greenhut, and Li (1980) that two-thirds of surveyed firms employ some form of spatial price discrimination.[2]

This sparse empirical literature is in contrast to a storied theoretical literature (e.g., normal. We also conduct an empirical application and demonstrate that (1) estimation is markets.[4] These assumptions preclude inference regarding spatial di erentiation 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 Hortaçsu (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, and consumption, each at the regional level (e.g., we observe total consumption separately for northern California, southern California, Arizona and Nevada). We find that the estimation procedure produces impressive in-sample and out-of-sample fits despite parsimonious demand and marginal cost specifications. For instance, 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. The quality of these fits is underscored by the rich time-series variation in these data due to macro-economic fluctuations.

tion.) 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 sub-markets for each candidate parameter vector, given prices and an observed geographic distribution of consumers, and aggregate these predictions to construct predicted firm-level sales. The estimation procedures then selects the parameters that bring the predicted firm-

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 $\{z_1; z_2; :::; z_J\}$, where $z_j \in \mathbb{C}$. 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 (w) characterizes the geographic distribution of consumers and we define $M = {}_{\mathbb{C}} \; (dw)$ to be the potential demand of the space. We denote the distance between any two points in the geographic space, say a and b, as the Euclidean distance $\|a - b\|$.

Without loss of generality, we partition the geographic space into N distinct geographic consumer *areas*, such that each area $\mathbb{C}_n \in \mathbb{C}$ is itself a connected set in $\mathbb{R}^2$. We conduct the partition such that $\mathbb{C}_1 \quad \mathbb{C}_2 \quad \cdots \quad \mathbb{C}_N = \mathbb{C}$

produced by plant j, and $c(Q;\mathbf{w}_j;\ _0)$ is a convex and di erential marginal cost function of known form. The vectors $\mathbf{x}_{jn}$ and $\mathbf{w}_j$ include demand and cost shifters, respectively, and the matrix $\mathbf{X}_j$ stacks the relevant $\mathbf{x}_{jn}$ vectors. Finally, $\ _0$ is a K-dimensional parameter vector.

We model consumer behavior using a conventional discrete-choice demand system. Each consumer observes the plant locations and the available mill prices, and either purchases from one of the J plants or foregoes a purchase altogether (i.e., selects the outside good). The indirect utility that consumer i receives from plant j is:

$$u_{ij} = \ ^c + \ ^p p_{nj} + \ ^d \|w - z_j\| + \mathbf{x}'_{jn}\ ^x + \ _{ij}; \tag{2}$$

where $\ _{ij}$ is an idiosyncratic preference shock that is observed to the consumer and uncorrelated with distance and prices. We provide motivation for the preference shock to Section 3.4. Following standard practice, we normalize the mean utility of the outside option to zero. Finally, $(\ ^c;\ ^p;\ ^d;\ ^x) \in \ _0$ are the demand parameters and the ratio $\ ^d = \ ^p$ represents the unit transportation cost incurred by consumers.

We assume that consumers select the plant that supplies the highest utility. Within an area, this assumption defines the set of consumer characteristics $(w \in \mathbb{C}_n;\ _i)$ that lead to the selection of plant j, and we denote this set

$$A_{jn}(\mathbf{p}_n;\mathbf{z};\mathbf{x}_{jn};\ _0) = \{(w \in \mathbb{C}_n;\ _i)|u_{ij} \geq u_{ik}\ \forall\ k = 0;1;:::J\};$$

where $\mathbf{z} = (z_1;z_2;:::;z_J)$. If ties occur with zero probability then the quantity produced by plant j and consumed in area $\mathbb{C}_n$ is given by:

$$q_{jn}(\mathbf{p_n};\mathbf{z};\mathbf{x}_{jn};\ _0) = M_n \int_{A_j} @P(w;\ ); \tag{3}$$

where $P(\cdot)$ denotes a population distribution function. We place two normalcy conditions on demand, namely that $q_{jn}(\mathbf{p_n};\mathbf{z};\mathbf{x}_{jn};\ _0)$ is twice continuously di erentiable and also downward sloping in p_{jn} (i.e., $\ ^p < 0$).

For clarity, we sketch one possible geographic space in Figure 1. The dashed lines delineate three consumer areas, $\mathbb{C}_1$, $\mathbb{C}_2$, and $\mathbb{C}_3$. Two plants operate in the space and are characterized by the locations z_1 and z_2. A distribution of consumers span the space, and both plants compete for each consumer. The plants imperfectly price discriminate by setting

and (2) demand in area $\mathbb{C}_n$ is una ected by mill prices in area $\mathbb{C}_m$ for $n \neq m$.

We now rearrange and stack the first-order conditions:

$$f(p;\ ;\ _0) \equiv p - c(Q(p;\ ;\ _0);\ ;\ _0) + \Omega^{-1}(p;\ ;\ _0)q(p;\ ;\ _0) = 0: \qquad (6)$$

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

3.4 Discussion

We o er 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 tradeo 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 a ect the plant's competitiveness in other areas.

p)o68(d) tf Second, the areas $\mathbb{C}_1; \mathbb{C}_2; ::: \mathbb{C}_N$ are best interpreted as determining the extentte9Td[(1)]TJa-wt tion (see Section 6.1.2). In the more general case, the idiosyncratic preference shocks can be motivated as capturing various plant- and consumer-level heterogeneity that, due to costly bargaining or other reasons, is not reflected in mill prices.

4 Estimation

The model generates a rich set of predictions on equilibrium prices, production, and shipments within the geographic space $\mathbb{C}$. Yet the parameters of the model can be recovered using relatively coarse data. In this section, we introduce a novel GMM estimation procedure that exploits variation in $t = 1;2;:::T$ time-series observations on aggregated price moments (e.g., observations on average mill prices). We show that the GMM estimator is consistent and asymptotically normal, given assumptions on the existence and uniqueness of equilibrium. We then demonstrate that the estimator can be extended in a straight-forward manner to incorporate observations on aggregated non-price moments (e.g., total production or total consumption). The flexibility of these data requirements makes the estimator widely applicable to economic settings characterized by spatial di erentiation.

4.1 GMM estimation

We first clarify the level of detail on consumer locations, $w \in \mathbb{C}$, needed to support estimation. In many instances, precise consumer locations may be unavailable or too costly to discover, so that the direct application of equation (2) is infeasible. We make the following assumption on the exogenous spatial data available to the econometrician:

Assumption A1: *The econometrician observes the mean distance between plant j and the consumers in area $\mathbb{C}_n$, for all j and n.*

We denote the mean distance between plant j and the consumers in area $\mathbb{C}_n$ as d_{jnt}. Under A1, we can rewrite the indirect utility equation as follows:

$$u_{ijt} = {}^{c} + {}^{p} p_{njt} + {}^{d} d_{njt} + \mathbf{x}'_{jnt}\,{}^{x} + {}^{*}_{ijt}; \tag{7}$$

where ${}^{*}_{ijt}$ is a composite error term that includes the idiosyncratic preference shock and the consumer-specific deviation from mean distance. Formally, ${}^{*}_{ijt} = {}_{ijt} + {}^{d}(\|w - z_{jt}\| - d_{njt})$. The composite error term is orthogonal to price and mean distance, given the assumptions already placed on the model. As long as the distributions of ${}^{*}_{ijt}$ 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* reasonable to further assume that the sampling error is independent of the \right-hand-side" data $_t$. This simplifies the construction of the estimator, and we impose the additional assumption here:

Assumption A4′: *The sampling error is mean zero conditional on* $_t$*:*

$$E[\mathrm{p}_t^d - \mathrm{S}(\mathrm{H}(\ _0;\ _t))|\ _t] = 0:$$

A4′ enables estimation with multiple equation nonlinear least squares, which is equivalent to GMM with the optimal instruments

$$\mathrm{Z}_t = -\frac{@\mathrm{S}(\mathrm{H}(\ _0;\ _t))}{@\ _0}\Lambda_0(\ _0)^{-1}; \tag{10}$$

where $\Lambda_0(\ _0) \equiv E[\mathrm{S}(\ _t)|\ _t]E[\mathrm{S}(\ _t)|\ _t]'$ is the variance matrix of the aggregated error terms. Thus, the sample moment equations that correspond to A4′ are

$$\frac{1}{T}\ _{t=1}^{T}\ -\frac{@\mathrm{S}(\mathrm{H}(\ ;\ _t))}{@}\mathrm{C}_T^{-1}(\mathrm{p}_t^d - \mathrm{S}(\mathrm{H}(\ ;\ _t))); \tag{11}$$

where C_T is some consistent estimate of $\Lambda_0(\ _0)$ and 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., $\mathrm{H}(\ ;\ _t)$). 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 ($\ _{ij}^*$ in equation (7)). In our application, we are typically able to numerically compute a vector, call it $\mathrm{H}(\ ;\ _t)$, 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:

$$= \arg\min_{\in\Theta} \frac{1}{T} \quad {}^{T}_{t=1} [\mathrm{p}_t^d - \mathrm{S}(\mathrm{H}(\ ;\ _t))]' \mathrm{C}_T^{-1} [\mathrm{p}_t^d - \mathrm{S}(\mathrm{H}(\ ;\ _t))]: \qquad (12)$$

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* $\in \Theta$ *there exists a vector* $\mathrm{p}_1 \in \mathbb{R}_+^{JN}$ *such that* $\mathrm{f}(\mathrm{p}_1;\ _t;\ ) = 0$. *Further,* $\mathrm{f}(\mathrm{p}_1;\ _t;\ ) = \mathrm{f}(\mathrm{p}_2;\ _t;\ ) = 0 \leftrightarrow \mathrm{p}_1 = \mathrm{p}_2$.

A5 ensures that the GMM objective function is well-behaved.[10] We suspect that uniqueness alone may su ce 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* $\mathrm{S}(\mathrm{H}(\ ;\ _t))$ *is continuously di erentiable over* Θ.

Proof. See appendix A. □

[10] Recent 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, su ciently increasing marginal costs, and multi-plant firms (Konovalov and Sándor 2010). The assumption is not satisfied generally (e.g., Caplin and Nalebu (1991)).

Assumption A6: *The parameter vector* $_0$ *is globally identified in* Θ*. Formally,* $E[\mathrm{p}_t^d - \mathrm{S}(\mathrm{H}(\ ;\ _t))|\ _t] = 0 \leftrightarrow\ =\ _0$*:*

A6 could be violated even if parameters of the model would be globally identified given disaggregate data (i.e., even if $E[\mathrm{p}_t - \mathrm{H}(\ ;\ _t)|\ _t] = 0 \leftrightarrow\ =\ _0$). 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) $\ \rightarrow^p\ _0$ *and*

ii) $\sqrt{T}(\ -\ _0) \rightarrow^d N(0; \mathrm{V})$*;*

where $\mathrm{V} = (\mathrm{G}_0'\mathrm{C}_0\mathrm{G}_0)^{-1}\mathrm{G}_0'\mathrm{C}_0\Lambda_0\mathrm{C}_0\mathrm{G}_0(\mathrm{G}_0'\mathrm{C}_0\mathrm{G}_0)^{-1}$ *and* $\mathrm{G}_0 \equiv -E[@\mathrm{S}(\mathrm{H}(\ ;\ _t))=@\ ']$.

Proof. See appendix A. □

4.3 Incorporating non-price moments

The estimation strategy can be extended to incorporate observations on non-price moments, such as aggregate production or aggregate consumption, that are often available to the econometrician. We focus on production data for expositional brevity; the other endogenous data can be incorporated analogously. We assume the data are generated by:

$$\mathrm{q}_t = \mathrm{q}(\mathrm{H}(\ _0;\ _t);\ _t;\ _0)) +\ _t^*; \tag{13}$$

where q_t is a vector of length JN, and

The GMM estimate that incorporates these data is:

$$^* = \arg\min_{\in\Theta} \frac{1}{T} \sum_{t=1}^{T} \begin{pmatrix} p_t^d - S(H(\; ; \; _t)) \\ q_t^d - R(q(H(\; ; \; _t); \; _t; \;)) \end{pmatrix}' D_T^{-1} \begin{pmatrix} p_t^d - S(H(\; ; \; _t)) \\ q_t^d - R(q(H(\; ; \; _t); \; _t; \;)) \end{pmatrix} ; \quad (15)$$

where D_T is some positive definite matrix, and the relevant contemporaneous variance matrix is $\Lambda_0^*(\; _0) \equiv E[S(\; _t) \; R(\; _t^*)| \; _t]E[S(\; _t) \; R(\; _t^*)| \; _t]'$. Under A5 and an appropriately modified A6, Theorem 1 extends and the estimate is consistent and asymptotically normal.

5 The Portland Cement Industry

5.1 The product

Portland cement is a finely ground dust that forms concrete when mixed with water and coarse aggregates such as sand and stone. Concrete, in turn, is an essential input to many construction and transportation projects because its local availability and lower maintenance costs make it more economical than substitutes such as steel, asphalt, and lumber (Van Oss and Padovani (2002)). The producers of portland cement adhere to strict industry standards that govern the production process. Aside from geographic considerations, product di erentiation in the industry is minimal.[11]

Producers negotiate private contracts with their customers, predominately ready-mix concrete firms and large construction firms. Most contracts specify a mill (or \free-on-board") price for portland cement at the location of production. Customers are responsible for door-to-door transportation, which is an important consideration because portland cement is inexpensive relative to its weight.[12] This fact is well understood in the academic literature. For example, Scherer et al (1975) estimates that transportation would account for roughly one-third of total customer expenditures on a hypothetical 350-mile route between Chicago and Cleveland, and a 1977 Census Bureau study reports that more than 80 percent is transported within 200 miles.[13] More recently, Salvo (2010) presents evidence consistent with the importance of transportation costs in the Brazilian portland cement industry.

[11] The standards are maintained by the the American Society for Testing and Materials Specification for Portland Cement, and exist to protect the quality and reliability of construction materials.

[12] The bulk of portland cement is moved by truck, though some is sent by train or barge to distribution terminals and only then trucked to customers.

[13] 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.

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 five 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 o ces through which foreign imports enters 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. For instance, more than 98 percent of cement produced in southern California was shipped within the U.S. Southwest over 1990-1999, and more than 99 percent of cement produced in California was shipped within the region over 2000-2003. Outflows from Arizona and Nevada are unlikely because consumption routinely exceeds production in those states. And since net trade-flows between the U.S. Southwest and other domestic regions are insubstantial, these data points also imply that gross domestic inflows must also be insubstantial.

5.3 Data

We collect our endogenous 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 cement plants that collects detailed information on consumption, production, and mill prices.

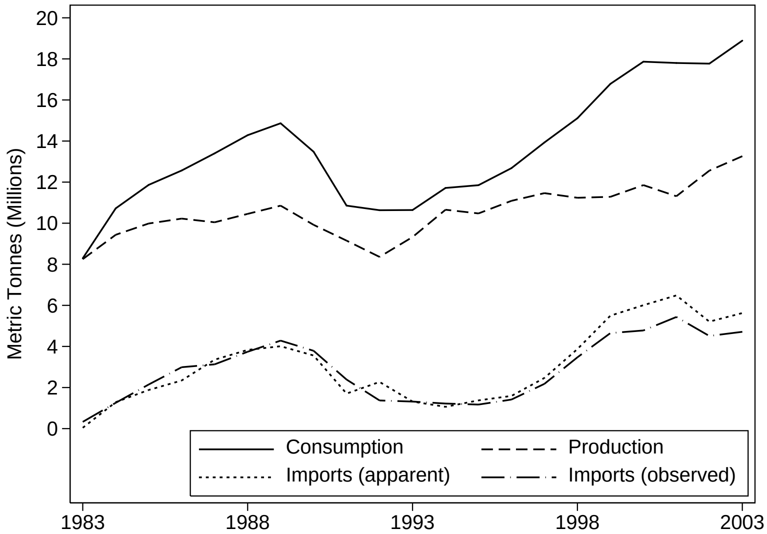

Figure 3: Consumption, Production, and Imports of Portland Cement. Apparent imports are defined as consumption minus production. Observed imports are total foreign imports shipped into San Francisco, Los Angeles, San Diego, and Nogales.

Census response rates are typically well over 90 percent, and the USGS sta imputes missing values for non-respondents based on historical and cross-sectional information.[15] 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

[15]The quality of the census has long generated interest among researchers. Other academic studies that feature USGS data include McBride (1983), Rosenbaum and Reading (1988), Jans and Rosenbaum (1997), Syverson and Hortaçsu (2007), and Ryan (2009).

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 o ces of the U.S. Southwest. Consumers pay the door-to-door cost of transportation from these customs o ces. 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 the stylized fact that foreign importers provide substantial quantities of portland cement to the U.S. Southwest when demand is strong.

6.1.2 Demand

We express the indirect utility that consumers i receives from domestic plant j as follows:

$$u_{ijt} = {}^{c} + {}^{p}p_{jnt} + {}^{d}MILES_{jn} * DIESEL_{t} + {}^{*}_{ijt}; \tag{17}$$

where p_{jn} is the mill price per metric tonne, $MILES_{jn}$ is the miles (in thousands) between the plant and the centroid of the consumer's area and $DIESEL_t$ is a diesel price index that equals one in the year 2000. Hence, transportation costs increase linearly in distance and fuel costs, and the combination ${}^{d}= {}^{p}$ is the cost per thousand tonne-miles given diesel prices at the 2000 level. We express the indirect utility received from the foreign importers as:

$$u_{i;J+1;t} = {}^{c} + {}^{i} + {}^{p}p_{J+1;t} + {}^{d}MILES_{J+1;n} * DIESEL_{t} + {}^{*}_{i;J+1;t}; \tag{18}$$

where $MILES_{J+1;n}$ is the miles (in thousands) between the centroid of the consumer's area and the nearest customs o ce. The import-specific intercept is needed because the USGS data on import prices exclude duties. To be clear, the domestic prices are not observed in the data and must be computed as the solution to equation (6), given import prices that are exogenously-determined, non-discriminatory, and observed in the data. Finally, we normalize the mean value of the outside good to zero, so that $u_{i0t} = {}^{*}_{i0t}$.

We assume that the distributions of the composite error terms (${}^{*}_{ijt}$) generate a nested logit demand system in which the inside options (the domestic plants and the foreign imports) are in a di erent nest than the outside option. That is, we assume the composite error terms have i.i.d. extreme value distributions and define a parameter that characterizes the degree to which valuations of the inside options are correlated across consumers (e.g., as in Cardell (1997)). Valuations are perfectlyp91erfs

collapses to a standard logit in the latter case. The demand parameters to be estimated are
(

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.[19] 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

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 $\frac{1}{J N} \| f(p ;\ _t;\ ) \| <$, where $\| \cdot \|$ denotes the Euclidean norm operator. The vector that defines numerical equilibrium given the 2003 data has $14 \times 90 = 1;260$ elements.[20]

We construct regional-level metrics based on the computed numerical equilibrium to compare the model predictions against the data. For notational convenience, we denote the elements of the equilibrium price vector as as $p_{jnt}(\ ;\ _t)$, and the corresponding quantities as $q_{jnt}(\ ;\ _t)$. We also define the sets $\aleph_r$ and $|_r$ as the counties and plants, respectively, located

[19] The regression of regional portland cement consumption on the demand predictors yields an ^{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 su cient to ensure that potential demand exceeds observed consumption in each region-year observation.

[20]

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

$$p_{rt}(\; ; \;\; _t) \quad = \quad \sum_{j \in |_r}$$

using the weighting matrix $\Lambda \otimes \mathrm{I}$: 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).[22]

6.3 Identification

The use of aggregated data precludes point identification if multiple candidate parameters produce identical aggregate predictions despite having distinct disaggregate predictions. (This would violate A6.) To check for aggregation problems, we pair a vector of \true" parameters with 40 randomly-drawn sets of exogenous data. Both the parameters and the data are chosen to mimic our empirical application. For each of set of exogenous data, we compute equilibrium, generate the relevant aggregated data, and estimate using GMM. We argue that the model is reasonably identified if the GMM estimates are close to the true parameters.[23]

Table 1 shows the results of the artificial data experiment. Interpretation is complicated somewhat because we use non-linear transformations to constrain the some of coe cients (e.g., $^{p} < 0$), and we defer details on these transformations to Appendix C. Nonetheless, it is clear that the means of the estimated coe cients are close to transformed true parameters. The means of the price and distance coe cients, which are particularly relevant for spatial models, 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 well identified.

[22] Estimation of the contemporaneous variance matrix is complicated by the fact that we observe prices, production, and consumption 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.

[23] 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. 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 and data that lead the estimator to nonsensical areas of parameter space. Throughout, we hold plant and county locations fixed to maintain tractability, and rely on the random draws of capacity, potential demand, and diesel prices to create variation in the distances between production capacity and consumers. Each artificial data set includes 21 draws on the exogenous data, with each draw representing a single time-series observation.

Table 1: Artificial Data Test for Identification

Variable	Parameter	Truth ( )	Transformed ( )	Mean Est	RMSE
Demand					
Cement Price	p	-0.07	-2.66	-2.51	0.66
Miles×Diesel Price	d	-25.00	3.22	2.86	0.59
Import Dummy	i	-4.00	-4.00	-6.07	1.23
Intercept	c	2.00	2.00	1.11	0.51
Inclusive Value		0.09	-2.31	-1.73	0.54
Marginal Costs					
Cost Shifter 1	1	0.70	-0.36	-0.88	0.51
Cost Shifter 2	2	3.00	1.10	0.54	0.45
Utilization Threshold		0.90	2.19	1.71	0.59
Over-Utilization Cost		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.

6.4 Multiple equilibria

We search for only a single equilibria in the inner loop, and problems may arise if multiple equilibria exist (this would violate A5). To provide some reassurance that the estimation procedure is reasonable, we conduct a Monte Carlo experiment and search for the existence of multiple equilibria. In particular, we compute equilibrium at eleven di erent starting points for thousands of randomly-drawn candidate parameter vectors. We then evaluate whether, for each given candidate parameter vector, the computed equilibrium prices are sensitive to the starting points.[24] More precisely, for each candidate parameter vector, we calculate 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

[24] We consider 300 parameter vectors for each of the 21 years in the sample, for a total of 6,300 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 Sándor 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 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.

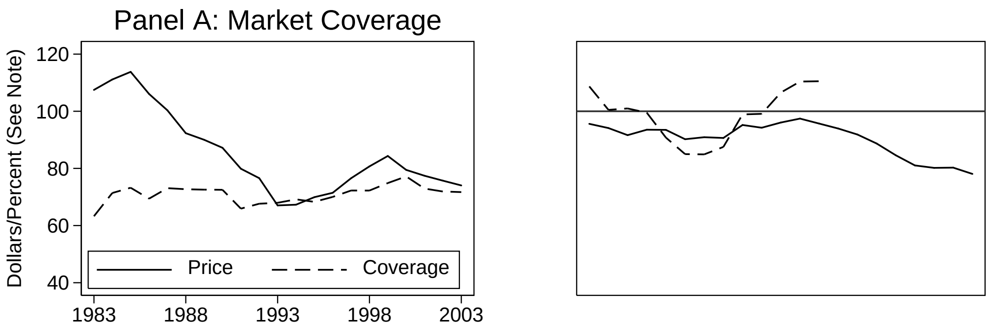

Figure 4: Empirical Relationships in the U.S. Southwest. Panel A plots average cement prices and market coverage. Prices are in dollars per metric tonne and market coverage is defined as the ratio of consumption to potential demand (times 100). Panel B plots excess production in each region, which we define as the gap between 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).

7 Empirical Results

7.1 Demand estimates and transportation costs

Table 2 presents the parameter estimates of the GMM procedure. The price and distance coe cients are the two primary objects of interest on the demand side; both are negative and precisely estimated.[25] The agg69etsnd plotsofnandbed7.Tf13.3(ted7.Tftes)-31 es-D2r

Table 2: Estimation Results

Variable	Parameter	Estimate	St. Error
Demand			
Cement Price			

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.

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.

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.[28]

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 measurement error in the data.[31] Alternatively, they may be induced by the implicit assumption that plant productivity is fixed over the sample period.

7.3 Regression fits

One measure of an econometric model's viability is in its ability to fit the data.[32] In Figure 6, we plot observed consumption against predicted consumption (panel A), observed production against predicted production (panel B), and observed prices against predicted prices (panel C). Univariate regressions of the data on the predictions indicate that 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. Thus, the model performs reasonably well in accounting for the variation in the endogenous data.

It is also telling to examine the model's out-of-sample predictions. In panel D, we plot observations on cross-region shipments against the corresponding model predictions. We use 14 of these observations in the estimation routine { the shipments from plants in California to consumers in northern California over 1990-2003 { but the remaining 82 data points are withheld from the estimation procedure and do not influence the estimated parameters. Even so, the model explains 98 percent of the variation in these data.

The quality of these fits is underscored by the rich time-series variation in the data due to macro-economic fluctuations. To illustrate, we aggregate the data and the model predictions across regions, and plot the resulting time-series in Figure 7. 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 predictions 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.[33]

[31] 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).

[32] We are unaware of any statistical specification tests that are suitable for GMM with optimal instruments, which is exactly identified by construction.

[33]

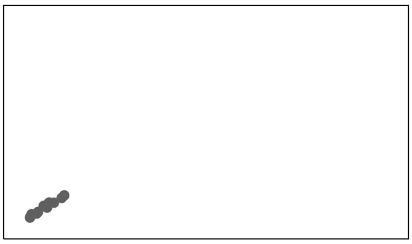

Figure 6: GMM Estimation Fits for Regional Metrics. Consumption, production, and cross-region shipments are in millions of metric tonnes. Prices are constructed as a weighted-average of plants in the region, and are reported as dollars per metric tonne. The lines of best fit and the reported ^{2} values are based on univariate OLS regressions.

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 Gi ord-Hill in 1986, when the firms together operated six plants and accounted for 43 percent of industry capacity in the U.S. Southwest.[34]

[34] We follow standard practice to perform the counterfactuals. We define an ownership matrix $\mathbf{\Omega^{post}}(\mathbf{P})$ 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 $\mathbf{\Omega^{post}}(\mathbf{P})$ for $\mathbf{\Omega}(\mathbf{P})$. Following McFadden (1981) and Small and Rosen (1981), the change in consumer surplus due to the merger is:

$$\Delta C = \sum_{n=1}^{N} \frac{\ln(1+\exp(\beta^c+\lambda I_{nt}^{pre})) - \ln(1+\exp(\beta^c+\lambda I_{nt}^{post}))}{\beta^p} M_n,$$

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.

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 ^{2} values are calculated from univariate regressions of the observed metric on the predicted metric.

We map the distribution of consumer harm over the U.S. Southwest in Figure 8. In 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 Phoenix. Indeed, Maricopa County and Los Angeles County alone account for 60 percent of consumer harm and 10 counties account for more than 90 percent of the harm. We focus on potential remedies in panel B. We find that it is possible to eliminate 56 percent of the harm through the divestiture of a single plant. The divestiture of either the “Gifford-Hill 2” plant or the “Calmat 2” plant accomplishes this. As shown, however, these divestitures mitigates consumer harm in Southern California but do little to reduce harm in Maricopa

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

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.[35]

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

[35] The discrepancy does not diminish the substantial contribution of Ryan (2009), which estimates an innovative dynamic discrete choice game and focuses primarily on the dynamic parameters; market delineation is used simply to determine the payoffs at different realizations of the state space.

choices { the constant elasticity demand system produces an aggregate elasticity of -0.15 once housing permits are included as a control.[36] 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 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. Further, Ryan cannot use the nested logit system to divorce the firm elasticities from the aggregate elasticity (as we do) because because logit models assume di erentiated 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

The literature of the \new empirical industrial organization" focuses on the structural estimation of competition models and the recovery of the underlying parameters that guide firm and consumer decisions. Econometric innovations and greater computer power have improved our ability to link empirical correlations with sensible theoretical models of behavior. One area of particular interest has been the estimation of product di erentiation models, as in Berry, Levinsohn, and Pakes (1995) and Nevo (2001). Yet geographic considerations { often critical drivers of di erentiation { have received relatively little attention.

In providing an estimator for economic models of spatial price di erentiation and spatial price discrimination, we hope to extend the reach of researchers to a number of questions that have long been emphasized in the theoretical literature. For instance, researchers could study the relationship between transportation costs and the intensity of competition, the welfare e ects of spatial price discrimination, or the proper construction of antitrust markets. Though our empirical application is static, the estimator also could be used to define

References

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 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 Sta 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.

Gallant, A. R., H. Hong, and A. Khwaja (2010). Dynamic entrytrytrycrosstandariu6(w)27(a)-53

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.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurchner. *Econometrica 55*, 993{1033.

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.

$\substack{T \\ t=1}$

A Proofs

Proof of Lemma 1: The proof is by contradiction. Suppose that S(H( )) is not continuously di erentiable at some parameter vector $_1 \in \Theta$, i.e., that S(H( ))=@ $'$ is discontinuous at $_1$. Then, by the linearity of S and the definition of discontinuity,

$$\lim_{*\to\ _1^-} \frac{@H(\ )}{@\ '}\bigg|_{=\ *} \neq \lim_{*\to\ _1^+} \frac{@H(\ )}{@\ '}\bigg|_{=\ *}$$

However, the function f(p;) is continuously di erentiable in p and by the assumptions placed on q(p;) and c(p;). Thus, for the arbitrary price vector H( ), it follows that @f(H( );)=@ $'$ is continuous, i.e.

$$\lim_{*\to\ _1^-} \frac{@f(p;\ )}{@H(\ )}\frac{@H(\ )}{@\ '}\bigg|_{=\ *} + \frac{@f(p;\ )}{@\ '}\bigg|_{=\ *} = \lim_{*\to\ _1^+} \frac{@f(p;\ )}{@H(\ )}\frac{@H(\ )}{@\ '}\bigg|_{=\ *} + \frac{@f(p;\ )}{@\ '}\bigg|_{=\ *};$$

and it is straight-forward to show (e.g., via log transformation) that this implies

$$\lim_{*\to\ _1^-} \frac{@H(\ )}{@\ '}\bigg|_{=\ *} = \lim_{*\to\ _1^+} \frac{@H(\ )}{@\ '}\bigg|_{=\ *}:$$

□

Proof of Theorem 1: The needed regularity conditions are:

i) $\frac{1}{T}\ \ {}_{t=1}^{T}[p_t^d - S(H(z_t;\ ))] \to^p E[p_t^d - S(H(z_t;$

Table 3: 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 [excludes duties and grinding costs]				
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.

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. Finally, imports are less expensive than domestically produced portland cement. This discrepancy exists for two reasons: First, imports typically come in the form of clinker, which observes water from the air more slowly than cement. The clinker is ground into cement only after it clears customs. The import price does not include the grinding cost. Second, the import price does not include tari s and duties, which are substantial. We include the import dummy in the demand specification to adjust for these factors.

librium. We use the following constraints: the price and distance coe cients ($_1$ and $_2$)

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.