Competitive Price Discrimination in a Spatially Differentiated Intermediate Goods Market*

Walter Beckert* Howard Smith† Yuya Takahashi‡

April 19, 2015

Abstract

Intermediate product markets are distinct in several ways, including the large size of transactions and the ability to price discriminate using buyer-specific prices. We study price determination in an intermediate goods market where products are differentiated by geographic locations of sellers/buyers as well as product characteristics. Using a rich dataset of transactions from the UK brick industry, we estimate a bargaining model in which prices are negotiated between the buyer and seller for each transaction. We analyze the effect of bargaining power, location, and transaction size on prices. In a counterfactual analysis, we measure the welfare impact of price discrimination based on the size of transaction and location of the buyer.

Preliminary and Incomplete. Comments welcome.

*We thank Nathan Miller and conference participants of the DC IO Conference, and seminar participants at ESSEC, Helsinki Center of Economic Research and Penn State for their helpful comments and suggestions. We are grateful to the John Fell Fund at Oxford University for financial assistance. *Birkbeck College, University of London †Department of Economics, Oxford University ‡Johns Hopkins University. We are grateful to a number of people for research assistance: David Elliott, Jack Rowse, Stephen Tozer, and Ben Vanderspar.

1 Introduction

Intermediate product markets are distinct from final product markets because of the greater sophistication of the buyers, the large size of the transactions, and the prevalence of price discrimination using buyer-specific prices. Price discrimination often depends on the volume sold or the location of the buyer. These features of intermediate goods markets are reflected in competition policy: the protection of small downstream firms was the motivating factor behind the Robinson-Patman Amendments (to Section 2 of the Clayton Act) in the US, which prohibit price discrimination, while discrimination between buyers based on their location has been a recurrent issue (for products with high weight-to-value ratio) in Europe and the US. Despite its importance for public policy, there has been very little empirical analysis of intermediate goods pricing.

We study intermediate pricing empirically using a unique dataset of transactions for intermediate products. We use data from the UK brick industry and focus on sales to large construction firms. Each transaction requires the bricks to be delivered to a specific location (building site) so that the choice set varies by transaction. Given the bulky nature of the product, transport costs are important and we exploit the exogenous variation in buyer size and location. We consider demands by the large national housebuilders, who negotiate directly with the brick manufacturers on an order-by-order basis. The data comprise transactions and cost information from the largest brick manufacturers over the period 2001-06. There are more than 2 million transaction records, containing prices, volumes, brick characteristics, manufacturing plant location, and delivery location. The cost data are monthly over the same period, at the plant level.

We estimate a bargaining model in which prices are negotiated between the buyer and seller specifically for each transaction. Our model nests as a special case a model in which the buyers have no negotiating power (similar to Thisse and Vives, 1988) and sellers set prices to each buyer individually de-

pending on their spatial location and order size (i.e., Bertrand-Nash model). We analyze the effect of location and transaction size on the prices that are negotiated. We estimate the importance of buyer size and location effects on the distribution of prices.

Our estimation results suggest that the data reject the Bertrand-Nash model against our bargaining model. Using the estimated parameters, we perform counterfactual analyses to highlight the importance of competition and location effects. Specifically, we solve the model with the following changes to the environment: (i) eliminating joint ownership of plants, (ii) setting the transport cost parameter to zero, thereby eliminating geographic differentiation, and (iii) banning price discrimination based on the size of transaction and location of the buyer. In the first scenario, we find that the largest seller enjoys significant market power, as prices fall by 10.6% when there is a de-merger of its plants. This is approximately 20% of the average markup. In the second scenario, for all four sellers the price increases when transport costs go to zero. One possible reason for this is that firms that are located further from consumers are able to raise their prices when transport costs are eliminated. In the third scenario, we find that the uniform price restriction results in higher equilibrium prices, but sellers' profits are reduces because their market shares significantly drop. Overall, the total surplus decreases by 23.6% due to the uniform price restriction. Most of this decrease in the total surplus comes from the loss of market shares of inside products and misallocation among inside products does not play a major role.

There is now a large theoretical literature on intermediate goods pricing (see surveys by Katz, 1989; Rey and Tirole, 2007). Some models use a leader-follower interface between sellers and buyers, in which the downstream firms are price-takers (e.g. Rey and Tirole, 1986; Katz, 1987). Others use a bargaining interface (see Dobson and Waterson, 1997; Chipty and Snyder, 1997; Inderst and Wey, 2007; Chen, 2003; Smith and Thanassoulis, 2012; de Fontenay and Gans, 2013). Many of these models feature prices set specifi-

cally to individual buyers. Of particular interest to us is discrimination by volume sold or location of the buyer, as studied in Katz (1987), Inderst and Valetti (2009), and Thisse and Vives (1988).

Compared to the theoretical literature, there is relatively little empirical literature on intermediate goods pricing. Those papers that have studied intermediate prices directly are in three groups. The first uses regression analysis rather than structural modelling (see Ellison and Snyder, 2010; Sorensen, 2003); these papers establish the existence of buyer power, particularly in cases when there is upstream competition. The second group studies spatial competition and market power arising from geographic differentiation (see Houde, 2012; Chicu, 2013; Miller and Osborne, 2014); they provide evidence that transport (travel) costs and consumer location are important determinants of demand and market power exercised by producers. A third group of papers estimate structural bargaining models, typically assuming bilateral oligopoly (see Crawford and Yurucoglu, 2010; Grennan, 2013; Gowrisankaran, Nevo and Town, 2013). These all use a specific bargaining model: the "Nash-in-Nash" bargaining solution proposed in Horn and Wolinsky (1988). Our paper develops this approach to allow greater focus on price dispersion in bargaining equilibria, most notably due to buyer level location and size differences.

In Section 2 we discuss the industry. In Section 3 we develop the theoretical model of price determination, and derive the likelihood functions for estimation of parameters in Section 4. In Section 5 we describe the data. Section 6 discusses results and performs several counterfactual analyses. Section 7 concludes.

2 The Market for Bricks

The main use for bricks in Great Britain is as an external cladding material for houses.[1,2] This use accounts for 80%-90% of all bricks produced. House-building firms have some alternatives to bricks as options for external cladding materials, such as timber, stone, and plaster. The decision to use bricks, as opposed to one of these alternatives, depends partly on aesthetic considerations and partly on the cost of the bricks. If the builder decides to use bricks, the total volume required for any building project is determined by the number of houses or apartments in the project. When bricks are used they are on average only a small proportion, about 3%, of the overall cost of the buildings.

There are 48 brick manufacturing plants in Great Britain, owned by four main firms. Brick production consists of three main steps: extracting clay from the ground; grinding, shaping and drying the bricks; and then firing them in kilns at temperatures of 1000 degrees centigrade. The main costs of production are labor and gas; the latter comprises 17%-26% of production costs.

Brick plants are always located next to naturally occurring clay deposits, as any quantity of bricks is cheaper to transport than the corresponding quantity of clay input. The bricks produced by any plant are always made from local clay, which in turn affects the appearance and in particular the shade of the brick. Thus, for example, a red brick produced at two diffrent plants are a different shade of red, as a result of the different clay at each plant, and a buyer seeking a red brick may prefer one plant to another for this reason.

There is also heterogeneity in marginal costs across plants due to the

[1] For a detailed discussion of the brick market see Competition Commission (2007), which is the source for the evidence presented in this section.

[2] To preserve the anonymity of firms we do not reveal the identities of the individual house builders or brick manufacturers.

extent of robotic systems and more energy-efficient kilns.There is very limited entry and exit of plants: for example there was no new entry of plants during the six years of our data (2001-2006), though a few plants closed in this period because firms had too much capacity.

There are significant inventories of bricks stored at each plant. At any point in time stocks are equivalent to about one third of the annual flow of production. Manufacturers hold brick inventory to allow an order to be supplied quickly from stock and to allow smoothing of production relative to demand, which is at its greatest between March and September when housebuilding is most active. House building firms generally require a just-in-time delivery and therefore rely on a manufacturer's ability to hold sufficient stock and deliver bricks at short notice.

Geographic location is very important in the brick market. Delivery costs constitute about 25% of the cost of delivered bricks, including costs of loading and unloading. As a result, brick deliveries tend to be quite localized: about 50% of bricks are sold within a 110 km radius of the plant, and about 80% within a 200 km radius. Figure 1 shows the locations of the plants, where the largest circles drawn have a radius of 200 km. Delivery of bricks to the constriction site is carried out by third party hauliers using purpose-built vehicles, arranged either by the manufacturer or the buyer.

The manufacturers offer a range of brick products that vary in a number of characteristics. Some are relevant for aesthetics, such as the color and the manufacturing method. Others affect the performance of the brick. A more detailed discussion of brick characteristics is in Section 5. Typically the color of brick for any housing project is determined in plans that must be submitted to the local authority for planning permission, well before the choice of brick manufacturer is determined; local authority planners may stipulate brick cladding of a particular color to fit with local conditions.

The major house builders source their bricks directly and do not buy from intermediaries such as brick retailers. Table 1 shows that house building is

a concentrated market—the top 15 firms account for 85% of volumes sold to house building firms—and that the size of these buyers varies greatly. In this paper we will model the transactions involving the top 20 house bulding firms.

The table also shows that most builders buy from more than one seller. While some major developers deal with all the four sellers, a few place the great majority of their business with only one of them. This could be because of the geographic match between the buyer and the seller, or a preference to buy from some firms rather than others. There is no sign of switching costs: switching between manufacturers is relatively easy, with no contractual barriers. For any given project location, however, the buyer typically only sources from a single firm. Developers do not change brick supplier part way through a development, but they will source from different manufacturers for different house building developments.

Prices are negotiated with buyers individually, with prices depending on the volume, the historic relationship with the manufacturer, buyer size, and distance from plant. Manufacturers frequently vary prices for locations further from the plant to compete for business with more local plants. Some buyers agree annual "framework agreements" that set out a detailed matrix of prices for different brick products at different locations. However, there is generally no firm commitment to buy at these prices.

3 A Model of Price Formation

We model price formation at the level of the individual brick transaction. We define a brick transaction as a unique combination of brick product, plant, buying firm, delivery location, and year. Thus if several deliveries of any product in any year are made from the same plant to the same construction site, this counts as a single transaction. A given buying firm is therefore associated with many transactions in any year. We assume these transactions

are conducted independently, so in the notation that follows it is convenient to use the same index (i) for "buyer" and "transaction".[3]

3.1 Utility and Cost

A buyer i has a construction project whose scale q_i and location are determined outside the model. The buyer must choose a cladding material j for the project, which may be a brick product or a non-brick cladding material. We let $j = 0$ represent the outside option of a non-brick cladding choice.

Each brick product j has a unique plant $a(j)$ and manufacturer $g(j)$. The distance from plant a and buyer i's delivery location is $d_{ia(j)}$. The color of bricks $c(i)$ is determined for buyer i by the local authority planner. Conditional on this colour, the choice set for buyer i is $\mathcal{J}_i \equiv \mathcal{J}_{c(i)}$. The products in this choice set that are offered by seller g are denoted $\mathcal{J}_{ig}$, so that $\mathcal{J}_i = \bigcup_g \mathcal{J}_{ig}$.

Suppose that buyer i makes the payment T_{ij} to pay for the brick product j. Then we assume that the indirect net utility u_{ij} of buyer i from product j is given by

$$u_{ij} = \left(\beta_j - \tau d_{ia(j)} + \epsilon_{ij}\right) q_i + \lambda_{ig(j)} - \alpha T_{ij} \tag{1}$$

where α is the marginal utility of money, β_j is mean utility per-unit of product j, and ϵ_{ij} is an unobserved per-brick taste disturbance from this mean. We assume linear transportation costs, with parameter τ. The transport costs are proportional to volume q_i, as the number of trucks needed to deliver the bricks is approximately proportional to the volume of the bricks.

We assume that the buyer has a net fixed utility $\lambda_{ig(j)}$ of transacting with each seller that is independent of the volume of bricks in the transaction. This captures the idea that a buyer may prefer to deal with some sellers rather than others because of, for instance, historical buying relationships. For a relatively large order (a large value of q_i), this effect is of lower importance

[3] For a builder, bricks are on average only about 3% of the overall cost of buildings for developers, and it is reasonable to assume that a change to the terms that one house builder agrees with a manufacturer does not have externalities on other buyers.

relative to the effects that appear inside brackets in equation (1).

We allow the consumer to have a per-unit disturbance ϵ_{i0} to their utility from the outside option and normalize the mean utility of the outside good to zero. The utility from the outside option is therefore $u_{i0} = (0 + \epsilon_{i0})\, q_i$. This is the utility that the buyer gets per unit from a non-brick choice of cladding for the buildings.

The total cost of production at plant a in period t is given by

$$C_{at} = C_a(Q, w_{at}) \tag{2}$$

where w_{at} is the factor prices faced by plant a at time t and Q is the total quantity of bricks produced at plant a. We assume that all brick products produced at plant a have the same cost of production.

The per-unit cost of producing quantity q_i from plant a is derived from the total plant-level cost function, assuming total output Q, is as follows:

$$\frac{C_{at}(Q, w_{at}) - C_{at}(Q - q_i, w_{at})}{q_i}.$$

As q_i is very small relative to total production Q we approximate the per unit cost of producing q_i by the marginal cost, i.e.

$$c_{a(j)} = \frac{\partial C_{at}(Q, w_{at})}{\partial Q}.$$

The q_i units of supply required for buyer i also incur a per-transaction fixed cost $F_{a(j)}$. The cost term $F_{a(j)}$ is a per-transaction fixed cost, incurred for each transaction. Thus the total cost of supply for buyer i's q units is given by

$$c_{a(j)} q_i + F_{a(j)}. \tag{3}$$

The seller supplies the buyer's demand from the inventory at the plant, so that $c_{a(j)}$ is therefore assumed to be the marginal cost of production of the

plant when operating at average total output levels for the year. As discussed in Section 2, there is sufficient inventory for us to assume that the seller can treat the cost of supply of transaction i as being independent of other transactions. Examples of activities that cause a per-transaction cost include the labor time costs of contracting with the seller as well as loading costs that are incurred by the seller.

The profit that the seller receives from the transaction is

$$\pi_{ij} = T_{ij} - c_{a(j)} q_i - F_{a(j)}$$

The joint surplus that is generated when i buys from j is given by adding net utility and profits:

$$S_{ij} = u_{ij} + \alpha \pi_{ij} \equiv \left(\beta_j - \tau d_{ia(j)} - \alpha c_{a(j)} + \epsilon_{ij}\right) q_i - \alpha F_{a(j)} + \lambda_{ig(j)},$$

where we have scaled profits using the marginal utility of money α. The socially efficient match is the choice of j that generates the highest surplus S_{ij}.

3.2 Product Choice and Equilibrium Price

We assume that all players have complete information, including knowedge of the tastes $\epsilon_i = \{\epsilon_{ij}\}_{j \in \mathcal{J}}$ for all i.

We consider a model that allows buyers to have bargaining power, but which nests the case where the buyer has no bargaining power. When the buyer has no bargaining power, sellers simultaneously post buyer-specific prices, and there is an implied Bertrand-Nash equilibrium. When endowed with bargaining power, the buyer is able to reduce the price below this Bertrand-Nash level towards the seller's cost of supply.

We define j^* and j^{**} as follows: j^* is the product that maximizes social surplus. j^{**} is the product that maximizes social surplus if all products

produced by the seller producing j^* are eliminated, i.e.

$$j^* = \arg\max_{j' \in \mathcal{J}_i} S_{ij'} \qquad\qquad j^{**} = \arg\max_{j' \in \mathcal{J}_i \setminus \mathcal{J}_{g(j^*)}} S_{ij'}.$$

We call these the first and second products.

Consider first the special case, nested within our general model, where the buyer has no bargaining power. It is a Nash equilibrium for the buyer to buy product j^* and for prices to be as follows: the seller of j^{**} sets price to incremental cost and the payment T for product j^* is given by the solution to the following indifference condition:

$$\begin{aligned} & \left(\beta_{j^*} - \tau d_{ia(j^*)} + \epsilon_{ij^*}\right) q_i + \lambda_{ig(j^*)} - T \\ = \; & \left(\beta_{j^{**}} - \tau d_{ia(j^{**})} - c_{a(j^{**})} + \epsilon_{ij^{**}}\right) q_i + \lambda_{ig(j^{**})} - F_{a(j^{**})}. \end{aligned}$$

This condition equates the utility of the buyer buying q_i units of j^{**} at T to the utility it would get if it bought j^{**} and paid the incremental cost of production. If the indifference condition did not hold then seller of j^* could raise the price and still sell the product. When the price of j^* is as defined by this equality, and the the seller cannot gain by deviating from the price. The other seller cannot gain by raising price or by cutting it below marginal cost. Therefore this is an equilibrium.

We denote the equilibrium payment $T_{ij^*|j^{**}}$ which recognises that product j^{**} is relevant. Rearranging $T_{ij^*|j^{**}}$ in terms of price per unit $p_{ij^*|j^{**}}$ gives

$$\begin{aligned} p_{ij^*|j^{**}} &= \frac{T_{ij^*|j^{**}}}{q_i} \\ &= \left(\beta_{j^*} - \beta_{j^{**}}\right) + \tau\left(d_{ia(j^{**})} - d_{ia(j^*)}\right) + \epsilon_{ij^*} - \epsilon_{ij^{**}} \\ &\quad + c_{a(j^{**})} + \frac{\left(\lambda_{ig(j^*)} - \lambda_{ig(j^{**})}\right) + F_{a(j^{**})}}{q_i}. \end{aligned}$$

The Bertrand-Nash price per brick in this equation is increasing in (i) the

mean utility difference between the two products $\left(\beta_{j^*} - \beta_{j^{**}}\right)$, (ii) the extra distance to the second best plant $(d_{ia(j^{**})} - d_{ia(j^*)})$, (iii) the cost per unit of the second best product $c_{a(j^{**})} + \frac{F_{a(j^{**})}}{q_i}$ and the difference in fixed preference for the two sellers $\frac{(\lambda_{ig(j^*)} - \lambda_{ig(j^{**})})}{q_i}$.

The extent to which an increase in volume q_i affects unit price depends on the size of the fixed effects $(\lambda_{ig(j^*)} - \lambda_{ig(j^{**})}) + F_{a(j^{**})}$.

We now discuss the general version of our model, which nests this Bertrand-Nash special case, and allows the buyers to have some active role in price determination, provided they have bargaining skill.

From seller g's set $\mathcal{J}_{ig}$ we assume that i would choose to transact product j^* as defined above as this offers the highest joint surplus. The payoffs for the buyer and seller, respectively, from transfer payment T, are given by

$$\left\{ \left(\begin{array}{c} \left(\beta_{j^*} - \tau d_{ia(j^*)} + \epsilon_{ij^*}\right) q_i \\ +\lambda_{ig(j^*)} - T \end{array} \right), \left(T - c_{a(j^*)} q_i - F_{a(j^*)}\right) \middle| \, T \geq 0 \right\}$$

and the disagreement points to each agent are as follows

$$\left(\left(\begin{array}{c} \left(\beta_{j^{**}} - \tau d_{ia(j^{**})} - c_{a(j^{**})} + \epsilon_{ij^{**}}\right) q_i \\ +\lambda_{ig(j^{**})} - F_{a(j^{**})} \end{array} \right), 0 \right)$$

where product j^{**} is as defined above. The assumption in the buyer's disagreement point is that the buyer can always obtain the product at the Bertrand-Nash price from the second best seller as a disagreement point in negotiations with the best seller. The seller makes no profit if it sells no units to buyer i so its disagreement point is zero.

The Nash product between buyer i and seller $g(j^*)$ for product j^*, given

next-best product j^{**}, is

$$N_{ig(j^*)|j^{**}}(T) = \left[\begin{array}{c} \left(\left(\beta_{j^*} - \tau d_{ia(j^*)} + \epsilon_{ij^*}\right) q_i + \lambda_{ig(j^*)} - T\right) \\ -\left(\begin{array}{c} \left(\beta_{j^{**}} - \tau d_{ia(j^{**})} - c_{a(j^{**})} + \epsilon_{ij^{**}}\right) q_i \\ +\lambda_{ig(j^{**})} - F_{a(j^{**})} \end{array} \right) \end{array} \right]^{1-\theta_i} \times \left[\left(T - c_{a(j^*)} q_i - F_{a(j^*)}\right)\right]^{\theta_i}, \tag{4}$$

where θ_i is the bargaining power of the seller against buyer i. The buyers are allowed to have different bargaining strengths to reflect their heterogeneous size and sophistication.

The solution to this maximization problem is given by

$$T_{ij^*|j^{**}} = \left[\begin{array}{c} (1-\theta_i)\left(c_{a(j^*)} q_i + F_{a(j^*)}\right) \\ +\theta_i \left(\begin{array}{c} \left(\beta_{j^*} - \beta_{j^{**}} - \tau\left(d_{ia(j^*)} - d_{ia(j^{**})}\right) + c_{j^{**}} + \epsilon_{ij^*} - \epsilon_{ij^{**}}\right) q_i \\ +F_{a(j^{**})} + \left(\lambda_{ig(j^*)} - \lambda_{ig(j^{**})}\right) \end{array} \right) \end{array} \right] \tag{5}$$

When $\theta_i = 0$ the buyer has all the bargaining power and can push the price to the seller's cost. When $\theta_i = 1$ the buyer has no bargaining power and in this case we can see from (5) that the buyer obtains exactly the same transfer price $T_{ij^*|j^{**}}$ as the Bertrand-Nash case above.

It is useful to derive the price per brick, p_{ij}:

$$\begin{aligned} p_{ij^*|j^{**}} &\equiv \frac{T_{ij^*|j^{**}}}{q_i} \\ &= \left[\begin{array}{c} (1-\theta_i)\left(c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i}\right) \\ +\theta_i \left(\begin{array}{c} \left(\beta_{j^*} - \beta_{j^{**}} - \tau\left(d_{ia(j^*)} - d_{ia(j^{**})}\right) + c_{j^{**}} + \epsilon_{ij^*} - \epsilon_{ij^{**}}\right) \\ +\frac{F_{a(j^{**})} + \left(\lambda_{ig(j^*)} - \lambda_{ig(j^{**})}\right)}{q_i} \end{array} \right) \end{array} \right] \end{aligned} \tag{6}$$

To understand how the negotiated price varies depending on the buyer, consider the expression for the price per brick (6). We note that a buyer with a low value for θ_i (and hence a high bargaining power) attains a lower

price, other things equal. However there are some other variables that determine the price, conditional on any value for θ_i. First consider the effect of volumes q_i. If fixed costs per transaction $F_{a(j^*)} > 0$ and $F_{a(j^{**})} > 0$ then there will be buyer discounts, because the average cost of supply falls as q_i increases, and the buyer is able to appropriate some of this. In a similar way, if $\left(\lambda_{ig(j^*)} - \lambda_{ig(j^{**})}\right) > 0$, then provided $\theta_i > 0$ there will be further buyer discounts, because the average benefit of supply from a preferred supplier falls as q_i increases, so the favorite seller is able to appropriate less per unit. Second, consider the effect of distance $d_{ia(j)}$. Provided $\theta_i > 0$, buyers who have to go a greater distance $\left(d_{ia(j^*)} - d_{ia(j^{**})}\right)$ to get to the second best plant will pay a higher price as the seller can extract some of the surplus it generates because of its favorable location.

We now explain why the buyer chooses product j^*. If we substitute (6) into (1) and rearrange, we can show that the utility of buyer i is given by

$$u_{ij^*|j^{**}} = (1 - \theta_i)S_{ij^*} + \theta_i S_{ij^{**}}.$$

We now ask if the buyer i can do any better than this. Let us suppose that the buyer chooses product j to negotiate over and chooses product k as its disagreement product (i.e. the product it would buy at a Bertrand-Nash price if the bargaining broke down). Then the utility of the buyer is given by

$$u_{ij|k} = (1 - \theta_i)S_{ij} + \theta_i S_{ik}. \tag{7}$$

It is natural to impose that the buyer selects j and k subject to $S_j > S_k$, i.e. the product k chosen as a disagreement point offers less surplus than the product j the consumer wishes to buy. Then from (7) we see that for any given k the consumer will maximize $u_{ij|k}$ if it chooses

$$j = \arg\max_{j' \in \mathcal{J}_i} S_{ij'}.$$

4 Estimation

4.1 Step 1: Cost Function Estimation

The cost function (2) is specified as follows

$$\ln C_{at} = \gamma_{0a} + \gamma_1 \ln Q_{at} + \gamma_2 \ln G_t + \gamma_3 \ln W_{at} + \gamma_D' D_t + \eta_{at}, \tag{8}$$

where C_{at} is the total cost of plant a at time t, Q_{at} is the total production of bricks in plant a at time t, G_t is the (national) price of natural gas at time t and W_{at} is regional wage data at time t for the region of plant a. D_t is a quarterly dummy and η_{at} is unobserved cost. To allow for the possible endogeneity of the quantity variable Q, we use instrumental variables. The instrumental variables are demand shifters which affect quantity demanded but not costs. The demand shifters are (i) the number of new houses that builders have started to build (housing starts) and (ii) number of houses completed (completions) in the region of the plant $a(j)$ for that quarter. These are available from official government sources. We allow for plant fixed effects which control for unobserved productivity heterogeneity across plants.

4.2 Step 2: Transaction Estimation

In this step we use the implied marginal cost from the estimated cost function in the first step everywhere $c_{a(j)}$ appears in the model. In the discussion we will treat $c_{a(j)}$ as though it is observed.

We have N observations of transactions. Each transaction consists of two observed components: a chosen product j and a transaction price per unit at which the chosen product was sold, p_i. We also observe the brick characteristics x_j, for all $j \in \mathcal{J}_i$, the distance to all the plants $d_{ia(j)}$, the transaction volume q_i, and the seller associated with the chosen product $g(j)$.

4.2.1 Parameterization

We assume that the mean utility β_j for product j produced in plant $a(j)$ with characteristics x_j and colour $c(j)$ is given by

$$\beta_j = \beta_{a(j)}^{c(j)} + \beta x_j.$$

The term $\beta_{a(j)}^{c(j)}$ picks up differences between plants in the quality of products, for each colour c.

We assume that there is a buyer-seller pair specific term $\lambda_{ig(j)}$ in the buyer's utility function. It is difficult to estimate this parameter for every single pair, so we specify that

$$\lambda_{ig(j)} = \lambda_1 \varkappa_{ig(j)} + \lambda_2(1 - \varkappa_{ig(j)}),$$

where $\varkappa_{ig(j)}$ is a binary indicator variable for whether firm g is one of buyer i's regular suppliers.[4] This allows the buyer to prefer to deal with some sellers rather than others because of historical buying relationships, so differences in $\lambda_{ig(j)}$ reflect the cost of using an unfamiliar business relationship. We also assume that the per-transaction fixed cost F is seller-specific, rather than plant-specific.

We define the percentage difference between the observed transaction price and the predicted transaction price as the prediction error ν_{ij^*}. We assume that ν_{ij^*} follows a normal distribution with mean zero and variance σ_v^2.[5] In addition, we assume that ϵ_{ij} is IID according to a Type-1 Extreme Value distribution with the scale parameter σ_ϵ. We estimate both σ_v and σ_ϵ.

As was discussed in Section 2, the color of the brick is set by the building specification plans. We focus on two different colors, red and buff. Across these two colors, the market shares of plants are very different. This is

[4] This is defined as sellers that contribute more than 10% of the buyer's transactions.

[5] This prediction error involves various components. We approximate its distribution by this normal distribution. Details of this approximation are provided in the Appendix.

probably because some plant is better at producing bricks with a particular color than others. To capture this, we estimate plant-level fixed effects for each color separately.

Thus, the parameters to be estimated in our bargaining model are

$$\{\lambda_1, \lambda_2, \tau, \{\beta_a^{red}\}_{a=1}^{A}, \{\beta_a^{buff}\}_{a=1}^{A}, \{\beta_h\}_{h=1}^{H}, \{F_g\}_{g=1}^{4}, \{\theta_i\}_{i=1}^{20}, \sigma_v, \sigma_\epsilon\},$$

where H is the number of characteristics in vector $x_j = (x_{j1}, .., x_{jh}, .., x_{jH})$ and A is the total number of plants.[6]

4.2.2 Likelihood Function

We assume that all transactions are independent. The likelihood contribution of a transaction is given by the joint probability of the price and the choice of product. This is given by the density of the prediction error ν_{ij^*} conditional on choice of j^* multiplied by the probability that j^* maximizes the surplus from the transaction:

$$f_{\nu_{ij^*}}(\nu_{ij^*}|j^* = \arg\max_{j' \in \mathcal{J}_i} S_{ij'})\mathcal{P}_{ij^*}, \tag{9}$$

where $\mathcal{P}_{ij}$ is the probability that buyer i chooses seller j, defined as

$$\begin{aligned} \mathcal{P}_{ij} &= \Pr\left[j = \arg\max_{j' \in \{0, \mathcal{J}_i\}} S_{ij'}\right] \\ &= \Pr\left[\frac{S_{ij}}{q} \geq \frac{S_{ij'}}{q}, \forall j' \in \{0, \mathcal{J}_g\}\right]. \end{aligned}$$

[6] There are several plants that do not produce bricks with particular colors. The total number of β_a parameters is 77. The number of observable characteristics (H) is 6 including the constant term. Thus, the total number of parameters estimated is 112.

Letting

$$
\begin{aligned}
\tilde{S}_{ij} &= \frac{S_{ij}}{q_i} = \left(\beta_{g(j)} + \beta x_j - \tau d_{ia(j)} - \alpha c_{a(j)}\right) + \frac{\lambda_{ig(j)} - \alpha F_{a(j)}}{q_i} + \epsilon_{ij} \\
&= \delta_{ij} + \epsilon_{ij},
\end{aligned}
$$

we can write the choice probability as follows:

$$
\mathcal{P}_{ij} = \frac{\exp(\delta_{ij}/\sigma_\epsilon)}{\sum_{j' \in \{0, \mathcal{J}_i\}} \exp(\delta_{ij'}/\sigma_\epsilon)},
$$

where

$$
\delta_{ij} \equiv \beta_{g(j)} + \beta x_j - \tau d_{ia(j)} - \alpha c_{a(j)} + \left(\lambda_{ig(j)} - \alpha F_{a(j)}\right) / \, q_i
$$

and $\delta_{i0} = 0$.

Next, to compute the density of the observation error ν_{ij^*} conditional on choice of j^*, we derive expressions for the expected transaction price conditional on choice of product. From (6), it is easy to show that

$$
p_{ij^*|j^{**}} = c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \frac{\theta_i \left[S_{ij^*} - S_{ij^{**}}\right]}{q_i}.
$$

One difficulty is that the price depends on the buyer's second best option j^{**}, which is not observed. However, by integrating out the runner up j^{**}, we can obtain the following closed-form solution (see Brannman and Froeb, 2000):

$$
\begin{aligned}
E\left[p_{ij^*}\right] &= c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i E\left[\widetilde{S}_{ij^*}\right] - \theta_i E\left[\widetilde{S}_{ij^{**}} | j^*\right] \\
&= c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \left(\delta_{ig(j^*)} - \sigma \ln\left(\mathcal{P}_{ig(j^*)}\right)\right) - \theta_i \left[\delta_i^* + \sigma \frac{\ln\left(1 - \mathcal{P}_{ig(j^*)}\right)}{\mathcal{P}_{ig(j^*)}}\right] \\
&= c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \left[\delta_{ig(j^*)} - \delta_i^* - \sigma \ln\left(\mathcal{P}_{ig(j^*)}\right) - \sigma \frac{\ln\left(1 - \mathcal{P}_{ig(j^*)}\right)}{\mathcal{P}_{ig(j^*)}}\right],
\end{aligned}
$$

where

$$\delta_i^* = \sigma \ln \left(\sum_{g=0}^{4} \exp\left(\frac{\delta_{ig}}{\sigma}\right)\right)$$

and

$$\delta_{ig} = \sigma \ln \left(\sum_{j\in\mathcal{J}_{ig}} \exp\left(\frac{\delta_{ij}}{\sigma}\right)\right)$$

and

$$\mathcal{P}_{ig} = \sum_{j\in\mathcal{J}_{ig}} \mathcal{P}_{ij}.$$

We assume that the observed price $p_{ij^*}^{\text{obs}}$ is equal to the predicted price p_{ij^*}, multiplied by a prediction error, denoted by ν_{ij^*}; i.e.

$$p_{ij^*}^{\text{obs}} = p_{ij^*} e^{\nu_{ij^*}}. \tag{10}$$

The prediction error has a density $f_{\nu_{ij^*}}(\cdot)$ with a closed form which results from the distribution of ϵ_{ij^*}.[7] The contribution to the likelihood function of this transaction (9) is therefore given by the density of ν_{ij^*}, evaluated at difference between log of observed and predicted prices, multiplied by $\mathcal{P}_{ij^*}$:

$$f_{\nu_{ij^*}}\left(\ln p_{ij^*}^{\text{obs}} - \ln\left(\begin{array}{c} c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i\left[\delta_{ig(j^*)} - \delta_i^*\right] \\ -\theta_i\sigma\left[\ln\left(\mathcal{P}_{ig(j^*)}\right) + \frac{\ln\left(1-\mathcal{P}_{ig(j^*)}\right)}{\mathcal{P}_{ig(j^*)}}\right]\end{array}\right)\right)\mathcal{P}_{ij^*}.$$

5 Data

The data consist of detailed information on transactions and production costs for the four largest UK brick manufacturers for the period 2001-2006.[8] The sales of these firms are about 85% of the UK brick market.[9]

The transactions data are observed at the level of the individual delivery of bricks to the buyer. The observations include information on the price

[7] For this approximation, see the Appendix.

[8] Competition Commission inquiry (CC, 2007) into the merger of two large UK brick manufacturers.

[9] Market shares reported here are as given in CC (2007).

paid, the volumes delivered, brick characteristics, the manufacturing plant and its location, and the delivery location. Bricks bought by house building firms are delivered to the site that the houses are to be built. Each house builder builds houses on multiple locations so for each buyer there are many different delivery locations. The data also include the cost of transporting the bricks between the brick plant and the delivery location. The production cost data are recorded at the month-plant level, and give the operating cost of producing the plant's total output that month.

We supplement this data with other information from public sources. We have two input prices namely the monthly price of natural gas, the main energy used to heat clay to high temperatures, and monthly labor costs for the regional location of each plant. We also use public data on the total number of new houses that began construction in each quarterly time period and county. We have geographic coordinates for the plants and delivery locations which allow us to compute transportation distances for each transaction.

For each transaction we observe all the brick characteristics of the delivered brick.

One of the most important characteristics is the color of the brick. Most bricks in our data are red, about two thirds by volume. The second most common color is yellow (also known as "buff"). As mentioned in Section 2 the color of bricks used is restricted by the local authority government planner to suit local environmental considerations.

The second characteristic that relates to a brick's appearance is its texture, which depends on the manufacturing process. There are two main manufacturing processes. The first uses a mould to shape the clay, which can yield an irregular appearance which is considered appealing. Bricks produced in this way are known as soft mud bricks. The second process uses wires to slice the clay into bricks. Bricks produced in this way are known as wirecut bricks, and these have a much more regular shape.

Other observed brick characteristics relate to technical standards. The

brick's strength measures the load-bearing capacity of the brick (in Newton per mm^2). There are three strength categories. Water absorption—the ability to release and re-absorb moisture (a "breathing" process)—helps to regulate the temperature and humidity of atmosphere in a house. The desirable water absorption for clay bricks is between 12% and 20%, and we observe whether or not a brick falls into this category. Frost resistance is another consideration, and we use three categories: bricks suitable for low, moderate and severe frost exposure.

Each firm offers hundreds of individual brick brands, each with individual brand names such as "Durham Red Multi". For any color of brick we classify brick brands into a smaller number of brick product "types". A brick type is a specific combination of the discrete brick characteristics: manufacturing process (of which there are two possible values), strength (three possible values), and water absorption (three possible values). This classification leads to 18 possible different brick types for any color. Each of the 48 plants produces some subset of these 18 types and either one or both brick colors, resulting in 130 distinct combinations of plant and product brick type.

Using the location of the construction site and plants we can obtain the distances to the plants for any transaction. This allows us to construct choice sets for a given construction project. Each construction project is located at a different point in geographic space which implies a different distance to any plant for each project.

The transactions data are recorded at the level of the individual delivery of a brick brand. We aggregate these deliveries to the annual level and drop outliers for which the price per brick is in the top 1% of the sample, often these very high prices per brick are associated with very small transaction volumes that could be data errors.

As our dataset comprises brick transactions we do not directly observe the number of orders for other (non-brick) types of facing material. However we can indirectly compute a figure for this by assuming that the total market

size, the maximum potential demand for bricks, is determined by the number of new houses being built in any region. There are government statistics on the number of new houses built in each region and period, by type of dwelling (apartment or house). From this we compute the number of bricks that would be demanded if everyone used bricks, using estimates of how many bricks are needed for a typical house. As bricks are only a very small fraction of the cost of building a house we treat the demand for houses as independent of brick prices, so that if the brick price increases people substitute to other cladding choices rather than deciding not to build a house. Table 2 shows for each UK region the computed market size, the volumes of bricks delivered, and the implied market share of the outside option s_0. Figures (in millions) are totals for the period 2001-2006. We can see that market share s_0 varies from region to region. These shares are partly affected by variations across regions in the distance to a brick plant. For regions where there are few local brick plants the market share of other cladding materials is relatively high (an example of such a region is Scotland, as can be seen from Figure 1).

Table 3 describes the cost data. This is at the plant-quarter level. There are 48 plants and 6 years. The variables are: total cost (L), Quantity produced (Q), Gas Price (G), Regional Earnings (G). We can see that average cost per 1,000 bricks is about £157. We use these data to estimate a total cost function, from which marginal costs are derived.

Before estimating the model, we provide several raw data analyses. Figure 2 plots a density estimate for the transaction prices (the price per 1,000 bricks). The price variation is large, with the 90% of prices ranging from 142 to 250. This figure also shows the substantial heterogeneity among sellers.

Next, we restrict the sample to a few brick types that are most common in our data set.[10] Figure 3 plots the distribution of prices for this subsample. As expected, the variance of the price decreases significantly, but even within

[10]We choose red extruded/wirecut bricks with the medium-level water absorption. The number of observations is 5,143.

the same type, the transaction price still has a large variation. The 90% of prices ranges from 144 to 222.

Finally, to control for price differentials according to the transaction size, Figure 4 plots the distribution of prices for only transactions whose volume lies between 45th and 55th percentiles of the sample used in Figure 3.[11] The 90% of prices ranges from 143 to 213, and thus we confirm that the large part of the price variation remains.

To understand how the transaction price is determined we now look at the transactions data. Table 4 presents descriptive statistics. We use these data to regress the price per 1,000 bricks (in £100) on the volume, buyer size (defined as the total volume purchased by the buyer as a proportion of all volume), manufacturer dummies, plant fixed effects, other product characteristics, and on some spatial competition variables. Table 5 summarizes the results. The first specification only includes volume, and buyer size. This confirms the existence of a significant discount that depends on the volume demanded of bricks. If the number of bricks increases by 1,000 units, then the unit price decreases by £1.54. In addition, the size of the buyer is negative and significant. That is large buyers on average appear to buy at lower prices regardless of the size of their individual transaction.

The second specification includes plant fixed effects and controls for product characteristics. The fit of the regression measured by the adjusted R^2 shows that plant- and product-level characteristics are important explaining the price variation. The magnitude of the size discount becomes smaller, but it is still economically important and statistically significant. Marginal costs are included here, and are (unsurprisingly) an important determinant of price.

The third and fourth specifications include variables which aim to establish importance of spatial competition in price setting. For each transaction,

[11] The 45th and 55th percentiles of the sample volume are 39,600 and 58,000, respectively. The number of observations is 516.

we can identify the seller. Associated with the seller, we count the number of plants within 200km of the buyer that are owned by other competing sellers. This variable is "competition 200km" and is used in the third specification. We find that the parameter on this variable is significant and negative. This indicates that local competition puts a downward pressure on the transaction price.

We also calculate the distance between the nearest plant and the second nearest plant (dist diff). This variable is included in the fourth specification and it has a positive and significant parameter. This suggests that if the competing plants are located farther away, the seller can charge a higher price to the buyer. The table suggests upstream competition and location are part of what determines price setting in the industry.

6 Results

6.1 Parameter Estimates

Table 6 presents the cost parameters estimated in the first step. The signs on the gas price (G) and wage (W) parameters are as expected. The parameter on Q indicates that there are diminishing returns to scale. The parameter estimate on Q in the IV estimation with fixed effects is not very different from the OLS parameter, suggesting that there is no bias on the quantity parameter in the OLS results. For the second stage estimation, we use the third set of estimates which include fixed effects.

Table 7 and Figure 5 present the markups implied by the cost estimates. The markups are rarely negative, as we expect, and show considerable price dispersion after allowing for differences in production costs c.

Table 8 presents the estimates of selected parameters of utility and cost functions. The distance coefficient has the correct sign (i.e. a positive transport cost) and is statistically significant. The magnitude of the coefficient implies transportation costs of £110 (110 UK Pounds) per 1,000 bricks for

100 km. Since the average price is £192 per 1,000 bricks, haulage costs are 36% of the total cost to the buyer of delivered bricks (=110/(110+192)), assuming that the average delivered distance is around 100km. Our data include haulage costs too, so we use this to check how reasonable our estimates are. The direct estimate from the data is £80 per 1,000 bricks for 100 km, which implies that our model slightly overestimates the haulage costs. Furthermore, CC (2007) reports that haulage costs are about 25% of the overall cost to the buyer of delivered bricks. Our estimate suggests that haulage costs are about 36% of the overall costs. Thus, we regard our estimate of haulage costs as reasonable.

There is a large amount of heterogeneity across sellers particularly in the per-transaction fixed cost. The fixed cost of seller 2 is significantly higher than other sellers, which is plausible given the low market share of seller 2. One interesting result is that the fixed cost of seller 4 is higher than those of sellers 1 and 3, although seller 4 has the highest market share. Thus, there must be other advantages to rationalize the highest market share of seller 4. One advantage is that seller 4 has the largest number of plants, which gives seller 4 a high market power based on location advantages. Also, the plant-level fixed effects explain the high market share of seller 4: The average size of estimated fixed effects of plants owned by sellers 1, 2, 3, and 4 are 3.18, 2.30, 3.00, and 3.89, respectively. Thus, products produced by seller 4's plants have high quality on average. Other parameters $(\lambda_1, \lambda_2, \sigma_v, \sigma_\epsilon)$ are all estimated precisely.

The estimates of the bargaining power parameters are reported in Figure **??**. We sort buyers in an increasing order of the estimated bargaining parameter (left axis). The lower this parameter is, the stronger the buyer's bargaining position is. For five buyers, the bargaining parameter is not statistically different from one, implying that sellers post prices for those buyers. However, for the remaining 15 buyers, the bargaining parameter is significantly lower than one, ranging from 0.72 to 0.93. We perform a formal test

and reject the joint null hypothesis that all the bargaining parameters are one. The likelihood ratio (LR) test suggests that the Bertrand-Nash model is rejected against our bargaining model at the 1% significance level. That is, buyers have some power over the transaction prices.

There is some relationship between the estimated bargaining power and the buyer's characteristics. Figure **??** also plots the size of buyers (measured in the share of total volumes) on the right vertical axis. First, all buyers with the bargaining parameter of one are a very small buyer. It is natural that these buyers have little power against large sellers. Second, the largest buyer has the highest bargaining power against sellers (the bargaining parameter of 0.72). This relationship between the buyer size and bargaining power is consistent with our finding in Table 5 that the buyer size has a negative coefficient in the price regression.

To evaluate goodness-of-fit of the estimated model, for each transaction we randomly draw ϵ_i and ν_{ij} from the estimated distributions and solve the model. We repeat this process 200 times and average over 200 simulations. Then, we compute mean price, the standard deviation of the price, and the choice probability for the four sellers. Table 9 summarizes the results. The model slightly understates the average price. The standard deviation is predicted to be higher than the empirical counterpart. This may suggest that a more flexible distribution for the prediction error should be used. The model does a good job fitting the predicted choice probabilities.

6.2 Counterfactual Analysis

6.2.1 Source of Market Power

We first evaluate sources of market power. To do so, we solve the model with the following changes to the environment: (i) eliminating joint ownership of plants and (ii) setting the transport cost parameter to zero, thereby eliminating geographic differentiation. For each scenario, we simulate the model

200 times and average over those simulations. Table 10 presents results of these counterfactual analyses. The first and second columns give the average observed price and the average predicted price under the baseline market conditions, respectively.

The third and fourth columns show the effect of eliminating joint ownership. Recall that seller 4 is the largest of the four firms (as can be seen from the choice probabilities in Table 9). The counterfactual analysis suggests that it enjoys significant market power, as its prices fall by 10.7% when there is a de-merger of its plants. This is approximately 20% of the average markup.

The fifth and sixth columns show the effect of setting transport costs to zero. The idea that transport costs confer market power seems to imply that the price effect is negative. However, for all sellers the price increases when transport costs go to zero. One possible reason for this is that firms that are located further from consumers—who previously had to cut their prices to be competitive—are able to raise their prices. Another possibility is that surplus is now higher for all products, so buyers are willing to pay higher prices.

6.2.2 Effect of Uniform Price Restriction

In the terminology of Thisse and Vives (1988), there is price discrimination when sellers do not set uniform FOB prices in a geographical context. In a product differentiation context, on the other hand, there is price discrimination if two varieties are sold at different base prices. In our application, products are differentiated both in terms of geography and product characteristics. As our estimates indicate, price discrimination exists in both dimensions in the UK brick industry.

To investigate the effect of such price discrimination on welfare, we consider a counterfactual scenario where price discrimination is banned in the same spirit as Grennan (2013) and Miller and Osborne (2014) to analyze

how sellers' optimal prices will change. Under this restriction, there should be a single price for any given product (a product is a combination of product characteristics and plant). Sellers cannot price-discriminate based on the location of construction site, the size of the transaction, and the identity of buyers. To compute optimal prices, we proceed as follows. We solve the one-shot Bertrand-Nash pricing game. In this setting, when sellers set prices for their products, they do not observe the random shock ϵ. Once sellers choose their prices, they cannot change their prices during the course of the sample period.

Thus, the expected profit of seller g when the price vector is given by $\mathbf{p}$ is

$$\Pi_g(\mathbf{p}) = \sum_{i=1}^{N} \sum_{j \in \mathcal{J}_{ig}} \mathcal{P}_{ij}(\mathbf{p}) \left[p_j q_i - c_{a(j)} q_i - F_{a(j)} \right]$$

where $\mathcal{P}_{ij}(\mathbf{p})$ is the probability that buyer i chooses product j when the price vector is $\mathbf{p}$.

The first-order conditions are

$$\sum_{i=1}^{N} \left[\sum_{j' \in \mathcal{J}_g} \frac{\partial \mathcal{P}_{ij'}(\mathbf{p})}{\partial p_j} \left[p_{j'} q_i - c_{a(j')} q_i - F_{a(j')} \right] + \mathcal{P}_{ij}(\mathbf{p}) q_i \right] = 0 \text{ for } j = 1, ..., J \tag{11}$$

where

$$\mathcal{P}_{ij}(\mathbf{p}) = \frac{\exp(\beta_{a(j)} + \beta x_j - \tau d_{ia(j)} + \lambda_{ig(j)} / q_i - p_j)}{1 + \sum_{j' \in \mathcal{J}_i} \exp(\beta_{a(j')} + \beta x_{j'} - \tau d_{ia(j')} + \lambda_{ig(j')} / q_i - p_{j'})} \tag{12}$$

and

$$\frac{\partial \mathcal{P}_{ij'}(\mathbf{p})}{\partial p_j} = \begin{cases} -\mathcal{P}_{ij}(\mathbf{p})(1 - \mathcal{P}_{ij}(\mathbf{p})) & \text{if } j = j' \\ \mathcal{P}_{ij}(\mathbf{p}) \mathcal{P}_{ij'}(\mathbf{p}) & \text{if } j \neq j' \end{cases}.$$

We numerically find a J-dimensional vector $\mathbf{p}$ that satisfies J first-order conditions in (11).

Table 11 summarizes the results. Under the uniform price restriction, the

seller-level average price increases by 11.8% to 15.6%. The average markup is around £117 per 1,000 bricks, which is substantially higher than the factual markup (see Table 7 and Figure 5). On the other hand, all sellers lose substantial market share. The share of the outside option increases by 27.7%. This is because sellers cannot freely adjust prices for each transaction under the uniform price restriction, so buyers often choose the outside option even if the surplus for some inside product is higher than that of the outside option. The non-cooperative nature of the pricing game leads to higher average prices, which results in the loss of inside product shares even further. Since the effect of lower market shares of inside products dominates the effect of higher average prices, all sellers have lower profits under the uniform price restriction.

One interesting finding is that the change in profit differs widely across sellers. The profit earned by seller 1 and seller 2 decreases by around 15%, while the decrease in profit earned by seller 4 is more than 30%. This difference is mostly due to the change in market shares. Under the counterfactual scenario, seller 4 loses 19.8% of its market share, while sellers 1 and 2 both experience only a slight decrease in their market shares. Remember that seller 4 has the largest number of transactions. Therefore, we can argue that in the factual scenario, seller 4 can exploit its ability to price discriminate the most. Once price discrimination is banned, however, seller 4 loses competition against its competitors for a fraction of transactions and significantly decreases its market share.

Higher equilibrium prices imply that buyers' utility decreases. Since profits decrease for all sellers, the total surplus in the market unambiguously decreases. Table 11 shows that the total surplus in the market decreases by 23.6%. To better understand this figure, we decompose this into two factors. The total surplus can decrease due to two different reasons under the uniform price restriction. First, some buyers switch to the outside option due to higher prices of inside products. Second, there may be misallocation among

inside products. In the baseline model, the product with the highest surplus is chosen. However, under the uniform price restriction, buyers' choices depend on their utility, which may deviate from the efficient outcome. To isolate the second effect from the first one, we simulate the model holding s_0 constant. That is, for every transaction where the outside option is chosen under the baseline case, we let the outside option being chosen in the counterfactual scenario too. In the same way, for transactions where some inside product is chosen in the baseline case, we exclude the outside option from their choice set. The last column of Table 11 summarizes the result. Although there is heterogeneity across sellers, overall, the misallocation effect among inside products is very small (a 0.1% decrease in surplus). That is, most of the decrease in surplus comes from the loss of market shares of inside products and misallocation plays a minor role.

7 Conclusions

Intermediate product markets are distinct from final product markets because of the greater sophistication of the buyers, the large size of the transactions, and the prevalence of price discrimination using buyer-specific prices. We develop a model of intermediate product market with price setting in which the prices are negotiated between the buyer and seller. We estimate the model using a rich dataset of transactions from the UK brick industry to analyze the effect of competition, location, and transaction size on the prices that are negotiated.

The estimation results show that our bargaining model does a better job explaining the data than the Bertrand-Nash model. Using the estimated parameters, we perform counterfactual analyses, solving the model with the following changes to the environment: (i) eliminating joint ownership of plants, (ii) setting the transport cost parameter to zero, thereby eliminating geographic differentiation, and (iii) banning price discrimination based on the

volume of transaction and location of the buyer. In the first scenario, we find that the markup that the largest seller charges falls approximately by 20%. In the second scenario, for all four sellers the price increases when transport costs go to zero. One possible reason for this price increase is that sellers that are located further from buyers are able to raise their prices when transport costs are eliminated. In the third scenario, we find that the uniform price restriction leads to higher equilibrium prices, while it significantly reduces sellers' profits and total surplus. Overall, the total surplus is reduced by 23.6% due to the uniform price restriction. Most of this decrease in the total surplus comes from the loss of market shares of inside products.

8 Appendix

8.1 Prediction Errors

This appendix characterizes three components of approximation that are implicit in the error specification (10). For the sake of argument, we work with a simple bargaining model in which buyer's outside option is zero. The argument can be easily extended to our full bargaining model.

The predicted unit price paid by buyer i to winning supplier j^* resulting from Nash bargaining is

$$\begin{aligned} p_{ij^*} &= c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \mathbb{E}[\tilde{S}_{ij^*}] \\ &= c_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \left[\delta_{ij^*} + \sigma_\epsilon \left(\gamma - \ln \mathcal{P}_{ij^*}\right)\right], \end{aligned}$$

where γ is Euler's constant, $\mathbb{E}[\tilde{S}_{ij^*}] = \mathbb{E}[\max\{\tilde{S}_{ij}, j \in \mathcal{J}_i\}]$, and $\tilde{S}_{ij} = \delta_{ij} + \epsilon_{ij}$, $j \in \mathcal{J}_i$, for $\epsilon_{ij} \overset{iid}{\sim} EV(1)$. This implies that the observed price $p^{\text{obs}}_{ij^*}$ is equal to the predicted price p_{ij^*}, plus a prediction error, denoted by u_{ij^*}; i.e.

$$p^{\text{obs}}_{ij^*} = p_{ij^*} + u_{ij^*}.$$

The distribution of the prediction error results from the distribution of $\epsilon_{ij^\star}$. It is the same as the distribution of the maximum order statistics of the $EV(1)$, expect for the shift by the mean $[\delta_{ij^\star} + \sigma_\epsilon (\gamma - \ln \mathcal{P}_{ij^\star})]$ and scaling by θ_i.

Consider the marginal cost term $c_{a(j^*)}$ in the above expression. If costs are measured with error $\xi_{a(j^*)}$, then

$$c_{a(j^*)} = \bar{c}_{a(j^*)} + \xi_{a(j^*)},$$

where $\bar{c}_{a(j^*)}$ is predicted cost. Therefore,

$$p_{ij^*}^{\text{obs}} = \bar{c}_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \left[\delta_{ij^*} + \sigma_\epsilon \left(\gamma - \ln \mathcal{P}_{ij^*}\right)\right] + \xi_{a(j^*)} + u_{ij^*},$$

where the combined residual term $\nu_{ij^*} = \xi_{a(j^*)} + u_{ij^*}$ has a distribution that is a mixture of the (shifted) distribution of the maximum order statistics of the $EV(1)$ and the distribution of $\xi_{a(j^*)}$. With a convenient choice for the distribution of $\xi_{a(j^*)}$, this distribution, $F_{\nu_{ij^*}}$, is analytically tractable, but computationally cumbersome. Formally, the contribution to the likelihood function of this transaction is then given by the density of ν_{ij^*}, evaluated at the difference between observed and predicted prices, multiplied by P_{ij^*}:

$$f_{\nu_{ij^*}} \left(p_{ij^*}^{\text{obs}} - \bar{c}_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \left[\delta_{ij^*} + \sigma_\epsilon \left(\gamma - \ln \mathcal{P}_{ij^*}\right)\right] \right) \mathcal{P}_{ij^*}.$$

For all practical purposes, it may be reasonable to approximate $f_{\nu_{ij^*}}$ by a practical alternative, e.g. the pdf ϕ of $N(0, \sigma_\nu^2)$. (Approximation 1)

Denote the predictable part of $p_{ij^*}^{\text{obs}}$ by $\bar{p}_{ij^*}$; i.e.

$$\bar{p}_{ij^*} = \bar{c}_{a(j^*)} + \frac{F_{a(j^*)}}{q_i} + \theta_i \left[\delta_{ij^*} + \sigma_\epsilon \left(\gamma - \ln \mathcal{P}_{ij^*}\right)\right].$$

If the skewness of the distribution of prices is better matched by considering

the logarithmic transform of prices, then

$$\begin{aligned}\ln(p_{ij^*}) &= \ln(\bar{p}_{ij^*} + \nu_{ij^*}) \\ &= \ln(\bar{p}_{ij^*}) + \frac{1}{\bar{p}_{ij^*}}\nu_{ij^*} + HOT,\end{aligned}$$

where HOT denotes higher order terms. This suggests two further approximations: (Approximation 2) replaces the scale factor on ν_{ij^*}, $\frac{1}{\bar{p}_{ij^*}}$, by a constant (subsumed in σ_ν); and (Approximation 3) ignore the HOT.

8.2 Outside Option

When a buyer chooses the outside option, the choice is not observed in the dataset. Therefore, we augment the dataset as follows. For each region, we randomly draw a transaction from its empirical distribution and add the transaction to the data as a sample (observation) in which the outside option is chosen. We repeat this until the share of added transactions equals to the outside option share s_0 for the region, which we observe in the data (see Section 2). Consider the following example. Suppose s_0 is 0.2 for a region and the number of transactions in the region was 800 in the data. Then, we randomly draw a transaction with replacement from the empirical distribution in the region. We repeat this 200 times and assume that buyers of these 200 transactions chose the outside option.

References

[1] Brannman, L. and L.M. Froeb (2000): "Mergers, Cartels, Set-Asides, and Bidding Preferences in Asymmetric Oral Auctions," Review of Economics and Statistics, 82(2), 283-290.

[2] Chen, Z. (2003): "Dominant retailers and the countervailing-power hypothesis", RAND Journal of Economics, 34(4), 612-625.

[3] Chicu (2013): "Dynamic Investment and Deterrence in the U.S. Cement Industry", Working Paper.

[4] Chipty, T. and C.M. Snyder (1999): "The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry", The Review of Economics and Statistics, 81(2), 326–340.

[5] Competition Commission (2007): *Wienerberger Finance Service BV / Baggeridge Brick plc.*

[6] Crawford, G. S. and A. Yurukoglu (2012): "The Welfare Effects of Bundling in Multichannel Television Markets", The American Economic Review, 102(2), 643-85.

[7] de Fontenay C and J Gans (2013): "Bilateral Bargaining with Externalities" Working Paper.

[8] Dobson, P.W. and M. Waterson (1997): "Countervailing Power and Consumer Prices", The Economic Journal, 107(441), 418-430.

[9] Ellison, S. F. and C. M. Snyder (2010): "Countervailing Power in Wholesale Pharmaceuticals", The Journal of Industrial Economics, 58(1), 32–53.

[10] Gowrisankaran, G. Nevo, A. and R.J. Town (2013): "Mergers When Prices are Negotiated: Evidence from the Hospital Industry", NBER Working Paper No. 18875.

[11] Grennan, M. (2013): "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices", The American Economic Review, 103(1), 145-177.

[12] Horn, H. and A. Wolinsky (1988): "Bilateral monopolies and incentives for merger", The RAND Journal of Economics, 19(3), 408-419.

[13] Houde, J-F, (2012): "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." American Economic Review, 102(5), 2147-2182.

[14] Inderst, R. and T. Valetti (2009): "Price discrimination in input markets", The RAND Journal of Economics, 40(1), 1-19.

[15] Inderst, R. and C. Wey (2007): "Buyer Power and Supplier Incentives", European Economic Review , 51(3), 647-667.

[16] Katz, M.L. (1987): "The Welfare Effects of Third-Degree Price Discrimination in Intermediate Good Markets", The American Economic Review, 77(1), 154-167.

[17] Katz, M.L. (1989): "Vertical Contractual Relations", in: Schmalensee, R. and R. Willig (eds.), Handbook of Industrial Organization, Vol. I, Amsterdam: North Holland, 665-721.

[18] Miller, N.H. and M. Osborne (2014): "Spatial differentiation and price discrimination in the cement industry: evidence from a structural model", The RAND Journal of Economics 45(2), 221-247.

[19] Rey, P. and J. Tirole (1986): "The Logic of Vertical Restraints", American Economic Review 76(5), 921-939.

[20] Rey, P. and J. Tirole (2007): "A Primer on Foreclosure", in: Armstrong, M. and R. Porter (eds.), Handbook of Industrial Organization, Vol. III, Amsterdam: North Holland, 2145-2220.

[21] Smith, H. and J. Thanassoulis (2012): "Upstream uncertainty and countervailing power", International Journal of Industrial Organization, 30(6), 483-495.

[22] Sorensen, A. T. (2003): "Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut", The Journal of Industrial Economics, 51(4), 471–492.

[23] Thisse, J.F. and X. Vives (1988): "On The Strategic Choice of Spatial Price Policy", The American Economic Review, 78(1), 122-137.

	Share of	Average # of	
	Volume	Sellers	Transactions
Top 5 buyers	61.2%	3.8	533
Top 10 buyers	77.6%	2.7	337
Top 15 buyers	85.8%	2.6	227
Top 20 buyers	91.0%	1.8	154
Top 100 buyers	99.9%	1.4	90

Table 1: Size Distribution of Top Builders and Developers

UK Nation	Region	Market Size	Brick Deliveries	s_0
England	North East England	554	435	0.216
	North West England	1,490	1,040	0.303
	Yorkshire & Humber	1,190	728	0.388
	East Midlands	1,350	1,250	0.077
	West Midlands	1,150	1,020	0.114
	East Anglia	1,540	1,080	0.297
	London	840	659	0.215
	South East	1,850	1,420	0.230
	South West	1,320	602	0.545
Wales		709	366	0.484
Scotland		1,920	573	0.702

Table 2: Outside Option Market Share by Region
Note: The unit for market size and brick deliveries is millions of bricks.

	Mean	St Dev	Min	Max
L (£/1000)	490.66	259.41	102.86	1715.05
Q (Million)	3.63	2.49	0.56	24.21
Gas Price (G)	1.03	0.40	0.65	2.23
Regional Earnings (W)	8.77	0.72	7.45	10.48
C/Q (Ave. Cost Per 1,000 bricks)	157.73	61.16	42.74	537.72
#obs = 1,063, #years = 6, #plants = 48				

Table 3: Data for Cost Regression

Variable	Mean	St Dev	Min	Max
Price (1,000 bricks) (£100)	1.91	0.96	0.12	11.33
Dist diff (100km)	0.30	0.22	0.00	3.08
buyersize	0.12	0.07	0.00	0.22
Volume (1,000,000 bricks)	0.07	0.08	0.01	1.18
Water Absorption level 1	0.37	0.48	0	1
Water Absorption level 2	0.33	0.47	0	1
Strength Category 1	0.62	0.48	0	1
Strength Category 2	0.14	0.35	0	1
Hand Made Brick	0.26	0.43	0	1
Color Buff	0.30	0.45	0	1
Competition 200km	1.79	1.23	0	4
Marginal Cost (1,000 bricks) (£100)	0.84	0.24	0.43	2.49
#obs = 26,631				

Table 4: Transactions Data

	Est	Std Err	Est	Std Err	Est	Std Err	Est	Std Err
constant	2.17	1.03	1.76	0.24	1.79	0.24	1.74	0.24
Vol (1,000 bricks)	-1.54	0.07	-1.03	0.07	-1.04	0.07	-1.04	0.07
Buyer size	-1.23	0.03	-1.06	0.07	-1.05	0.07	-1.06	0.07
Water Abs 1			0.01	0.02	0.01	0.02	0.01	0.02
Water Abs 2			-0.04	0.03	-0.04	0.03	-0.04	0.030
Hand Made			0.40	0.06	0.39	0.06	0.40	0.06
Color Buff			0.02	0.01	0.01	0.01	0.02	0.01
Strength 1			0.49	0.03	0.49	0.03	0.49	0.03
Strength 2			1.35	0.03	1.35	0.03	1.35	0.03
Marginal Cost			0.46	0.04	0.46	0.04	0.47	0.04
Competition 200km					-0.15	0.04		
Dist diff (km)							0.08	0.02
plnt fixd effct	No		Yes		Yes		Yes	
Adjusted R^2	0.026		0.185		0.185		0.186	
#obs = 26,631								

Table 5: Price Regression
Note: The marginal cost is the cost to produce 1,000 additional bricks.

Parameter	Pooled OLS Estimate	Pooled OLS Std. Err.	Pooled IV Estimate	Pooled IV Std. Err.	IV with Fixed Effects Estimate	IV with Fixed Effects Std. Err.
lnQ	0.66	0.01	0.59	0.06	0.65	0.21
lnG	0.12	0.03	0.13	0.03	0.15	0.05
lnW	1.30	0.11	1.21	0.14	1.21	0.18
Quarter 2	-0.07	0.02	-0.06	0.02	-0.06	0.02
Quarter 3	-0.08	0.02	-0.07	0.02	-0.07	0.02
Quarter 4	-0.03	0.02	-0.03	0.02	-0.03	0.01
Constant	2.59	0.24	2.86	0.34	2.80	0.29
R-squared		0.77		0.76		0.76
#obs		1,063		1,063		1,063

Table 6: Estimates of Cost Parameters

Note: Instruments for ln Q in the IV: regional housing starts, regional housing completions .

#obs = 26,631	Mean	St Dev	Min	Max
MC	84.34	24.39	45.94	233.40
Price	191.69	96.61	12.50	1133.06
Markup	107.36	94.27	-66.15	1051.55

Table 7: Price and Estimated Markup

Parameter	Bargaining Model Estimate	Bargaining Model Std. Err.
constant	-5.672	0.179
τ	11.037	0.166
F_1	0.593	0.033
F_2	1.145	0.096
F_3	0.473	0.054
F_4	0.726	0.027
λ_1	2.390	0.040
λ_2	-19.079	0.308
σ_ϵ	0.925	0.011
σ_v	0.226	0.001
Log-likelihood	-110,039.8	

Table 8: Estimates of Utility and Cost Parameters

	Price		Choice Probability				
	Mean	Std.	Outside	Seller 1	Seller 2	Seller 3	Seller 4
Observed	191.7	96.6	0.342	0.060	0.032	0.176	0.389
Bargaining	189.9	107.5	0.352	0.063	0.036	0.192	0.358

Table 9: Model Fit

	Observed	Model Prediction	Eliminate Ownership		Setting $\tau = 0$	
			Average Price	% Change in Price	Average Price	% Change in Price
Seller 1	179.0	160.7	157.4	-2.1	161.8	0.7
Seller 2	174.2	157.2	154.8	-1.5	157.3	0.1
Seller 3	179.4	176.6	163.5	-7.4	184.4	4.4
Seller 4	200.7	205.3	183.3	-10.7	224.5	9.4

Table 10: Source of Market Power

	% Change in Average Price	Change in Market Share (%)	Change in Profit (mil. £)	% Change in Profit	Change in Surplus (mil. £)	% Change in Surplus	% Change in Surplus with s_0 fixed
Seller 1	15.5	-1.8	-30.1	-15.4	-70.1	-17.9	4.0
Seller 2	15.6	-0.8	-15.7	-13.1	-41.6	-17.0	4.9
Seller 3	13.3	-6.2	-126.7	-23.5	-229.2	-23.1	0.3
Seller 4	11.8	-19.8	-292.5	-31.3	-410.6	-26.5	-2.2
Total			-465.0	-26.0	-751.5	-23.6	-0.1

Table 11: Effect of Uniform Price Restriction

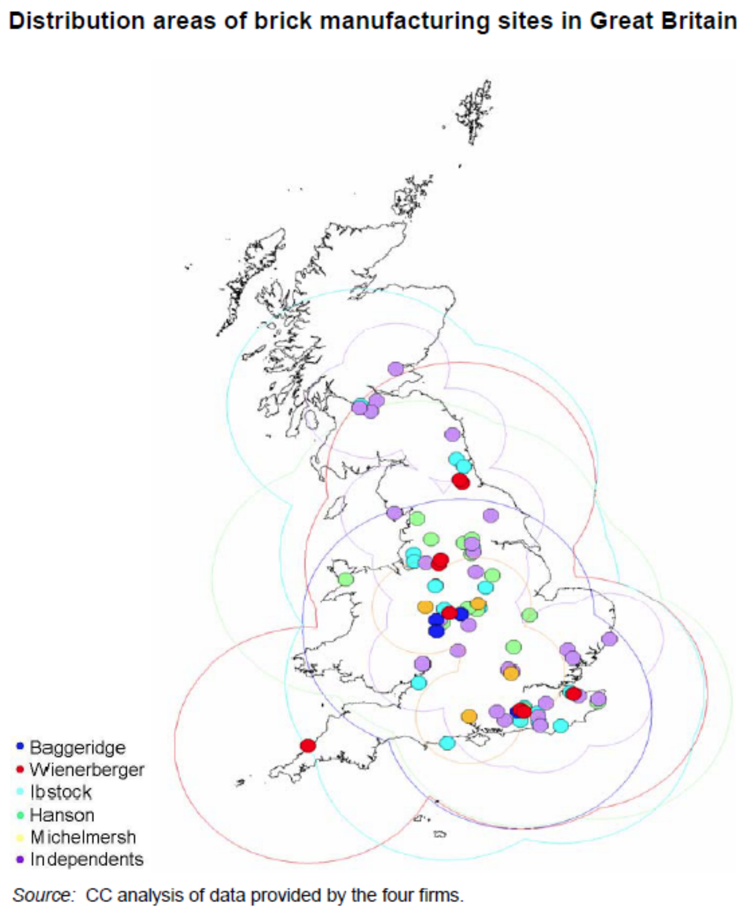

Figure 1: Location of Brick Manufacturing Plants

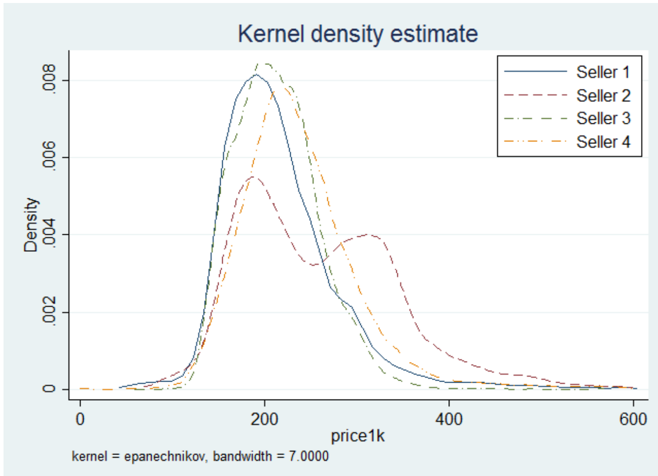

Figure 2: Price Distribution by Sellers

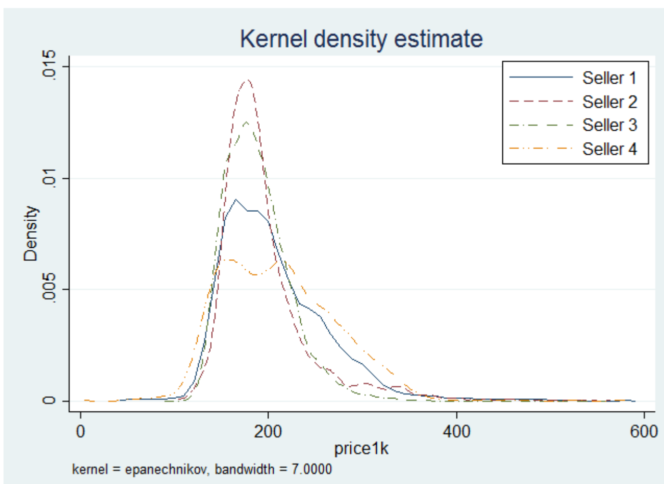

Figure 3: Price Distribution of Selected Product Types

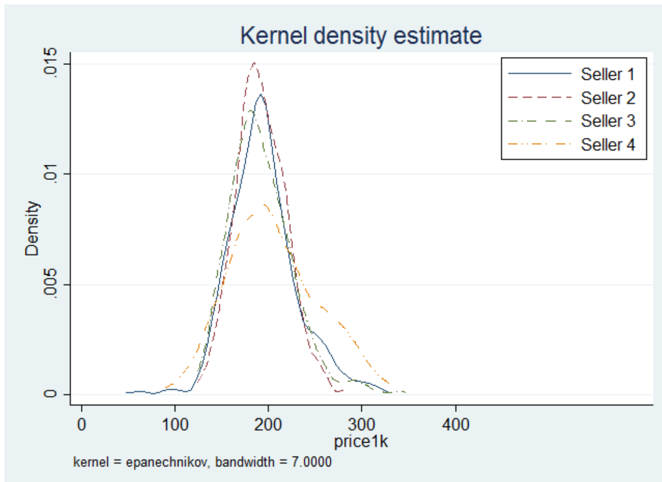

Figure 4: Price Distribution Conditional on Volume

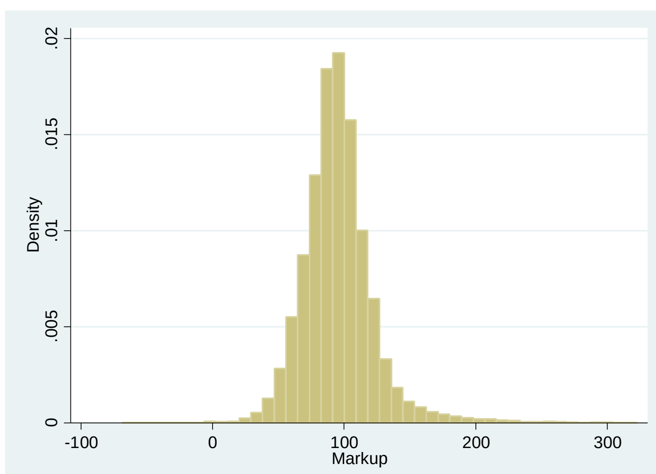

Figure 5: Frequency of Markups between Price and Estimated Marginal Costs