

# Geographic Market Segmentation and Productivity Heterogeneity: A Concrete Example

Chad Syverson  
University of Maryland

June 2000

## ***Abstract***

I contend that imperfect output substitutability explains part of the observed persistent plant-level productivity dispersion. In a case study of the ready-mix concrete industry, I examine the impact of one manifestation of this effect: geographic market segmentation caused by transport costs. A theoretical foundation is presented that characterizes how differences in local demand conditions impact equilibrium productivity levels across regions. I also introduce a new method of obtaining plant-level productivity estimates that is especially well suited to this application and avoids the potential shortfalls of commonly used procedures. I use these estimates to empirically test the presented theory, and the results support my contention. Local demand density (sales per square mile) has a significant influence on the shape of plant-level productivity distributions, and can account for part of the observed within-industry variation in productivity levels and in the dispersion of those levels within a given market.

This is preliminary work; comments are welcome. Please contact author before citing.

The research in this paper was conducted while the author was a research associate at the Center for Economic Studies, U.S. Bureau of the Census. Research results and conclusions expressed are those of the author and do not necessarily indicate concurrence by the Bureau of the Census or the Center for Economic Studies.

## Geographic Market Segmentation and Productivity Heterogeneity: A Concrete Example

Recent empirical explorations have left little doubt about the magnitude of plant-level productivity variation: it is enormous. This heterogeneity is also persistent. Perhaps surprisingly, most of the variation cannot be explained by differences between (even narrowly defined) industries. Haltiwanger (1997), for example, finds that 91.5% of plant-level productivity growth variation exists within four-digit SIC industries. An assortment of theoretical work has arisen attempting to explain the sources of such diversity. The great majority of this research focuses on supply-side (production) causes, such as idiosyncratic technology shocks, management influences, R & D efforts, or investment patterns.<sup>1</sup>

In this paper I turn my attention to the demand (i.e., output market) side, and look at how demand forces can cause such within-industry heterogeneity to persist.<sup>2</sup> I argue that across-plant differences in output market conditions are partially responsible for observed persistent productivity dispersion—and in fact, the dispersion of that dispersion. The specific channel through which this posited influence flows is variable output substitutability in a world of product differentiation. The more difficult it is for consumers to switch between various suppliers, the greater the amount of dispersion that can be sustained. I will focus here on a particular component of substitutability, geographic market segmentation created by transport costs, and examine its impact on productivity dispersion within a single industry. The purpose of this paper, however, is not to give the final word on transport costs and productivity in a particular industry. Instead, I hope to show through a detailed case study—where many potentially confounding factors are held constant—how transport costs as well as other substitutability factors might impact productivity levels, dispersion, and dynamics throughout the economy.

The rationale for using output market effects, and substitutability specifically, to explain the degree of within-industry productivity heterogeneity is more readily apparent when we

---

<sup>1</sup> Just a sampling includes Jovanovic (1982), Ericson and Pakes (1995), and Caballero and Hammour (1998). See Bartelsman and Doms (2000) for a review of this literature.

<sup>2</sup> There are supply-side stories that can explain persistent dispersion as well. I simply want to highlight another piece of the puzzle, to my knowledge not previously formalized, from the demand side.

consider how such wide efficiency variation can exist in equilibrium. After all, output should tend to be reallocated to more productive plants over time. High-productivity plants are able to produce output at lower cost than industry rivals, allowing them to grab additional market share by undercutting their opponents' prices without sacrificing profit rates. We might expect this shakeout process to redistribute most or all of an industry's production to a few high productivity plants. Output and productivity patterns like this are not usually observed in the data, however; the overwhelming weight of empirical evidence indicates widely varying producer productivity levels within nearly every industry.

What prevents this output reallocation process from occurring? Many possible explanations, such as demand booms (when nearly anybody can operate profitably), are short-run stories. They cannot explain why we see large within-industry productivity dispersion throughout the business cycle. Imperfect output substitutability, however, is a long-run explanation. For example, microbrewers may not produce their output at nearly as low a unit cost as Miller or Anheuser-Busch, but they can survive (and even thrive) in the long-run marketplace because segments of the population prefer microbrews to mass-produced beer and are willing to pay the higher unit prices necessary to support the microbrewers.

In Syverson (2000) I examine the role across-industry differences in measurable output substitutability factors play in determining the equilibrium distribution of plant-level productivity within an industry. The intuitive and testable premise of that work is that industries with less output market segmentation (i.e., greater substitutability) should have plant-level productivity distributions that have less dispersion and higher central tendency than distributions in industries with more segmented markets. The intuition behind this notion is simple. Greater substitutability makes it easier for customers to shift purchases to more efficient producers, driving plants at the bottom end of the productivity distribution out of business and raising the bar for successful entry. This truncation of the distribution lowers productivity dispersion and increases the average efficiency level.

This paper builds upon that theme, but takes a much more directed approach by looking at the influence of a single source of market segmentation within one four-digit SIC industry. Instead of relying on output substitutability variation across industries to explain differences in industry productivity distributions, this study investigates how within-industry market

segmentation creates productivity dispersion within and between these market segments. A single-industry case study is a useful complement to my interindustry work. Its narrow focus implies that any results found here do not have the comprehensiveness of the across-industry study, but it benefits from the fact that the influence of technological differences on productivity heterogeneity has largely been removed. The industry choice for the case study also isolates the effect of a single source of market segmentation that has inherent interest to a significant body of economic research: transport costs.

The industry I focus on is ready-mix concrete (SIC 3273). It has a number of characteristics that make it very favorable for this case study. First, of course, industry production is subject to substantial transport costs. This creates a series of quasi-independent geographically segmented markets—all potentially subject to idiosyncratic demand movements. I look at how output substitutability differences across these local markets affect the plant-level productivity distributions within them. The high transport costs also result in an industry characterized by a large number of geographically dispersed establishments, which will be useful in the empirical portion of the paper. Finally, industry output is more or less physically homogeneous. This largely eliminates the influence of physical product differentiation on the plant-level productivity distribution (which I have found to be a very substantial effect across industries). This will help sharpen the focus on the influence of geographic market segmentation rather than sources of aspatial market heterogeneity.

I model and test a competitive structure that results in local concrete markets with high demand density (measured as ready-mix sales per square mile) having less-disperse, higher-average productivity distributions than those in low density markets. The mechanism through which this happens will be explained in detail below, but can be summarized as follows. A larger market (measured by total concrete sales) requires more producers to serve it. The larger number of concrete establishments in a fixed market area leads to greater output substitutability for concrete buyers. High substitutability in turn implies that low-efficiency plants cannot operate profitably given the ability of customers to switch suppliers, forcing them out of business and truncating the low end of the plant-level productivity distribution. The resulting long-run equilibrium yields productivity distributions in larger markets that have less productivity dispersion, higher average productivity, and a greater share of output produced by high-efficiency

plants. This also causes a curious between-plant form of returns to scale: producers in larger markets are more efficient on average, but not because plants become more productive as they themselves become larger. Instead, the observed scale effect is the product of selective survivorship: less productive establishments are eliminated as markets grow.

It is easy to imagine how geographic market segmentation consequences can extend beyond the ready-mix industry, especially into other manufacturing industries with low value-to-weight output and the retail sector, but also into other industries to a lesser degree. Imperfect substitutability created by transport costs can thus explain a portion of the persistent productivity heterogeneity throughout the economy.

The paper is organized as follows. I begin by constructing a theoretical framework that formalizes my intuitive premise. A discussion of plant-level productivity estimation methodology and a review of the data follow this. I next present the empirical results and test them for robustness to several identification assumptions. A conclusion follows.

## I. Theory

To formalize the story linking demand density, the number of producers, output substitutability, and the local productivity distribution, I require a theoretical framework that incorporates heterogeneous producers and contains some notion of substitutability. Further, it should allow the endogenous determination of the equilibrium plant productivity distribution, and offer testable implications as to the nature of the distribution as exogenous factors vary. The primary exogenous variable that I am interested in is demand density, of course, so the model should incorporate both market sales and area into equilibrium determination. I can meet these requirements by modifying a model employed by Melitz (1999) in another context. This model, which I describe below, nicely incorporates these items and will serve as a theoretical foundation for my empirical work. Because a great deal of the model has been thoroughly discussed in other work, I forgo much of the formal analysis here and focus on intuitive discussion.

### *Model*

The model frames the allocation of market demand across producers in the familiar Dixit-Stiglitz (1977) structure. This setup allows interplant output substitutability to enter explicitly

into the model, and offers an analytically tractable way to incorporate heterogeneous producers into an equilibrium market structure. The Dixit-Stiglitz model does not have producers competing in a strategic way, which prevents the model from becoming unsolvable as greater numbers of heterogeneous producers exist in market equilibrium. This is a convenience mathematically but an analytical disadvantage. However, as will be seen, the framework does allow a non-strategic source of competitive pressure: larger numbers of plants in a market will increase the ease with which consumers can switch between suppliers, lowering margins and driving out less efficient producers. A reasonable interpretation is that the competitive framework here is a “black-box” model. That is, it simply relates key inputs (market size, area, and the number of producers) to key outputs (the equilibrium elasticity of substitutability), while glossing over more complex fundamental interactions driving the process that may be very interesting, but need not be fully characterized to serve my present purposes.

The size of the modeled market is an important consideration. It must be chosen with knowledge of what is theoretically sensible and what the data holds, and inevitably involves tradeoffs between conflicting needs. The best quality fit to the intuitive story outlined above implies that the modeled market unit should be an autonomous local geographic region. Local producers satisfy the entire demand in their own market; there is no trade between distinct geographical areas. At the same time, all local production units compete with each other; consumers in effect may purchase from any plant within their own locality. The assumption of no intermarket trade may at first perhaps seem unreasonable, but in the case of the ready-mix concrete industry is not necessarily far from the truth. As I will discuss later, careful definition of empirical market areas yields market groups that can be (in the high transport cost ready-mix industry) reasonably assumed to exhibit these characteristics.

I must incorporate a notion of within-market transport costs into the model. I do so here in a non-location-specific manner, so that an increase in the number of local producers reduces transport barriers for all customers regardless of their location within that market. I do not explicitly take into account specific placements of producers or consumers. Analyzing detailed geographic competition would doubtlessly be informative, but is more than needed to demonstrate the key processes I wish to describe. This may best be left to future work.

A representative concrete consumer in the modeled local market has C.E.S. preferences

over a continuum of goods (producers) indexed by  $i$ :

$$U = \left[ \int_{i \in I} q(i)^{\frac{\rho}{1+\tau}} di \right]^{\frac{1+\tau}{\rho}} \quad (1)$$

This is the familiar Dixit-Stiglitz utility function with one exception: the presence of the variable  $\tau$ . Its value reflects the influence of within-market transport costs on the ability of concrete purchasers to substitute the output of one ready-mix establishment for another. A larger value of  $\tau$  implies higher transport impedance within a locality and lower substitutability. The notion that there would be product differentiation in the concrete market, a product having very little physical product differentiation, may at first glance appear to be a contradiction. However, in industries where transport costs are relevant, products from different locations are often considered different goods even if each plant makes physically identical products. This is the familiar dual interpretation of spatial and product-space differentiation discussed in a number of studies, such as Hotelling (1929), Salop (1979), and Weitzman (1994). This framework is consistent with the intuition—important to the story here as well—that customers can more easily substitute between producers when there are more producers in a given market area. To capture this notion I make  $\tau$  a function of the ratio of producer mass  $M$  (the continuous product space analog to the number of producers) to the total market area  $A$ ; i.e.,  $\tau = \tau(M/A)$ . I will return to the specific form of this function later; the only requirement I impose for now is that it be nonnegative. I also allow plants to differ in product space, not because ready-mix concrete varies physically by any great degree from one producer to another (it does not), but because in reality there are always a host of plant-specific abstract goods, such as customer service and delivery reliability, bundled with the physical product. The imperfect substitutability resulting from this type of product differentiation is captured by the familiar parameter  $\rho$ , assumed to lie on the interval  $(0,1)$ .

As is well known, the utility function above implies an elasticity of substitution between any two goods equal to

$$\sigma = \frac{1}{1 - \frac{\rho}{1+\tau}} = \frac{1+\tau}{1+\tau-\rho} \quad (2)$$

This value is always greater than one because of the assumptions that  $\tau$  is nonnegative and that

different varieties are substitutes ( $0 < \rho < 1$ ). It is easy to show that the quantity ratio between any two varieties is determined completely by their price ratio:

$$\frac{q(i_1)}{q(i_2)} = \left[ \frac{p(i_1)}{p(i_2)} \right]^{-\sigma} \quad (3)$$

Production requires a single input, labor, which is supplied elastically to the industry at a wage  $w$ . The production function is linear in labor and includes an overhead labor fixed cost  $f$ . Plants differ only in their marginal product of labor, which is embodied in the productivity value  $\phi$ :

$$q = \phi(l - f) \quad (4)$$

As Dixit and Stiglitz show, the demand structure outlined above leads to each producer pricing their output at the same markup over marginal cost. Normalizing  $w = 1$ , the optimal price for each producer is

$$p = \mu \cdot MC = \frac{\sigma}{\sigma - 1} \frac{w}{\phi} = \frac{1 + \tau}{\rho \phi} \quad (5)$$

Therefore more efficient producers sell at a lower price.

Given that plant profits are equal to  $\pi = r - l = pq - l$ , substitution of the production function and optimal pricing rule into this expression yields establishment profits of

$$\pi = \frac{r}{\sigma} - f \quad (6)$$

Note that the above production structure will yield an output and revenue dispersion across establishments governed by plants' relative efficiencies:

$$\frac{q(\phi_i)}{q(\phi_j)} = \left( \frac{\phi_i}{\phi_j} \right)^\sigma \quad \text{and} \quad \frac{r(\phi_i)}{r(\phi_j)} = \left( \frac{\phi_i}{\phi_j} \right)^{\sigma-1} \quad (7)$$

so more productive plants produce more and have higher total sales. Notice, too, that as the elasticity of substitution  $\sigma$  increases, the output and revenue distributions skew further toward high-efficiency (low-price) producers.

Equilibrium in the local ready-mix concrete market will be characterized by a mass  $M$  of establishments with a productivity distribution over some portion of  $(0, \infty)$ . Now define a market quantity index  $Q = U$  and a price index  $P$  such that

$$P = \left[ \int_0^\infty p(\phi)^{1-\sigma} M \mu(\phi) d\phi \right]^{\frac{1}{\sigma-1}} \quad (8)$$

where  $\mu(\phi)$  is the probability density function of operating plants' productivity distribution. This can be rewritten by using the optimal pricing rule as

$$P = M^{\frac{1}{1-\sigma}} p(\tilde{\phi}) \quad (9)$$

where

$$\tilde{\phi} = \left[ \int_0^{\infty} \phi^{\sigma-1} \mu(\phi) d\phi \right]^{\frac{1}{\sigma-1}} \quad (10)$$

This moment of the productivity distribution, which I shall refer to as average productivity (it can actually be shown to be a quantity-weighted average of plant productivity levels), has convenient properties for aggregation. As Melitz (1999) shows, total market revenue  $R$  (and profits  $\Pi$ ) can also be expressed as simple relationships between the mass of producers and the revenue (profits) of the establishment with this average productivity level. That is,

$$R = \int_0^{\infty} r(\phi) M \mu(\phi) d\phi = M r(\tilde{\phi}) \quad \text{and} \quad \Pi = \int_0^{\infty} \pi(\phi) M \mu(\phi) d\phi = M \pi(\tilde{\phi}) \quad (11)$$

It is easy to see from these expressions that the plant with average productivity will have revenue and profits equal to the market averages. Also, it is apparent that the equilibrium mass of producers is pinned down by aggregate revenue (i.e., concrete sales in the local market) and the average productivity level. This fact will be important later.

Entry and exit processes take place in a dynamic framework. An infinite number of potential entrants exist, and can enter at any time if they pay a sunk labor cost of entry  $f_e$ . All entrants discover their productivity value  $\phi$ , drawn from a common distribution  $g(\phi)$  over  $(0, \infty)$ , upon paying the entry cost. Plant productivity levels are constant over time. After discovery of their efficiency level, plants can choose to either produce—in which case they are subject to the additional fixed cost of production  $f$ —or exit immediately at no additional cost. A steady-state equilibrium with a constant productivity distribution is maintained by the assumption that any producing plant is subject to an exit-forcing shock with exogenous probability  $\delta$ , which does not vary with the productivity level. In equilibrium, the mass of successful entrants equals the exiting mass.

Both  $g(\phi)$  and  $\delta$  are exogenous, but there is considerable flexibility allowed in their form and magnitude, improving the generality of results. Slightly more restrictive is the assumption that the probability of exit is irrespective of an establishment's productivity level. This notion

seems to be contrary to the plant-level evidence of entry and exit (see Dunne, Roberts, and Samuelson (1989), for instance). However, if one interprets the initial (endogenous) entrant's produce/exit choice as occurring over a short period of production in reality, rather than immediately as in the model, it can explain the tendency seen in empirical work for younger and less efficient plants to be more apt to exit.

While the productivity distribution prior to entry is exogenous, the distribution of producing plants is determined within the model. This is so because, as can be shown, there is always a positive productivity level where plant profits are zero. And because plant productivity levels are constant, no producer less efficient than this cutoff productivity level will produce in equilibrium. Therefore the equilibrium productivity distribution of plants choosing to operate,  $\mu(\phi)$ , will be a truncation of  $g(\phi)$ :

$$\mu(\phi) = \begin{cases} \frac{g(\phi)}{1 - G(\phi^*)} & \text{if } \phi \geq \phi^* \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

where  $\phi^*$  is the productivity level such that  $\pi(\phi^*) = 0$ , and  $G(\phi^*)$  is the cumulative distribution function of  $g(\phi)$  evaluated at the cutoff productivity value. We can therefore rewrite the earlier expression for average productivity as

$$\tilde{\phi} = \left[ \frac{1}{1 - G(\phi^*)} \int_{\phi^*}^{\infty} \phi^{\sigma-1} g(\phi) d\phi \right]^{\frac{1}{\sigma-1}} \quad (13)$$

Notice that given  $g(\phi)$ , the average productivity level depends only upon the cutoff.

The equilibrium value of  $\phi^*$  is determined by two conditions that must hold in the steady state. The first, as alluded to above, is that the plant with cutoff productivity  $\phi^*$  must make zero profits. The second is a free entry condition, requiring that the expected payoff from entry be zero in equilibrium. Both of these conditions can be used to derive expressions of average plant profits as a function of the cutoff level; the intersection of these two functions determines  $\phi^*$  in equilibrium.

The expressions (6) for plant profits, (7) for revenue distribution, and (11) for average revenue and profits can be used to derive an expression relating average profit and cutoff productivity  $\phi^*$ :

$$\bar{\pi} = \frac{\Pi}{M} = \pi(\tilde{\phi}) = \frac{r(\tilde{\phi})}{\sigma} - f = \left[ \frac{\tilde{\phi}}{\phi^*} \right]^{\sigma-1} \frac{r(\phi^*)}{\sigma} - f \quad (14)$$

The equilibrium condition that the marginal plant earn zero profit implies that  $r(\phi^*) = \sigma f$ . Thus we have the equilibrium expression relating average profits and the cutoff productivity level for the first condition:

$$\bar{\pi} = \left[ \left( \frac{\tilde{\phi}(\phi^*)}{\phi^*} \right)^{\sigma-1} - 1 \right] f \quad (15)$$

I explicitly write average productivity as a function of the cutoff level to stress that  $\phi^*$  enters into this function in two places.

The second condition requires that the expected value of entry be equal to zero. Recall that the entering firm, upon learning its productivity level, must decide whether to produce or exit before producing. Because productivity levels do not change over time, if it produces the establishment will make the same profit every period until it receives a killer shock (which has a probability of  $\delta$  each period). Thus after the sunk cost, the value of entry is either zero if  $\phi < \phi^*$ , or (assuming no discounting other than knowledge of the possibility of an exogenous shock):

$$v(\phi) = \sum_{t=0}^{\infty} (1-\delta)^t \pi(\phi) = \frac{\pi(\phi)}{\delta} \quad (16)$$

The expected value of entry is then the product of the probability and average value of successful entry minus the sunk entry cost.

$$v_e = (1 - G(\phi^*)) \frac{\bar{\pi}}{\delta} - f_e \quad (17)$$

Given that this must be zero in equilibrium, the second expression relating average profit and  $\phi^*$  is obtained:

$$\bar{\pi} = \frac{\delta f_e}{1 - G(\phi^*)} \quad (18)$$

This equation shows that the free entry condition requires average profits be increasing in  $\phi^*$ .

This makes sense; as the cutoff productivity rises, the probability of successful entry decreases, requiring that the expected profit from successful entry increases to compensate for this fact.

The impact of changing  $\phi^*$  on the other equilibrium condition (15) is not so immediately apparent. Recall, however, that the first term in the bracketed expression is equal to the revenue ratio between the average productivity plant and the marginal plant. This climbs toward infinity

as the cutoff productivity level goes to zero, because average productivity must always be positive. The behavior of this ratio as the cutoff productivity level increases depends on the properties of  $g(\phi)$ , but for common distribution functions, it will monotonically decrease toward one as  $\phi^*$  rises. Thus the requirement that the marginal plant make zero profits implies a negative relationship between average profits and the cutoff productivity level. This, too, is as expected. Since the marginally efficient plant earns zero profits, lower average profit levels imply that the productivity level required to operate in the black must increase.

One can demonstrate that the properties of these two functions ensure a unique equilibrium cutoff productivity level  $\phi^*$  (see Melitz (1999)). Notice that the elasticity of substitution  $\sigma$  affects this level only through one function: that which is derived from the requirement of zero profits at the cutoff productivity level (15). An increase in  $\sigma$  shifts this function up (average profit increases for all possible cutoff productivity levels). Because the free-entry condition implies average profits are an increasing function of the cutoff productivity, an increase in substitutability raises both the cutoff productivity and the average profit level in equilibrium. This in turn further truncates the ex-ante productivity distribution, thereby decreasing productivity variation across producers and raising their average productivity level. The density of producers within a market influences the market's productivity distribution because  $\tau(M/A)$  determines the elasticity of substitution. This, of course, is exactly the process outlined in the introduction.

The equilibrium mass of producers  $M$  must still be determined. Recall that this is pinned down by the total size of the local market and the average establishment revenue (a function of the average productivity level), the latter being independent of the former:

$$M = \frac{R}{\bar{r}} = \frac{R}{\sigma(M)[\bar{\pi}(\sigma(M)) + f]} \quad (19)$$

Here, all values that functions of  $M$  are noted. The implicit solution for  $M$  from this relation yields the equilibrium mass of producers.

The model is now completely specified. Given assumptions about the form of the ex-ante productivity distribution  $g(\phi)$ , the size of fixed costs, and the function  $\tau(M/A)$ , I can compute equilibrium values of the cutoff productivity level, number of producers, and elasticity of substitution for any given value of demand density  $D = R/A$ . These derived relationships will

yield testable implications regarding the nature of the plant-level productivity distribution within a local market as the density of demand changes.

### *Numerical Comparative Statics*

It is not possible to derive analytical expressions for all of the relevant variables in the model above, but comparative statics exercises relevant to my empirical work can be performed computationally with an assumption about the form of  $g(\phi)$ . I describe this process and present the results below.

The computational algorithm to find equilibrium values in the model begins by solving an implicit function in  $\phi^*$  constructed from equating the zero cutoff profit (15) and free entry (18) conditions:

$$\left[ \left( \frac{\tilde{\phi}}{\phi^*} \right)^{\sigma-1} - 1 \right] (1 - G(\phi^*)) = \frac{\delta f_e}{f} \quad (20)$$

Assumed parameter values are used to compute the right hand side of the equation, and an ex-ante productivity c.d.f.  $G(\phi^*)$  is chosen. The expression is then solved numerically to find the cutoff productivity value corresponding to a posited value for the elasticity of substitution,  $\sigma_o$ . Using this  $\phi^*$  value and its equilibrium average profit value, I next compute average productivity and revenue levels according to (13) and (6). Given the exogenous market size (total revenues)  $R$ , average plant revenue allows computation of the mass of producers  $M$ , as in (11). This value is divided by the market area  $A$  and  $\tau = \tau(M/A)$  is computed. Finally the value  $\tau$  along with an assumed  $\rho$  is put into expression (2) to obtain a value for  $\sigma$ . If this value equals the posited value  $\sigma_o$ , the series of computations characterizes the equilibrium. If the posited and derived substitution elasticities do not match, other posited values are tried until a match is found. This process is repeated for various values of demand density  $R/A$ , allowing me to derive empirically testable implications about the nature of the productivity distribution in equilibrium.

It is easy to demonstrate that the complicated function  $\sigma(\sigma_o)$ , which relates the posited elasticity of substitution  $\sigma_o$  to the derived value  $\sigma$ , has at most one fixed point, and thus any computed equilibrium is unique. I have already shown how an increase in the elasticity of substitution increases the cutoff productivity and average profit levels. As seen in (19), an

increase in  $\sigma$  decreases  $M$  both directly and indirectly through its influence on average profits. As long as  $\tau$  is a weakly decreasing function for all values of producer density  $M/A$  (I shall discuss the case for this below), an increase in  $\sigma_0$  weakly decreases the computed value of  $\sigma$ . The fact that  $\sigma(\sigma_0)$  is everywhere weakly decreasing ensures that any fixed point is unique. A fixed point *exists* as long as the limit of this function is greater than one as  $\sigma_0 \rightarrow 1$  from above (recall that the utility function requires that  $\sigma > 1$ ). This condition depends on values of exogenous parameters and the specific form of  $\tau = \tau(M/A)$ .

For my simulations, I assume that ex-ante productivity distribution is exponential with mean one; i.e.,  $g(\phi) = e^{-\phi}$ . The shape of the distribution seems reasonable and it yields an easily calculated expression for the average productivity level:

$$\tilde{\phi} = e^{\phi^*} \int_{\phi^*}^{\infty} \phi^{\sigma-1} e^{-\phi} d\phi = e^{\phi^*} \Gamma(\sigma) [1 - F(\phi^*)] \quad (21)$$

where  $\Gamma(\sigma)$  is the gamma function evaluated at  $\sigma$ , and  $F(\cdot)$  is the value of the gamma- $(\sigma, 1)$  cumulative distribution function at the cutoff productivity level. I set  $f = 0.04$ ,  $f_c = 1$ , and  $\delta = 0.2$ . These values are arbitrary, of course, but are unlikely to affect the model's qualitative implications. They only change the value of the cutoff productivity level, and have no bearing on the direction of the other equilibrium influences. While the magnitudes of the variables in the model are surely affected by the choice of these parameters, these parameter values do not influence the qualitative behavior of the values of interest. I am not seeking quantitative implications from this highly stylized model.

The choice of the functional form for  $\tau(M/A)$  requires more careful attention. It should be everywhere nonincreasing in its argument to rule out multiple equilibria; that is, an increase in the number of producers in a given area should not decrease the elasticity of substitution between producers. This is a very reasonable proposition. It is hard to imagine how in any market the addition of another producer could make it more difficult (utility-wise) for consumers to switch between sellers. For the equilibrium to depend upon the number of producers,  $\tau(M/A)$  must be decreasing over at least some range. If  $\tau$  is constant over all producer densities, the equilibrium productivity distribution will be unaffected by the local market sales density. When this derivative is negative, however, a larger market (higher  $R$ ) increases the elasticity of substitution (and changes the productivity distribution) because the increased number of producers in a given

market area lowers the value of  $\tau$ . I contend a negative derivative is likely whenever there is some degree of geographic market segmentation within an industry. If markets are at least partially localized, having an additional plant locate within one's own market should make it easier to substitute between producers. Because larger markets require a greater number of producers to satisfy demand, consumers in these markets enjoy greater freedom in choosing suppliers.

For guidance as to the specific form of  $\tau(M/A)$ , I look to other theory on producer density and substitutability. Specifically, I employ the implications of a Salop (1979)-type model where homogenous producers distributed around a circular local market compete in prices. This model does not lend itself directly toward obtaining a relationship between producer density and the elasticity of substitution. However, it does offer an expression linking producers' markup over marginal cost and producer density. Since the markup in the present model is a unique function of the elasticity of substitution, I can back out the implied relationship between plant density and  $\sigma$ .

In a Salop-type competitive structure with  $N$  technologically identical producers evenly spaced around a circular market with circumference  $A$  and unit-length transport costs  $t$ , one can show (assuming every consumer buys a single unit of output) that the implied optimal markup is

$$\mu^* = 1 + \frac{t}{c \cdot \frac{N}{A}} \quad (22)$$

where  $c$  is the marginal cost of production common to all plants. This expression indicates that the markup varies with the reciprocal of the producer density within the market. Taking this as a benchmark for my model implies that we should choose the form of  $\tau(M/A)$  so that our markup behaves similarly. That is,

$$\frac{\partial \mu}{\partial \left(\frac{M}{A}\right)} = \frac{\partial}{\partial \left(\frac{M}{A}\right)} \left[ \frac{1 + \tau\left(\frac{M}{A}\right)}{\rho} \right] \propto \left(\frac{M}{A}\right)^{-2} \quad (23)$$

Clearly, the function  $\tau(M/A) = (M/A)^{-1}$  satisfies this condition. This is the form I use in the equilibrium computations. I perform robustness checks as well, computing equilibrium values for three other plausible other functional forms:  $\tau(M/A) = (M/A)^{-0.5}$ ,  $\tau(M/A) = [\ln(M/A)]^{-1}$  and

$\tau(M/A) = (M/A)^{-2}$ . All of these functions are everywhere positive for nonnegative producer densities and decreasing in their argument. I will compare results from each of these functions to each other to ensure that the crucial qualitative characteristics of the model do not change.

Figure 1 shows how the mass (number) of plants  $M$  and average revenue per plant vary with market size  $R$  in the computed equilibrium. Both are concave functions of market size, but average plant sales has considerably less curvature. Clearly, increases in total demand are taken up on both the intensive and extensive margins, but intensive expansion begins to dominate in larger markets. I will compare these implications of the computed equilibrium to the patterns seen in real-world ready-mix concrete markets below.

Figure 2 shows the primary testable implications of the model. As can be seen, the theory implies that the cutoff and average productivity levels are monotonically increasing in demand density. An increase in the cutoff productivity level implies a decrease in dispersion of the productivity distribution, of course. Thus the model exhibits characteristics of the central intuition of this paper: plant-level productivity distributions in regions with high sales (demand) densities should have less dispersion and higher central tendency than those in smaller markets. The secondary implication elucidated by the figure is the nature of the substitutability increase as demand density grows. Recall that in our framework, a higher  $\sigma$  will skew the output and sales distributions toward more productive producers. Thus we should expect to see in the data that variations in demand density across markets, besides causing cross-sectional variation in moments of local plant-level productivity distributions, will also shift the allocation of output across local plants. High-productivity plants should have larger output shares in high density markets.

I check the results for robustness to the functional form of  $\tau(M/A)$  in Figure 3. For the sake of clarity, I show only the cutoff productivity levels, but as is obvious from equation (13), the average productivity curves are similarly shaped. The figure indicates that while the cutoff productivity functions do differ slightly, the primary result demonstrated above—that the cutoff and average productivity levels monotonically increase in demand density—remains. The testable implications of the model that I am concerned with do not appear to depend on the chosen form of  $\tau(M/A)$ .

## II. Productivity Estimation and the Market Segmentation Method of Instrument Identification

The empirical portion of this paper requires plant-level productivity estimates. Typically, establishment productivity estimates are the residuals of an industry-wide production function estimated using plant data. This procedure implicitly assumes that all plants in the industry operate with the same production technology. It is a common assumption in such studies, and is likely appropriate in the case of ready-mix concrete, which is produced by largely the same process everywhere (U.S. Bureau of Labor Statistics (1979)).<sup>3</sup>

Methods of estimating the industry production function require some attention. A naive procedure would simply regress plant outputs on some functional form of inputs using ordinary least squares methods. However, as Marschak and Andrews (1944) first pointed out, simultaneity of productivity and inputs cause such methods to provide inconsistent estimators of production function parameters (and therefore productivity values as well). Researchers have struggled since then to circumvent the endogenous inputs problem through the use of various econometric techniques, some more successful than others.<sup>4</sup> Despite these efforts, no broadly applicable method could be found that could consistently and adequately address the simultaneity problem.

In an influential paper, Olley and Pakes (1996) proposed a three-step algorithm that has in a very short time since become the standard technique for estimating production functions with plant-level data because of its clever treatment of endogeneity and relative ease of implementation.<sup>5</sup> The thrust of their procedure is inversion of the plant investment function to back out a productivity proxy polynomial that contains only producer observables. They demonstrate this is mathematically consistent if plant investment is a monotonically increasing function of plant productivity, and if productivity is the only unobserved establishment-specific variable in the investment function. These assumptions insure that, given a plant's capital stock, there is a one-to-one mapping between plant productivity and investment, allowing one to control for unobserved plant productivity values with observed investment and capital stocks.

---

<sup>3</sup> I will test my results for robustness to technology differences across markets.

<sup>4</sup> Griliches and Mairesse (1995) survey these methods and their relative benefits and shortcomings.

<sup>5</sup> For examples of its application, see Griliches and Mairesse (1995), Aw, Chen, and Roberts (1997), and Levinsohn and Petrin (1999).

Unfortunately, the ready-mix concrete industry is a particularly poor fit for the Olley-Pakes (O-P) algorithm. As I argue in Syverson (1999), the required assumption that productivity is the only unobserved plant-specific state variable in the investment function is unlikely to hold when output markets are segmented. And because of transport costs, the ready-mix concrete industry is highly segmented geographically; industry establishments sell a majority of their output to buyers in their immediate vicinities. Under such conditions local markets can yield considerable spatial demand variation across producers. Because they operate so narrowly, geographically speaking, ready mix plants are very likely to take their idiosyncratic (region-specific) demand state into account when hiring inputs. Demand (or expected demand) is then an additional plant-specific variable in the input demand functions of these plants. As I demonstrate in the same paper, when other plant-specific state variables do affect investment, the O-P algorithm provides biased estimates of production function parameters.<sup>6</sup> The presence of additional unobserved plant-specific state variables in the investment function breaks down the one-to-one relationship between plant productivity and investment, so it is no longer possible to pin down productivity levels with investment observations.<sup>7</sup>

Instrumental variables techniques are a preferred alternative in such cases; they offer consistent estimates even with endogenous regressors. In practice, however, obtaining good instruments for plant-level production data can be a challenging task. Indeed, the call for methods such as the O-P algorithm grew out of a perceived lack of quality instruments. A suitable instrument must exhibit some variation across plants to gain any additional identifying power from the plant data. Aggregate or even industry-wide series will not suffice. It is this criterion that has caused many researchers who work with plant-level data to forsake the search for instruments as too difficult, if not hopeless. However, I believe that careful theoretical consideration of market structures allows identification of plant-specific instruments that fit my

---

<sup>6</sup> The Olley-Pakes algorithm also requires a similar assumption about the character of a plant's produce/liquidate decision which I contend can also lead to biases under market segmentation. This point is tangential to the discussion here, however, so I will not address it further. An interested reader should see Syverson (1999).

<sup>7</sup> Recent proposed modifications to the Olley-Pakes algorithm—such as Levinsohn-Petrin (1999), which advocates using materials rather than investment to back out productivity proxies—are also subject to problems when markets are segmented. In the case of the Levinsohn-Petrin modification, producers operating in segmented markets are also likely to account for local demand when making materials purchases. This eliminates any one-to-one mapping between establishment materials use and productivity and results in incorrect proxies.

needs, and incidentally, are obtained using an intuitive framework that can be used to identify instruments applicable to many other studies.

I contend that market segmentation—geographic segmentation here specifically—can be exploited to identify establishment-level instrument series. In this way, the very influence that limits the applicability of the O-P method in certain cases can be used to obtain consistent estimates. The key to identifying such instruments is recognizing how markets are segmented across the plants of interest. Market segmentation, for my purposes, refers to any way in which a seemingly industry- or economy-wide market is actually comprised of a collection of heterogeneous “local” market units. (“Local” does not necessarily imply that the market is geographically heterogeneous, although that is case here and for many other goods and industries.) That is, markets are segmented whenever there is some degree of plant-level separation in an industry's output or inputs markets. Recognizing such market heterogeneity allows identification of instrumental variables that will exhibit across-plant variation when measured along the same dimension as the segmentation is present.

The intuition behind the market segmentation method of instrument identification can be explained in the context of my current application. For production function estimation instruments, I need variables that are relevant to inputs in ready-mix plants (the explanatory variables in a production function) but uncorrelated to plant productivity levels (the residuals). Furthermore, as discussed above, the variables must exhibit some degree of across-plant variation.

Shea (1993) argues that measures of construction activity are relevant to inputs and orthogonal to productivity in many intermediate construction goods industries, at least at the industry level. Construction is relevant to the input levels of concrete plants because a large portion of industry output is used in final construction output; construction activity and ready-mix plant inputs are thus very likely to move together.<sup>8</sup> Furthermore, because construction projects generally require output from a wide array of industries, the percentage of total costs of final construction firms attributed to ready-mix alone is likely to be relatively small.<sup>9</sup> This small

---

<sup>8</sup> For example, firms engaged in new construction activity purchased 79.8% of 1977 ready-mix output. See Shea (1992).

<sup>9</sup> Looking at 1977 again, concrete accounted for 6.5% of new construction costs that year.

cost share makes it less likely that any productivity advances in the ready-mix concrete industry—which lower the relative cost of concrete—will alter the amount of construction activity, because idiosyncratic price drops in a single intermediate input will not greatly lower the total costs faced by final construction firms. Therefore, productivity movements in the industry are nearly (if not entirely) uncorrelated with final construction activity, satisfying the exogeneity criterion.<sup>10</sup>

My technique extends these instruments to the plant level by matching local construction activity measures to upstream-industry plants in the same geographic market. The high weight-to-value ratio of concrete makes it reasonable to assume that concrete plants sell the vast majority of their output locally, and thus make productive input decisions partly on the basis of their local demand state. Comprehensive shipments data from the 1977 Commodity Transportation Survey support this; ready-mix plants shipped 94.4 percent (by weight) of their total output less than 100 miles. Therefore local construction activity measures (which presumably reflect local demand for concrete) should be suitable plant-specific instruments. We can be reasonably confident that construction activity in, for instance, the Lincoln, Nebraska area will influence the input choices of concrete producers in Lincoln, but not those in, say, Tucson, Arizona. Conversely, fluctuations in Greater Tucson's construction business will not affect Lincoln plants. Therefore Lincoln- (Tucson-) area construction activity can be used to instrument for inputs in a Lincoln (Tucson) ready mix plant with reasonable assuredness that the relevance and exogeneity criteria are being met for each plant. If construction activity measures are spatially disaggregated enough, local activity measures will capture substantial interplant variance in the instrument series.

It is conceivable that, despite the small cost share of concrete in overall construction, productivity in ready-mix plants is still correlated with local construction activity if there are common local productivity shocks. For example, if there are urbanization spillovers affecting all industries in an area, these spillovers may boost overall construction activity while simultaneously increasing productivity levels in ready mix plants. Such a condition would of

---

<sup>10</sup> The idea of using input-output linkages to identify instruments was proposed by Shea (1993). His paper offers a more thorough discussion of how one can identify instruments at the industry level which are both relevant and exogenous using demand and cost shares.

course weaken the exogeneity of my instruments and lead to estimation biases. To eliminate this possibility, I do not use local construction activity measures directly to instrument for ready-mix inputs. Instead, I regress my construction sector activity measure (sector employment), on a measure of overall economic activity in the same region (total employment) and use the residual as my instrument. Thus I am instrumenting for ready-mix plant inputs with the component of local construction activity that is unrelated to overall activity in the region. This effectively removes the possibility of instrument endogeneity because of common regional productivity shocks.

### ***Local Markets in the Ready-Mix Concrete Industry***

To test the implications put forth by the earlier model, I must place the ready-mix plants in my sample into local markets. My chosen geographic market unit is the Bureau of Economic Analysis' Component Economic Area (CEA). CEAs are collections of counties usually—but not always—centered on Metropolitan Statistical Areas (MSAs). The BEA selects counties for inclusion in a given CEA based upon MSA status, worker commuting patterns, and newspaper circulation patterns (subject to the condition that a CEA contains only contiguous counties). This ensures that counties in a given CEA are substantially intertwined economically. These 348 markets are mutually exclusive and exhaustive of the land mass of the United States, so each is typically comprised of seven or eight counties.<sup>11</sup>

I choose the CEA as my unit of analysis because it is the best compromise between several conflicting requirements. The theoretical foundation of this study assumes that local concrete markets are essentially isolated geographic units, where plants in one market only competitively interact with other plants in their local area. Any interaction with ready-mix production units in other markets is assumed away. While there are bound to be some cross-border concrete sales in reality, the Commodity Transport Survey shipment data discussed previously testify to the high transport costs for the industry. These plants have very limited operations radii, so if I draw local markets sufficiently large enough, I can decrease the amount of cross-market sales occurring in my data. For this reason, defining individual counties as separate markets may be inappropriate; it is likely that a non-trivial fraction of ready-mix produced in the

---

<sup>11</sup> See U. S. Bureau of Economic Analysis (1997) for more detailed information about CEA creation.

county will be consumed outside of it. On the other hand, I do not want to make markets so large that there is very little competitive interaction between many of the included establishments. Plants placed in too large a market may not all respond to the same market forces (either external influences or the actions of industry competitors). CEAs are a suitable compromise between these two poles. Furthermore, because most are centered around MSAs or other population centers (and those that are not are composed of very rural counties), the counties on CEA borders are likely to be more sparsely populated with concrete plants. The bulk of ready-mix production in a market is then centrally located, decreasing the likelihood of between-market sales. CEAs are also not required to adhere to state boundaries, which would sometimes place unwarranted market boundaries in economically interconnected areas, such as exist in the Washington, D.C. area.<sup>12</sup>

To evaluate the effect of local market size on the number and size of area ready-mix plants, I estimate cubic functions of the number of establishments and average logged plant output in a market on the log of local market size (computed as the sum of all constant-dollar ready-mix shipments originating in the market). The data is taken from the 1982, 1987, and 1992 Census of Manufactures. The results are plotted in levels in Figure 4. As the figure shows, both the expected number and average size of ready-mix plants are monotonically increasing, concave functions of market size. Increases in the local demand for concrete are taken up on both the intensive and extensive margins, with the extensive margin becoming more dominant as market size increases. Thus suggests that the decreasing returns indicated in the production function estimates (seen below) do eventually set in for ready-mix plants. Curiously, these results are the converse of the theory's implications. Recall, as shown in Figure 1, that both the number (mass) and average size of plants grow as concave functions in market size, but the average plant size curve has the lesser curvature of the two functions and begins to dominate in large markets. Perhaps this is not surprising given the globally increasing returns in the theory's production technology. To compare these results to other theoretical work, a Salop (1979)-type circular market model with homogeneous producers implies that the number of producers increases with

---

<sup>12</sup> A further practical consideration favors the use of CEA-defined markets. Larger market areas increase the number of area establishments, allowing better estimation of productivity distribution moments, but decrease the total number of observations, decreasing the precision of estimates. Most CEAs contain an adequate number of ready-mix establishments to obtain plant-level productivity distribution moments, while still affording a sufficient

the square root of total market demand. This is a function more in accordance with the theoretical results than the data. This may be explained by the fact that marginal costs of production are constant at all production levels in that model as well. Findings on similar matters from other empirical research have depended on the industries studied. Campbell and Hopenhayn (1999) show that for a variety of retail industries, the number of establishments per capita decreases and average establishment sales increase with the local population. The former finding suggests that the number of establishments is a concave function of market size, and the latter implies a monotonically increasing function of average establishment sales in demand. These results support the general shape of the empirical and theoretical functions. In contrast, Dinlersloz (1999) finds that for two-digit manufacturing industries, the number of plants per capita and the employee size distribution of plants is invariant to local population. Both of these facts imply that the number of manufacturing plants is a linear function of population. Curiously, this would be an implication of the above model if the elasticity of substitution does not depend on the mass of producers in the market (see equation 19).

Comparing margins of adjustment as an industry's market grows has many interesting implications for the nature of competitive behavior (see Bresnahan and Reiss (1991), for example). Here, I restrict my attention to the implications of local market density effects on output substitutability. In my posited process, the ease of substituting one ready-mix plant's output for another increases as the number of plants in a given market area increases. Given the function for the number of plants as market size changes shown in Figure 4, an increase in demand density (i.e., a rise in total industry sales in a given area) should have the largest impact on substitutability—and thus the local plant-level productivity distribution—in smaller markets.<sup>13</sup> However, because most demand growth is taken up on the extensive margin in larger markets, this effect should still be present regardless of market size.

### III. Data

#### *Local Construction Activity Data*

The key to implementing the market segmentation principle of instrument identification

---

number of observations.

<sup>13</sup> As long as the growth of substitutability is not convex in plant density.

is instrument data that can be pored along the axis of market heterogeneity. The present case requires construction and all-industry activity data at a geographically disaggregate level. Such data does exist. I use local construction instruments derived from the Census Bureau's public-use County Business Patterns (CBP) annual data over the 1979-1993 period. The CBP contains Mid-March employment by major industry for every county in the United States. These employment values are my downstream activity measures. Public-use Census data at such a fine geographic resolution often have censored observations, but this is a very minor obstacle in the case of the construction sector (SICs 15-17). The sector's ubiquity and abundance of small firms allows full disclosure of summary statistics in all but the smallest of counties. For those counties with exact construction employment data withheld for the sake of confidentiality (roughly 1.5% of the county-year observations), a total employment range is reported. In those cases, I simply use the mean of the range as the imputed employment for the period. The impact of using imputed numbers is likely to be even less than their proportion indicates, as the typically small nondisclosure counties are less likely to contain sample plants.<sup>14</sup>

I also take advantage of the geographic dimension of the CBP survey to examine how changing the level of geographic aggregation of the construction activity data affects the instruments' relevance. I aggregate the instrument data at three geographic levels. The finest aggregation is at the county level, as the data are originally reported. In this case, downstream construction activity in a given county (the component independent of overall county activity) instruments for inputs at ready-mix establishments in that county. County activity is an extremely local measure, however, even for plants in locally focused industries. While many such plants do likely operate largely within one county, it is also highly probable that a significant fraction sell their output outside the boundaries of their county. This is especially true for larger establishments in multi-county metropolitan areas, and in the Northeast, where counties are simply smaller in area than their western counterparts. Multicounty activity measures may be more appropriate in such instances. I therefore also instrument using construction activity data aggregated at two broader levels. The first, and the smaller of the two geographically speaking, is at the CEA level. The third and highest geographic instrument

---

<sup>14</sup> CBP data also have annual industry payroll numbers that could also serve as activity measures. I found no systematic difference between results obtained using real payroll (not reported here) and those using employment.

aggregate I use is at the Economic Area (EA) level. The BEA combines CEAs that are considered themselves to be economically interconnected into 172 EAs. Construction sector and all-industry employment for these larger geographic divisions are simply the sums of the respective county-level values for all counties within the CEA or EA. I do lose some across-plant variation in the instrument set when I aggregate geographically, of course. The loss in identifying power may be a necessary trade-off in order to gain relevance in those industries with plants that largely operate beyond their counties' borders.

### ***Plant Level Production Data***

I take plant output and inputs data from the 1982, 1987, and 1992 versions of the Census of Manufactures (CM). The CMs (part of the Bureau's Longitudinal Research Database) contain a wealth of information on plant production activity. Importantly here, they also contain the state and county where the establishment is physically located, so it is possible to match each plant with local instrument values at all three geographic aggregation levels. My sample period was limited because of availability limitations of the annual CBP instrument data, which is only available for 1977 onward (I require three lags of instrument values for each input observation). Some small plants (typically with fewer than five employees) in the CM have imputed data for some variables. I exclude these plants from my sample, but do include their (non-imputed) shipments numbers when calculating total market size.

My benchmark estimated production function is expressed in terms of gross output. Yearly nominal gross output is a plant's reported total value of shipments plus an adjustment for changes in inventories of final goods over the year. Nominal output is converted to a real value by dividing by an output price deflator for the plant's corresponding four-digit industry, taken from the Bartelsman-Becker-Gray/NBER Productivity Database. Using deflated plant sales as measures of real output is subject to measurement error when establishment-level prices vary within an industry. In testing for any problems of this sort (which will be discussed below), I also use the physical product data available in the CM. This data offers physical production numbers broken down by seven-digit SIC products at the plant level.

Producer labor inputs are the sum of production worker hours (a reported value in the CM) and an imputed value for nonproduction worker hours. Nonproduction worker hours are

constructed using the method of Davis and Haltiwanger (1991), where the number of nonproduction workers at the plant is multiplied by the average annual hours worked by nonproduction employees within the corresponding two-digit industry and year. These latter values are based on Current Population Survey data.

Real investment for each plant is calculated simply by dividing reported equipment and structures investments (the CM contains separate capital data for both of these capital types) by the respective Bureau of Labor Statistics (BLS) two-digit type-specific capital deflator.

Equipment and structures capital stocks for each plant are the establishment's reported book value capital stocks multiplied by the ratio of book to real values for the entire three-digit industry in that year. Industry-level capital stocks are from published BEA data. The value of any reported machinery or building rentals is inflated to a capital stock by dividing by the BLS's rental cost of capital series for the respective capital type. The total capital stock used in production function estimation is constructed by summing the equipment and structures stocks.

Real materials usage is plant materials costs divided by a corresponding four-digit materials deflator. Energy input is the sum of electricity and fuel expenditures deflated using a four-digit energy cost index. Each of the industry-specific price deflators used in this process is taken from the Bartlesman-Becker-Gray Productivity Database.

Input cost shares used to construct a composite input are computed as follows. Establishment labor costs are the sum of total salaries, wages, and benefits paid to permanent workers plus any costs from hiring contract labor. Capital costs are the product of establishment capital stocks and the BLS capital rental cost series. Energy costs are the sum of electricity and fuel purchases, and materials costs are a separately reported item in the CM. I sum the four to obtain total costs, and calculate shares using this value. Each input is weighted in the composite input by the average cost share in the ready-mix industry over the current and previous CM years.

## IV. Empirical Results

### *Production Function and Productivity Estimation*

The prerequisite for my empirical work is estimation of an industry production function. As mentioned above, the commonly used Olley-Pakes procedure may not be appropriate for the ready-mix industry because the geographic market segmentation present in the industry makes it

very likely that concrete producers take their local demand state into account when making their investment decisions. This makes it impossible to back out accurate productivity proxies because there is no one-to-one mapping between plant productivity and investment. In Table 1, I present evidence to this effect. The table shows relevance statistics obtained from regressing the investment levels of ready-mix plants (those observations with nonzero investment) on the instrument set I use in estimating the production function. Each investment observation projected on the current value, three lags, and one lead of the annual construction activity measure (already cleansed of overall regional effects).<sup>15</sup> I also include year dummies to remove industry-wide and aggregate effects and a dummy indicating whether the plant belongs to a multiplant firm. The table shows F-statistics for the joint significance of the five construction activity terms and the  $R^2$  of the regression, estimated both with and without plant effects. The demand terms are hugely significant statistically and, I believe, relevant economically. Producers are clearly taking their local demand state into account when making investment choices. The influence is especially strong once I account for plant effects. Because the demand instruments should be orthogonal to plant productivity levels, the regressions imply that there is a substantial influence of downstream demand on investment that is independent of plant productivity. This breaks any one-to-one correspondence between productivity and investment that the O-P algorithm could exploit to obtain a productivity proxy. Instrumental variables estimates are preferred for the present application.

To obtain plant productivity values, I estimate the following production function:

$$y_{it} = \gamma_0 + \delta_t + \gamma_d d_{mult} + \gamma_x x_{it} + \omega_{it}$$

where

$$x_{it} = s_l l_{it} + s_k k_{it} + s_m m_{it} + s_e e_{it}$$

and  $s_j$  is the plant-level cost share of input  $j$  during period  $t$  (actually the average of the industry cost share in the present and preceding CM year). All continuous variables are measured in natural logarithms. The production function specification includes year dummies to estimate  $\delta_t$  and a multiplant dummy  $d_{mult}$  that captures effects from operating as part of a multiple-establishment firm. Instead of entering the four inputs (labor, capital stock, materials, and

---

<sup>15</sup> This lag/lead pattern was chosen based on two considerations. The first is my prior belief about the extent of management decision horizons, both forward- and backward-looking. The second consideration is Buse's (1992) demonstration that superfluous instruments in an instrument set lead to estimation biases. The resulting

energy) separately into the function, I use an industry-cost-share-weighted composite input  $x_{it}$ . Under the assumption of cost minimization, the estimate of  $\gamma_x$  is the degree of returns to scale. The plant-specific productivity level is therefore  $\omega_{it}$ .

I use a composite input rather than the four individual components because of practical estimation considerations. While the market-segmentation instruments can be theoretically applied toward estimation of any functional specification, there is an issue hampering such efforts. As Shea (1997) demonstrates, instruments should not only be relevant to each of the individual endogenous explanatory variables, they should have *linearly independent* relevance. This implies in this case that downstream activity measures should have an influence on each input (labor, capital, energy, and materials) that is independent of their influence on the other inputs. While some independence may be gained through the ability of the lag/lead structure of the instrument set to capture differing dynamic impacts across input demand functions, the high degree of comovement in the response of the inputs to downstream demand may overpower any such effect. This difficulty was realized in practice; attempts to estimate a Cobb-Douglas specification using IV methods yielded unacceptably high standard errors.<sup>16</sup> The necessity of linearly independent relevance is obviously not an issue when using a composite input. The composite input specification offers the further advantage of not imposing a specific functional form on the production function.

Table 2 shows production function estimates obtained using the local construction activity instruments in a two-stage least squares procedure. I present estimates for each of the three geographic aggregation levels. Through the remainder of the paper, I will use estimates derived from the CEA instruments so as to be consistent in my market area definition. It is unlikely this choice will much change the nature of my findings; as can be seen in the table, the production function estimates are remarkably consistent across instrument sets.

The first stage relevance statistics in Table 2 clearly indicate that even after controlling

---

lag/lead structure is a reconciliation of these two factors.

<sup>16</sup> This does not mean that instruments obtained via market segmentation can only be applied to certain functional forms. The current restriction results from the fact my instrument set here, downstream demand shifts, tends to move several inputs simultaneously. If I could obtain additional instruments that have influence on specific inputs, I could add these to the instrument set and possibly gain linearly independent influence across inputs, allowing separate technology parameter estimation by input. Time and data constraints leave me to only use downstream measures in the instrument set for now. I leave expansion of the set to future work.

for overall local economic activity and aggregate effects, local construction activity is germane to concrete plant inputs. The F-statistic for joint significance of the five construction activity terms is highly significant. The first stage results not only make the case for statistical relevance, but economic relevance as well. Local construction activity explains just under five percent of the input variation across concrete plants. This value compares favorably to results in other studies that use largely cross-sectional establishment panels. The remainder of the table shows the production function estimates, which are quite reasonable. Average industry productivity (indicated by the constant and the year dummies) increased around seven percent from 1982 to 1992, but most of this movement occurred by 1987. These values are almost surely affected by the trough-peak-trough nature of the observed years in the sample. The coefficient on the multiplant firm dummy indicates that plants affiliated with other corporate establishments have slightly higher productivity levels on average. The composite input coefficient is quite precisely estimated, and it indicates very slight decreasing returns to scale in the industry.

Plant-specific productivity estimates are simply the residuals of the estimated production function; I do not convert them to levels by taking their exponential. I use these estimates in the following exploration.

### ***Local Market Demand Density and Productivity Distributions***

The central question of this paper is whether substitutability factors in markets (as embodied in local demand density) affect plant-level productivity distributions. The two distribution moments that I am most concerned with are dispersion and central tendency. I will use measures of these moments selected to account for specific measurement concerns. I measure dispersion as the interquartile range of the local productivity distribution. An ordinal dispersion measure is employed to minimize spurious influence from outliers. This is not an uncommon practice; see Roberts and Supina (1997), for example. Outliers are a special concern in this study for two reasons. With establishment data, it is fairly easy for measurement and reporting error to creep into the data and create nonsensical observations. Additionally, the fact that some of the markets have a fairly small number of plants increases the vulnerability of traditionally calculated moments to outlier effects. I choose the interquartile range rather than another quantile span because broader percentile ranges would not be any different in small

markets, and because wide spans, despite being ordinal measures, are also more subject to outliers in small markets. I measure the central tendency of the productivity distribution in two ways. The first is the median productivity level in the market. Again, I choose an ordinal measure to minimize measurement problems. The second, the market's output-share-weighted average productivity, takes the output distribution explicitly into account. This measure is more vulnerable to the influence of outliers, of course, but captures whether output is reallocated to more productive producers as demand density increases, one of the implications of the theory.

The use of Component Economic Areas as local market units offers a potential number of observations of 348 CEAs x 3 years = 1042. In the benchmark results, I use only those market-year observations with at least five plants in order to improve moment measurement accuracy. I will test the results for robustness to this cutoff. Half of the CEA markets have at least ten plants in each of the three sample years (1982, 1987, and 1992). The 25<sup>th</sup>-percentile number of plants is six in 1982 and five in 1987 and 1992.

The empirical specification used to test for the impact of demand density on the local productivity distribution is as follows:

$$prodmom_{it} = \beta_0 + \beta_d dens_{it} + X_{cit} B_c + \varepsilon_{it}$$

This specification requires that the plant-level productivity distribution moment in market  $i$ , year  $t$  (either a measure of dispersion or central tendency) is a function of a constant, the local demand density  $dens_{it}$ , a vector  $X_{cit}$  of other influences on the moments, and an industry-specific error term. I estimate four versions of this general model. A simple univariate regression of the productivity moment on logged demand density (constant-dollar sales for all producers in the market divided by its land area) is run to characterize the nature of the correlation between these variables. I then add the vector of local demand influences  $X_{cit}$  to control for the influence of other variables that may affect local demand. Finally, both of these models are rerun with year dummies to remove any industry-wide influence on local productivity distributions.

The vector of local demand controls  $X_{cit}$  contains an assortment of variables that plausibly impact the state of the regional ready-mix concrete market, and hence may impact the local establishment-level productivity distribution.<sup>17</sup> I include a set of variables characterizing the

---

<sup>17</sup> Most of these control variables change over time in actuality; however, some are time-invariant measures here due to data limitations. For these controls, I have attempted to use values gathered as close to the middle of the sample period as possible.

demographics of the CEA: the percentage of the population that is nonwhite, the fraction over 25 years old, the proportion with at least a bachelor's degree, and the number of marriages per 1000 population. Each of these variables is taken from the 1988 version of the City and County Data Book, which compiles data from a number of Census Bureau surveys at geographically disaggregate levels. The race and the marriage variables are 1984 values, while the others are from the 1980 population census. I also include variables that are likely correlated with the concrete demand specifically, including the fraction of households with at least two automobiles, the fraction of housing units that are owner-occupied, the median value of owner-occupied housing, and median personal income (also from 1980 and 1984). To capture any urbanization effects that have not already been removed by regressing the construction activity instruments on overall employment, I add Ciccone and Hall's (1996) measure of local employment density for the CEA, computed using 1986 civilian employment numbers. I also include the growth rate of ready-mix output over the previous five years to control for short-term effects on the distribution (for example, a temporary boom might allow relatively inefficient producers to operate for a short while). Lastly, the average primary product specialization ratio (PPSR) of the ready-mix plants in the region and year is included. I found in Syverson (2000) that physical product differentiation strongly affects industry productivity distributions. Controlling for PPSR differences across market areas should remove much of any product differentiation impact.

The summary statistics in Table 3 indicate that there are nontrivial differences in productivity moments across local markets. For example, a market having a median productivity that is one standard deviation larger than another's has a median productive ability that is 13% higher (when characterized in terms of output levels). The standard deviation of productivity dispersion is two-thirds of its average.

The main results of the paper are presented in Table 4. The table shows, for each specification and local ready-mix productivity distribution moment (dispersion, median, and quantity-weighted average productivity), the estimated demand density coefficients and robust standard errors. I do not report covariate estimates in the interest of parsimony.

The results unambiguously support my assertions. Productivity dispersion declines with market density, and median productivity and the quantity-weighted productivity increase as

ready-mix sales per square mile increase. These results hold for every moment and for each of the four models, and the coefficients are in all cases statistically significant at the 5% level. The estimates are economically as well as statistically significant. The estimates imply that a one-standard-deviation increase in logged sales density will decrease expected dispersion by approximately 0.029 points—roughly 11% of the mean dispersion and one-sixth of its standard deviation (see Table 3). The same density increase raises the median productivity level in the market by roughly a quarter of a standard deviation (equivalently, a 3.4% productivity increase in terms of output levels) and quantity-weighted productivity by one-sixth of a standard deviation.

Adding market demand-condition controls to the regression does change the magnitudes of the estimated coefficients. The estimated magnitude of the negative effect on dispersion increases once local demand conditions are accounted for. The median productivity and the quantity-weighted productivity estimates decrease, the latter significantly so. Still, even after accounting for these influences, the impact of demand density remains. The coefficients remain significant in all cases.

Interestingly, it seems that transport-cost-driven substitutability explains only a modest portion of the differences in local productivity distribution moments across markets. The  $R^2$  for the univariate regressions indicate that demand density differences alone account for roughly 2% of the across-market dispersion of productivity dispersion. The ability of density to explain productivity levels is stronger, but still moderate. These modest values (despite being common in primarily cross-sectional studies) are surprising, given the perceived level of homogeneity in ready-mix output.

### ***Robustness Checks***

In this section I will test the robustness of the central findings to many of the empirical modeling assumptions made above.

*Minimum Number of Establishments.* Table 5 shows estimates obtained when using different minimum-establishment cutoff numbers for selecting the observations used in the regressions. Panel A contains the results using any CEA-year observation with more than one plant. The estimates are strikingly similar to those from the five-producer cutoff. The only estimates that

change appreciably are those from the models including demand controls for the dispersion and quantity-weighted productivity moments. The estimated influence on weighted productivity increases. The drop in the estimated dispersion-moment coefficient and loss of statistical significance can perhaps be explained by the fact that dispersion measures become less meaningful in markets with very few plants. Including these small markets may be introducing measurement problems into the dispersion moment observations.

This assertion is supported by the results in Panel B of Table 5. Using a higher cutoff number of plants—ten in this case—returns the two wayward dispersion coefficients to their earlier levels, suggesting that dispersion measurement does suffer somewhat in markets with few plants. Two each of the median and quantity-weighted productivity estimates—those from the regressions that include controls for demand conditions—drop from their levels in the benchmark estimation. Given this and the fact that coefficients from the specifications using only density (with and without year effects) remain largely the same, it could suggest that some of density’s impact on these two moments declines in larger markets. However, if this story is to be believed, we would have to account for the fact that the dispersion influence remains strong.

While some changes in estimates occur with other minimum plant number criteria, it is apparent that the contents of Table 5 largely support the central premise of the paper. In all cases, the estimated sales density effects are of the expected sign. Furthermore, the majority of estimates are statistically significant.

*Output Measure.* An important issue arises in the use of plant-level production data that I must be mindful of here. The results reported above are obtained using the measure of plant output most commonly used in establishment-level studies: plant revenue deflated by an industry-wide price index. If there is price dispersion within the industry, it is readily apparent that such practice can lead to incorrect output measurement. Those establishments that sell output at a price above the industry average will have calculated outputs that are larger than the true value, while those selling at a lower price will have undercounted output. This deviation of measured from actual output enters into the error term of a production function; i.e., it is included in the plant’s estimated productivity level. In this way, price variation between plants can induce spurious productivity dispersion (unless plant-level prices completely and solely embody quality

differences, in which case we would want the departure from the average industry price to enter into productivity measures).

Possible implications stemming from output mismeasurement of this form, while recognized by many, have largely been ignored in other studies because of difficulties in properly accounting for their influence. There are some exceptions, however. Particularly relevant to this paper is the research of Klette and Griliches (1996), which demonstrates how price-induced measurement error can bias common production function estimates. If plant-level prices are correlated with my instruments, the instruments are no longer orthogonal to the production function residual (now containing the measurement error). This is a distinct possibility; it is quite plausible that plant-level prices are positively correlated with plant-specific demand. This could lead to biases in my production function estimates, and therefore in my estimated productivity values as well.

I address this concern in two ways. Both strategies take advantage of product data available in the Census of Manufactures, which includes information about production in physical units and sales of specific products. Ready-mix plants report their production in thousands of cubic yards. With the available product-specific sales data, I can compute unit prices for the plants in my data set. The first strategy compares the strength of the comovement between my demand instruments and plant-level prices to that between the instruments and plant inputs. The second (and more direct) tactic involves redoing the entire production function and productivity estimation procedures using physical output instead of deflated revenue.

Both exercises suggest quite conclusively that price dispersion is not driving the results above. One piece of evidence comes from comparing the statistical and economic impact of local demand on producer prices and inputs. In a regression of plant-level logged unit prices on the instrument set, the five local demand variables have a marginal  $R^2$  (the additional explanatory power gained by adding the demand terms to the year and multiplant dummies in the regression) of 0.004. The value of the F-statistic for joint significance of the demand terms is 10.76. Comparing these statistics to their corresponding values for the regression of plant inputs on the instrument set, a marginal  $R^2$  of 0.026 and an F-statistic of 63.74, it is clear that any influence from price dispersion is dwarfed by the correlation with inputs.<sup>18</sup> Thus, while local demand does

---

<sup>18</sup> This set of results is actually the weakest of several specifications. Using other geographic demand

statistically influence plant prices, this impact is small compared to that of demand on plant inputs. This highly elastic supply condition minimizes correlation between my instruments and any output measurement error that enters into the production function residual, serving to curtail possible estimation biases.

Strong confirmation of these findings from another methodology can be found in Table 6. This table shows results from the production function and productivity distribution moment regressions using physical rather than deflated revenue measures of output measures of output.<sup>19</sup> The salient feature of the table is its similarity to the benchmark results. The returns-to-scale estimate in Panel A is quite close to that shown in Table 2, suggesting that any price-induced bias in the production function and productivity estimations is small. (Note that the other coefficients, as intercept terms, will differ in magnitude from their benchmark counterparts because physical output is measured in different units than deflated revenue. This will be true of the demand density coefficients in the productivity distribution moment regressions as well). The coefficients reported in Panel B are all of the expected sign and statistically significant. This evidence and the other results reported in this section make it obvious that the benchmark findings are not simply a function of price-induced output measurement error.

*Calculated Productivity.* I also obtain a set of results, shown in Table 7, using calculated (rather than estimated) plant productivity measures. Productivity here is the log of deflated plant revenue minus a weighted sum of logged inputs. Weights for each of the four inputs (capital, labor, energy, and materials) are the average cost shares of the respective inputs over the current and previous CM years. Again the results bolster the robustness of the benchmark findings. The influence of local sales density on the productivity distribution moments is in the expected direction in each regression, and the coefficients are all statistically significant.

---

aggregations and adding plant effects highlighted the gap between local demand's influence over plant price and input levels. For instance, adding plant effects yielded marginal  $R^2$  and F-statistics of demand on inputs that were roughly ten times larger than their demand-price counterparts.

<sup>19</sup> Some SIC 3273 plants produce seven-digit products other than ready-mix concrete. This fact could lead to its own output measurement problems when plants differ in the percentage of their total output accounted for by ready-mix. Two factors minimize any such problem. First, ready-mix plants tend to be quite specialized; the average primary product specialization ratio (i.e., the percentage of total shipments accounted for by ready-mix) for industry plants is roughly 95%. Second, I divide plants' ready-mix production by their primary product specialization ratio to adjust all plants to a common output scale.

*Plant-Specific Effects.* I have purged my local concrete demand instruments of the influence of overall activity in the region to remove any region-wide aggregate productivity effects that may cause endogeneity between the instruments and concrete producers' productivities. However, it is possible that there may still be systematic differences in productivity levels across ready-mix plants that are correlated with local demand conditions, and hence introduce bias into my production function estimates. To address this concern, I repeat both the benchmark and physical product estimation sequences while including plant-specific fixed effects. The results are shown in Table 8.

Panel A contains the production function estimates for both output measures. The returns-to-scale estimate obtained using deflated output is notably larger than its counterpart in Table 2. On the other hand, the estimate from the physical output production function is virtually the same as its no-plant-effects analog in Table 6. The productivity impact of being in a multiplant firm is substantially diminished once plant effects are accounted for, suggesting perhaps that while more efficient plants tend to be in multiplant firms, the causal productivity benefit of belonging to a large firm is modest.

Panels B1 and B2 show results from the productivity distribution moment regressions.<sup>20</sup> The former panel shows the only sign seen so far that the benchmark results are less than completely robust. Local demand density is still found to have a statistically significant negative impact on productivity dispersion (of even slightly larger magnitude than the benchmark case), but the effect of market density on average productivity levels shrinks considerably. While all of the estimates have the expected sign, they are mostly insignificant statistically. The larger scale effects found in the corresponding production function estimates may be driving these results. Higher-density markets will tend to have greater total sales and thus larger ready-mix plants on average. Because of the estimated scale efficiency, the production of these large plants will be attributed to scale effects rather than plant-specific productivity. This may be removing much of the correlation between local sales density and average productivity seen above.

On the other hand, the results laid out in Panel B2 give no evidence weakening earlier

---

<sup>20</sup> Here, productivity values are still constructed as the production function residual using the output and input values in levels (not deviations from plant averages), but the production function coefficients are those from

findings. All estimates are strongly supportive of the theoretical conjectures when physical output measures are used. The productivity dispersion results are virtually identical to their analogs in Table 6, and the estimated impact on productivity levels is even greater than those found without accounting for plant effects.

The unusually weak correlation between local demand density and productivity levels found by adding plant effects to the benchmark case is curious. Considering the findings (consistent with the benchmark) from the physical output regressions with plant effects, perhaps fixed effects uncover an interaction between the deviation of prices and local demand from their plant means. Then again, the consistent support of the main results from the other sets of robustness checks may imply the weak correlation is simply a statistical anomaly.

*Technology Differences.* One of the advantages of using a single-industry case study to examine the link between substitutability and plant-level productivity distributions is the elimination of any impact from between-industry technology differences. However, even within a narrowly defined industry with largely similar production methods across plants, it is possible that some difference in production technologies exists. For example, the primary widespread technological innovation in the ready-mix concrete industry over the sample period was the transition from manual to automatic batching (the mixing of concrete orders according to a “recipe”). If this innovation spread unevenly over time or geography, plants may be operating simultaneously under different technologies.

To gauge the influence of any such effects, I run a specification where I have added three technology controls to all of the previously discussed demand controls. These control variables are specific to each region and year. The first is the fraction of plants in the region that operate as units of a multiplant firm. It is possible that plants that operate as part of a larger firm have additional advantage in securing investment financing and thus may be able to more easily obtain new production technology. Thus regions dominated by multiplant firms may be more likely to operate on the technological edge of the industry, which could affect moments of the local producer productivity distribution. The second control is the average capital-to-labor ratio of ready-mix producers in the region. This variable should capture much of any possible capital-

---

the fixed effects specification in Panel A.

embodied technology differences between regions. The third technology control is a measure of the average real wage in the region. This variable is constructed from County Business Patterns data for establishments in all industries, not just concrete. All-industry wages are used to capture whether an area is a high- or low-wage area without confounding the specific choices made by ready-mix plants in their labor purchases that may not be correlated with technology.

Table 9 compares the results for the full model (year dummies and demand controls included) with and without the technology controls. It appears that across-plant technology differences are not significantly affecting the results. Adding technology controls does not change the direction or significance of demand density's estimated influence. The magnitude of its impact diminishes slightly for dispersion and increases for average productivity levels, but neither of these changes is significant.

## V. Conclusion

I have posited that output substitutability differences across geographically segmented markets change plant-level productivity distributions in intuitively predictable ways. The results above seem to strongly support this assertion. Evidence from ready-mix concrete producer data shows that increases in demand density (which are arguably tied to substitutability increases) tend to decrease local productivity dispersion, increase average productivity levels, and reallocate output to more productive plants. These findings were quite consistent regardless of specific empirical modeling assumptions. The driving mechanism was explained theoretically as the effect of market density on producer density, and hence on within-market substitutability. The increased ease with which customers could switch suppliers truncated the lower end of the distribution of producer productivity levels.

The findings of this case study have several implications. Most directly, they suggest a role played by transport costs in accounting for some of the persistent productivity dispersion observed in the data. Further, transport costs induce a between-plant increasing returns-to-scale pattern based on selective survivability across market areas: larger plants tend to be in larger markets, and larger markets tend to have higher average productivity levels. The results may also say something about the impact of the continuing secular decrease in unit transport costs and its potential for decreasing productivity dispersion. This may be an interesting avenue for future

research.

Less directly but perhaps more broadly applicable, the paper's findings bolster results found using interindustry variation in measurable substitutability factors to explain differences in the productivity distributions of industry producers, as in Syverson (2000). Other factors that limit output substitutability besides transport costs, such as product differentiation, can be pointed to as sources of persistent efficiency differences. These other influences work through the same basic mechanism as transport costs to change producer productivity distributions.

The empirical evidence also suggests that much work remains to be done to completely characterize the nature and sources of productivity dispersion. Even in the "controlled" environment of an industry case study, observable factors still only account for a fraction of the observed variance of productivity distribution moments. There is still an enormous amount of productivity heterogeneity being caused by factors beyond observable technology and plausible output-market influences. Perhaps this is a strong statement about the role of unmeasured (and in many cases, unmeasurable) product differentiation, such as subjective product differentiation and bundled abstract goods, in explaining why we see such stark efficiency differences across plants.

## References

- Aw, Bee-Yan; Chen, Xiaomin and Roberts, Mark J. "Firm-Level Evidence on Productivity Differentials, Turnover and Exports in Taiwanese Manufacturing." NBER Working Paper no. 6235, October 1997.
- Bartelsman, Eric J. and Doms, Mark. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature* (forthcoming), 2000.
- Bresnahan, Timothy F. and Reiss, Peter C. "Entry and Competition in Concentrated Markets." *Journal of Political Economy*, 99(5), 1991, pp. 997-1009.
- Buse, A. "The Bias of Instrumental Variables Estimators." *Econometrica*, 60(1), 1992, pp. 173-80.
- Campbell, Jeffery R. and Hopenhayn, Hugo A. "Market Size Matters." Mimeo, University of Rochester, October 1998.
- Ciccone, Antonio and Hall, Robert E. "Productivity and the Density of Economic Activity." *American Economic Review*, 86(1), 1996, pp. 54-70.
- Davis, Steven J. and Haltiwanger, John. "Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-86." *Brookings Papers on Economic Activity: Microeconomics*, 1991, pp. 115-80.
- Dinlersloz, Emin M. "Agglomeration and Establishment Size in U.S. Manufacturing." Mimeo, University of Rochester, October 1999.
- Dunne, Timothy, Roberts, Mark J., and Samuelson, Larry. "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, 104(4), 1989, pp. 671-98.
- Griliches, Zvi and Mairesse, Jacques. "Production Functions: The Search for Identification." NBER Working Paper no. 5067, March 1995.
- Haltiwanger, John C. "Measuring and Analyzing Aggregate Fluctuations: The Importance of Building from Microeconomic Evidence." *Fed. Reserve Bank of St. Louis Review*, 79(3), pp. 55-77.
- Hotelling, H. "Stability in Competition." *Economic Journal*, 37(1), 1929, pp. 41-57.
- Klette, Tor Jakob and Griliches, Zvi. "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous." *Journal of Applied Econometrics*, 11(4), 1996, pp. 343-61.
- Jovanovic, Boyan. "Selection and the Evolution of Industry." *Econometrica*, 50(3), 1982, pp.

649-70.

- Levinsohn, James and Petrin, Amil. "When Industries Become More Productive, Do Firms? Investigating Productivity Dynamics." NBER Working Paper no. 6893, January 1999.
- Marschak, Jacob and Andrews, William H. "Random Simultaneous Equations and the Theory of Production." *Econometrica*, 12(3/4), 1944, pp. 143-205.
- Melitz, Marc J. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." Mimeo, University of Michigan, November 1999.
- Olley, Steven G. and Pakes, Ariel. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(4), 1996, pp.1263-97.
- Roberts, Mark J. and Supina, Dylan. "Output Price and Markup Dispersion in Micro Data: The Roles of Producer Heterogeneity and Noise." NBER Working Paper no. 6075, June 1997.
- Shea, John. "The Input-Output Approach to Instrument Selection: Extended Table III." Mimeo, University of Wisconsin, Summer 1992.
- Shea, John. "The Input-Output Approach to Instrument Selection." *Journal of Business and Economic Statistics*, 11(2), 1993, pp. 145-55.
- Shea, John. "Instrument Relevance in Multivariate Linear Models: A Simple Measure." *Review of Economics and Statistics*, 79(2), 1997, pp. 348-52.
- Syverson, Chad. "Production Function Estimation with Plant-Level Data: Olley and Pakes or Instrumental Variables?" Mimeo, University of Maryland, April 1999.
- Syverson, Chad. "Output Market Segmentation and Productivity Heterogeneity" Mimeo, University of Maryland, January 2000.
- U.S. Bureau of the Census. *1977 Census of Transportation, Commodity Transportation Survey*.
- U.S. Bureau of Economic Analysis. "Redefinition of the BEA Economic Areas." *Survey of Current Business*, February 1995, pp.75-81.
- U.S. Bureau of Labor Statistics. *Technology and Labor in Five Industries: Bakery Products, Concrete, Air Transportation, Telephone Communication, Insurance*.
- Weitzman, Martin L. "Monopolistic Competition with Endogenous Specialization." *Review of Economic Studies*, 61(1), 1994, pp. 45-56.

Figure 1: Mass of Producers and Average Plant Revenue vs. Market Size  
*Computed Equilibrium*

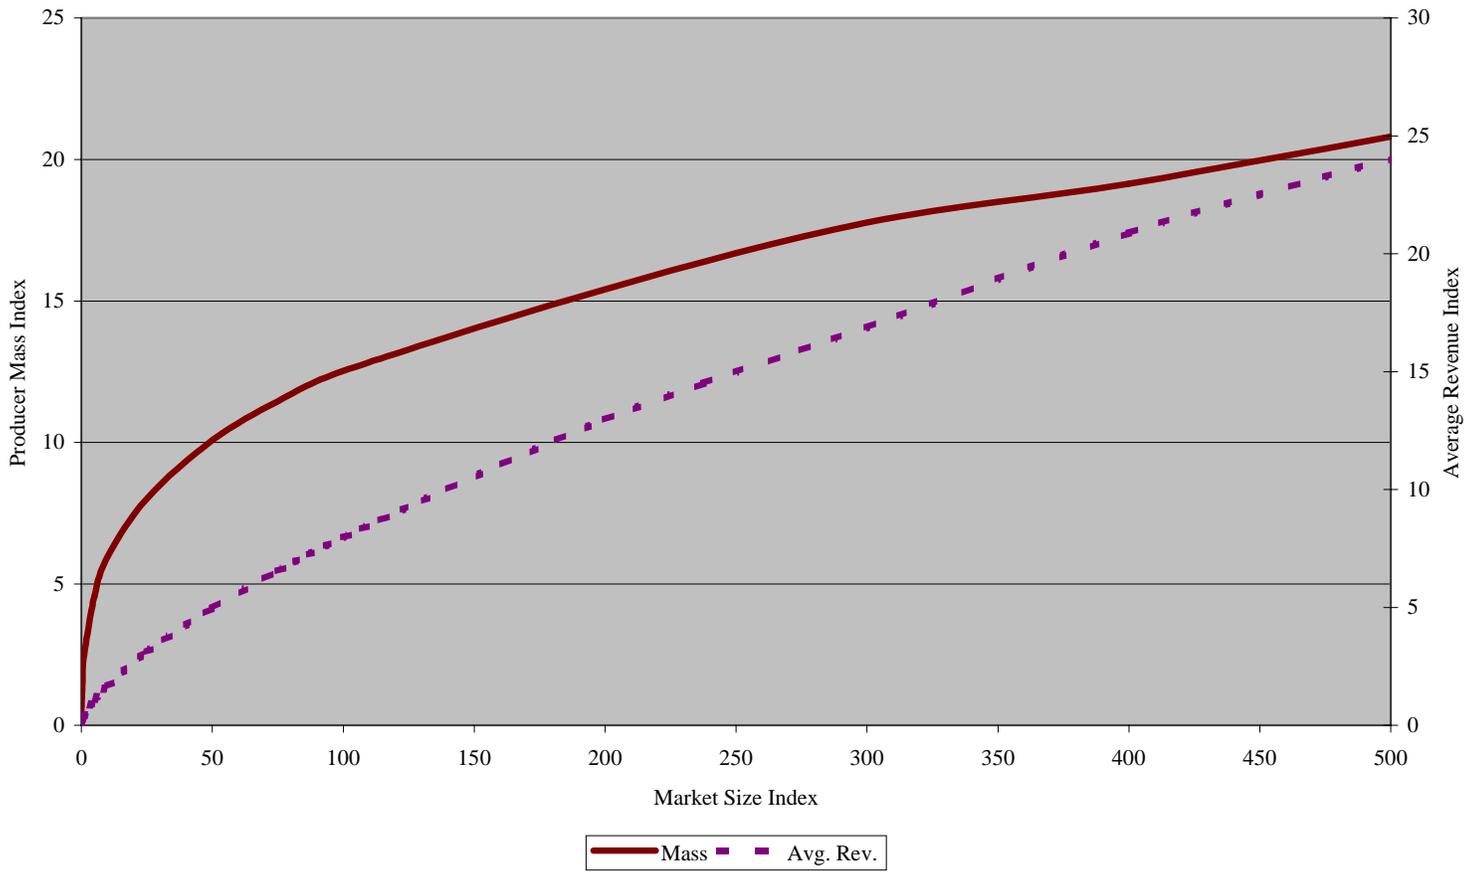


Figure 2: Cutoff, Average Productivity Levels and Elasticity of Substitution  
vs. Demand Density  
*Computed Equilibrium*

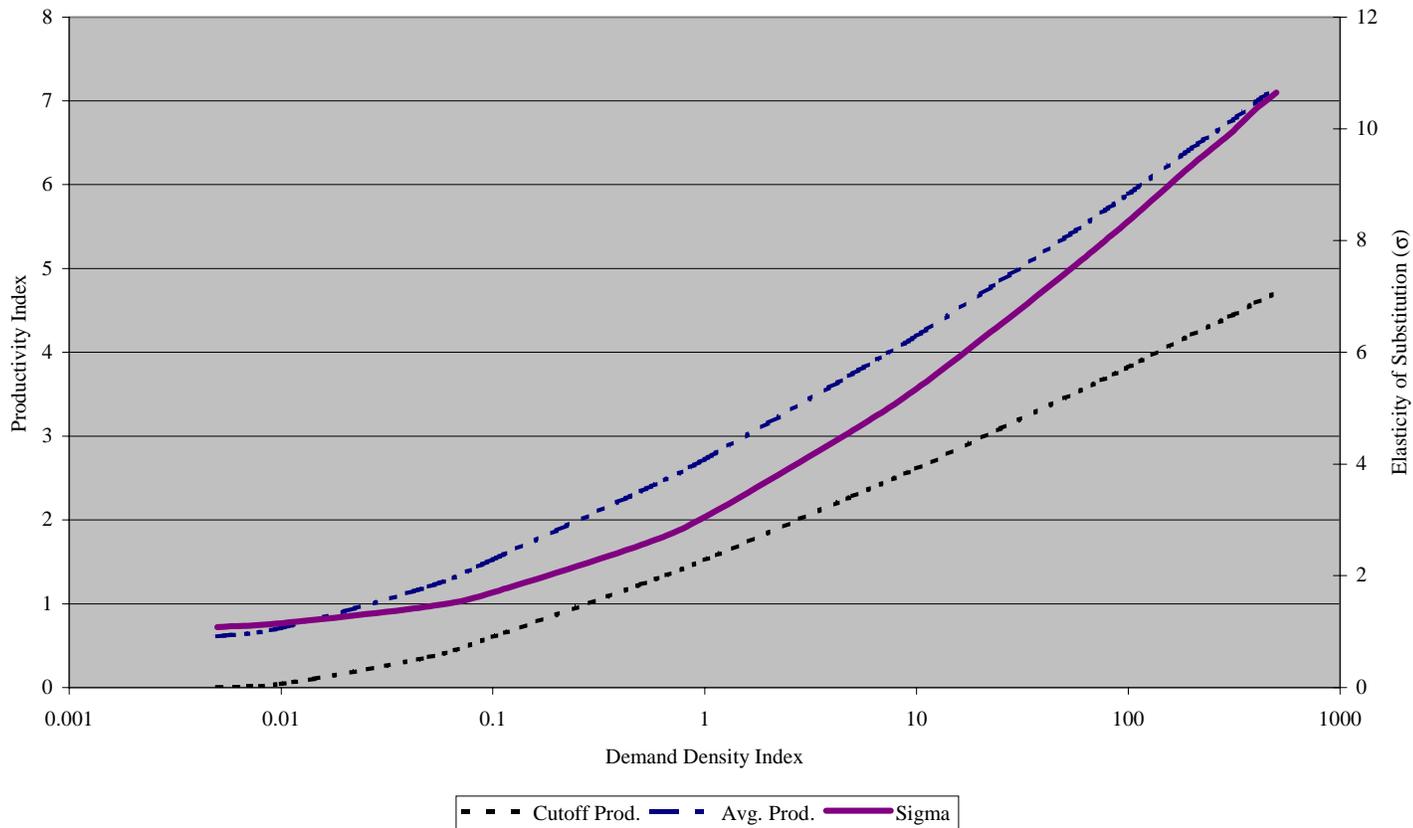
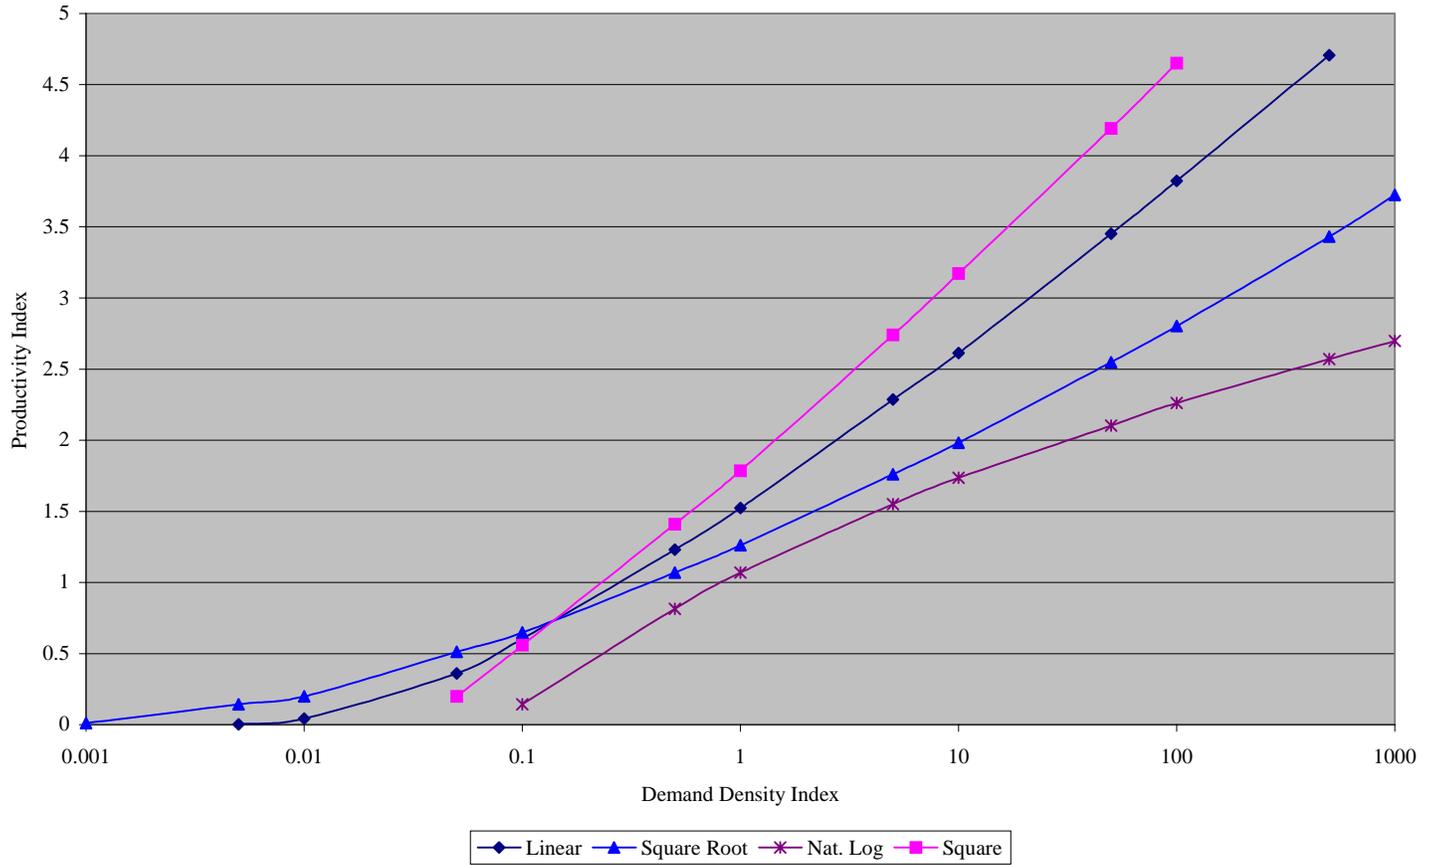
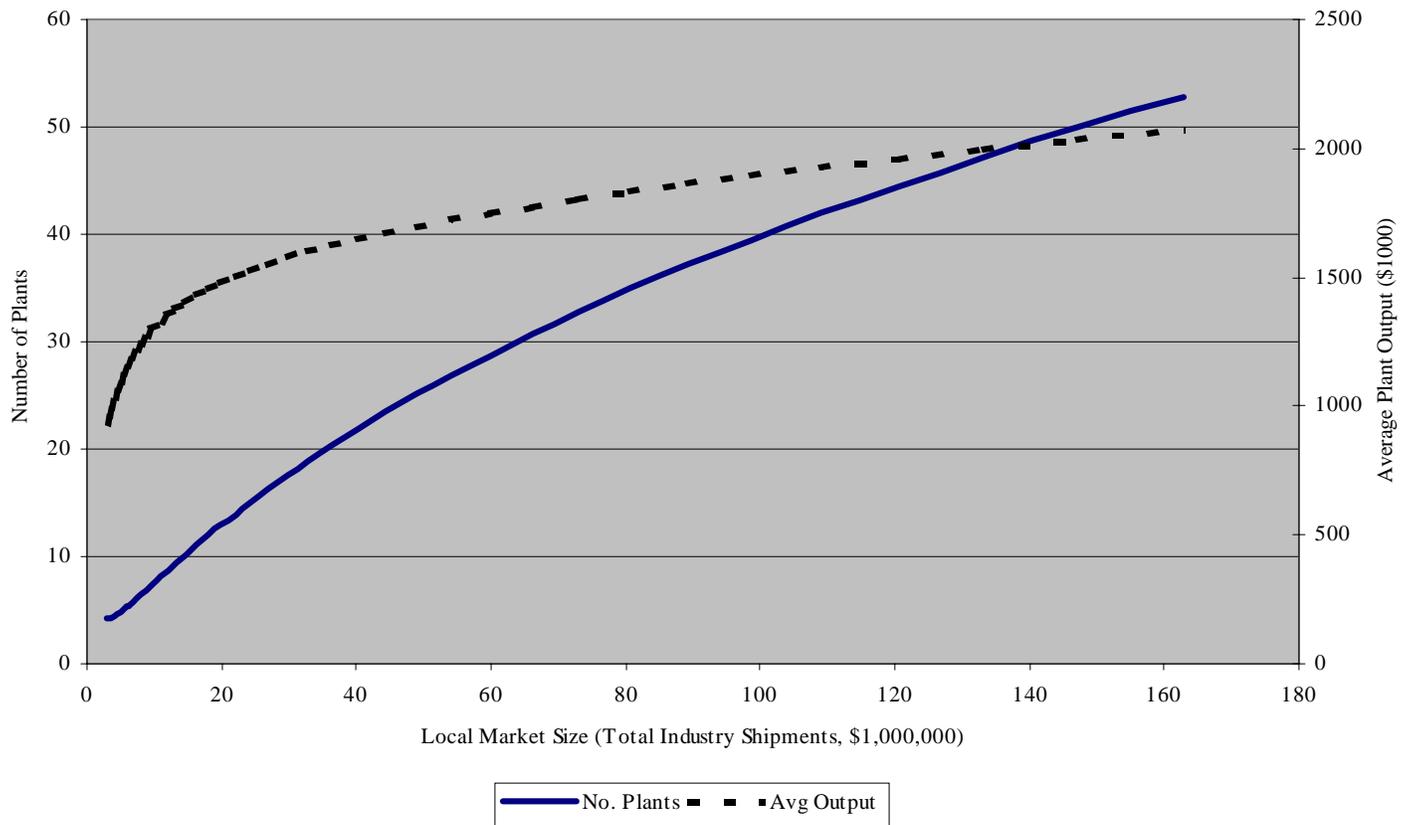


Figure 3: Cutoff Productivity Levels for Various Functional Forms of  $\tau(M/A)$



**Figure 4. Estimated Local Market Effects: Number of Plants and Average Plant Output**  
*(From Estimated Cubic Polynomials of Number of Plants and Logged Average Output in Local Market Size)*



**Table 1: Relevance of Local Demand to Plant-Level Investment**

Instrument Set	N	No Plant Effects		Plant Effects	
		F	R <sup>2</sup>	F	R <sup>2</sup>
County	8621	18.64	0.034	59.00	0.055
CEA	8695	28.28	0.039	100.4	0.077
EA	8695	31.18	0.041	98.93	0.076

**Table 2: Production Function Estimation Results**

All sample plants. Heteroskedasticity-robust standard errors in parentheses.

Instrument Set	N	1 <sup>st</sup> Stage		2 <sup>nd</sup> Stage Estimates					
		F	R <sup>2</sup>	Const	D82	D87	DMULT	LCOMP	R <sup>2</sup>
County	11,552	76.45	0.050	1.871	-0.104	-0.036	0.053	0.977	0.467
				(0.072)	(0.006)	(0.007)	(0.005)	(0.013)	
CEA	11,652	63.74	0.044	2.024	-0.105	-0.029	0.052	0.950	0.427
				(0.080)	(0.006)	(0.007)	(0.005)	(0.014)	
EA	11,652	77.99	0.050	1.997	-0.105	-0.030	0.052	0.956	0.456
				(0.072)	(0.006)	(0.007)	(0.005)	(0.014)	

**Table 3: Descriptive Statistics—All Region-Year Observations (N=1036)**

Variable	Mean	Std. Dev.	Skewness	IQ Range	90-10%ile
Prod. Dispersion (IQ Range)	0.265	0.177	2.176	0.174	0.397
Median Productivity	-0.014	0.133	0.443	0.117	0.289
Qty.-Weighted Avg. Prod.	0.043	0.188	-6.927	0.174	0.344
ln(Demand Density)	1.297	1.324	0.032	1.544	3.416
TFP (N=11652)	0.000	0.266	-0.709	0.262	0.561

**Table 4. Local Productivity Distribution Regressions—Main Results**

Location-year observations with at least 5 establishments, N=859. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.019	0.049	0.027	0.051
	Demand Density Coefficient	-0.015* (0.004)	-0.021* (0.007)	-0.015* (0.004)	-0.022* (0.007)
Median Productivity	R <sup>2</sup>	0.114	0.159	0.120	0.176
	Demand Density Coefficient	0.029* (0.003)	0.025* (0.005)	0.030* (0.003)	0.024* (0.005)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.070	0.113	0.074	0.121
	Demand Density Coefficient	0.036* (0.005)	0.023* (0.007)	0.036* (0.005)	0.021* (0.006)

**Table 5. Local Productivity Distribution Regressions—Alternative Minimum Required Observations**

Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

A. Observations with at Least 2 Establishments, N=1026

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.014	0.043	0.029	0.047
	Demand Density Coefficient	-0.016* (0.005)	-0.013 (0.009)	-0.016* (0.005)	-0.014 (0.009)
Median Productivity	R <sup>2</sup>	0.076	0.125	0.083	0.138
	Demand Density Coefficient	0.028* (0.003)	0.024* (0.006)	0.028* (0.003)	0.024* (0.006)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.064	0.109	0.070	0.118
	Demand Density Coefficient	0.036* (0.005)	0.028* (0.007)	0.037* (0.004)	0.028* (0.007)

B. Observations with at Least 10 Establishments, N=554

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.021	0.086	0.044	0.089
	Demand Density Coefficient	-0.014* (0.004)	-0.021* (0.008)	-0.014* (0.005)	-0.023* (0.008)
Median Productivity	R <sup>2</sup>	0.143	0.223	0.146	0.230
	Demand Density Coefficient	0.029* (0.003)	0.007 (0.005)	0.029* (0.003)	0.007 (0.005)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.106	0.205	0.122	0.230
	Demand Density Coefficient	0.030* (0.004)	0.008 (0.008)	0.031* (0.004)	0.008 (0.007)

**Table 6. Local Productivity Distribution Regressions—Physical Output Measure**

## A. Production Function Estimation Results

N	Const	D82	D87	DMULT	LCOMP	R <sup>2</sup>
11,114	4.843 (0.115)	-0.157 (0.009)	-0.144 (0.010)	0.095 (0.007)	0.995 (0.020)	0.263

## B. Local Productivity Distribution Regressions

Location-year observations with at least 5 establishments, N=859. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.034	0.083	0.050	0.095
	Demand Density Coefficient	-0.031* (0.006)	-0.030* (0.013)	-0.032* (0.006)	-0.028* (0.013)
Median Productivity	R <sup>2</sup>	0.044	0.071	0.048	0.075
	Demand Density Coefficient	0.032* (0.004)	0.018* (0.007)	0.032* (0.005)	0.019* (0.007)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.044	0.062	0.046	0.063
	Demand Density Coefficient	0.046* (0.007)	0.029* (0.011)	0.047* (0.007)	0.028* (0.011)

**Table 7. Local Productivity Distribution Regressions—Calculated Productivity Measure**

Location-year observations with at least 5 establishments, N=859. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.023	0.061	0.033	0.064
	Demand Density Coefficient	-0.017* (0.004)	-0.024* (0.008)	-0.017* (0.004)	-0.025* (0.008)
Median Productivity	R <sup>2</sup>	0.084	0.130	0.185	0.228
	Demand Density Coefficient	0.026* (0.003)	0.027* (0.006)	0.024* (0.003)	0.022* (0.005)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.036	0.063	0.064	0.093
	Demand Density Coefficient	0.029* (0.005)	0.036* (0.009)	0.028* (0.004)	0.030* (0.009)

**Table 8. Local Productivity Distribution Regressions—Productivity Estimates from Fixed Effects Specification**

A. Production Function Estimation Results (CEA Instrument Set)

Output Measure	N	1 <sup>st</sup> Stage		2 <sup>nd</sup> Stage Estimates					R <sup>2</sup>
		F	R <sup>2</sup>	Const	D82	D87	DMULT	LCOMP	
Deflated Output	11,652	185.8	0.112	0.020 (0.004)	-0.037 (0.005)	-0.024 (0.005)	0.012 (0.004)	1.093 (0.018)	0.340
Physical Product	11,114	n/a	n/a	0.054 (0.005)	-0.073 (0.006)	-0.087 (0.007)	0.008 (0.005)	0.992 (0.023)	0.201

B1. Local Productivity Distribution Regressions—Deflated Revenue Output Measure  
Location-year observations with at least 5 establishments, N=859. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.020	0.050	0.034	0.060
	Demand Density Coefficient	-0.017* (0.005)	-0.028* (0.008)	-0.018* (0.005)	-0.030* (0.008)
Median Productivity	R <sup>2</sup>	0.005	0.052	0.045	0.108
	Demand Density Coefficient	0.006* (0.003)	0.009 (0.005)	0.005 (0.003)	0.004 (0.005)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.001	0.031	0.018	0.054
	Demand Density Coefficient	0.004 (0.004)	0.017* (0.008)	0.004 (0.004)	0.012 (0.008)

B2. Physical Output Measure

Dependent Variable	Demand Controls: Year Dummies:	No	Yes	No	Yes
		No	No	Yes	Yes
Productivity Dispersion	R <sup>2</sup>	0.034	0.083	0.051	0.097
	Demand Density Coefficient	-0.030* (0.006)	-0.031* (0.013)	-0.032* (0.006)	-0.030* (0.013)
Median Productivity	R <sup>2</sup>	0.065	0.093	0.079	0.108
	Demand Density Coefficient	0.039* (0.004)	0.032* (0.007)	0.038* (0.005)	0.028* (0.007)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.021	0.050	0.043	0.065
	Demand Density Coefficient	0.059* (0.012)	0.072* (0.025)	0.062* (0.012)	0.065* (0.025)

**Table 9. Local Productivity Distribution Regressions—Technology/Input Controls Added**

Location-year observations with at least 5 establishments, N=859. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Dependent Variable	Controls:	Demand Only	Demand & Tech
Productivity Dispersion	R <sup>2</sup>	0.051	0.063
	Demand Density Coefficient	-0.022* (0.007)	-0.016* (0.007)
Median Productivity	R <sup>2</sup>	0.176	0.231
	Demand Density Coefficient	0.024* (0.005)	0.032* (0.005)
Q-Wt. Avg. Productivity	R <sup>2</sup>	0.121	0.161
	Demand Density Coefficient	0.021* (0.006)	0.029* (0.006)