

Product Substitutability and Productivity Dispersion

Chad Syverson*

University of Chicago and NBER

July 2003

* I am grateful to Gary Becker, Mark Duggan, John Haltiwanger, Ali Hortaçsu, Casey Mulligan, John Shea, Plutarchos Sakellaris, two anonymous referees, and seminar participants at the University of Chicago and University of Maryland for their comments. 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 and Chicago Census Research Data Center. 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. This paper has been screened to ensure that no confidential data are revealed. The author can be reached at: Department of Economics, University of Chicago; 1126 E. 59th St.; Chicago, IL 60637; e-mail: syverson@uchicago.edu

Abstract

Tremendous differences in producer productivity levels exist, even within narrowly defined industries. This paper explores the influence of product substitutability in an industry on this disparity. When consumers can easily switch between producers, relatively inefficient (high-cost) producers cannot profitably operate. Thus high-substitutability industries should exhibit less productivity dispersion and have higher average productivity levels. I demonstrate this mechanism in a simple industry equilibrium model and test it empirically using producer-level data from 443 U.S. manufacturing industries. I find evidence that substitutability—measured in several ways—is indeed negatively related to within-industry productivity dispersion and positively related to median productivity levels.

I. Introduction

Empirical explorations into the productivity levels of individual producers have consistently found large heterogeneity across plants. Perhaps surprisingly, a great amount of productivity variation between plants is observed within what may seem to be narrowly defined (e.g., four-digit SIC) industries. Table 1 shows statistics demonstrating this dispersion. Using plant-level data from the 1977 Census of Manufactures, I compute productivity distribution moments for four-digit manufacturing industries for each of four different productivity measures.¹ As can be seen in the first numerical column, the average *within-industry* interquartile range of logged plant-level labor productivity values is roughly 0.66. This corresponds to a nearly 2-to-1 ratio in value added per labor unit (employee or employee-hour) between the 75th and 25th-percentile plants in an industry's productivity distribution. Bear in mind that these differences are observed when restricting attention to the middle half of the distribution; including more of the tails amplifies intra-industry heterogeneity. The average 90-10 and 95-5 percentile productivity ratios within industries are over 4-to-1 and 7-to-1, respectively. Factor intensity variations are not solely responsible for these large differences, either. Intra-industry total factor productivity differences, while smaller, are still sizeable. The values in the bottom half of Table 1 indicate average interquartile total factor productivity (TFP) ratios between 1.34-to-1 and 1.56-to-1, depending on the measure. It is important to note that the heterogeneity observed here is a persistent phenomenon. Empirical studies using other (but perhaps less comprehensive) cross sections have found similar within-industry productivity differences.

A host of theoretical work has arisen in an attempt to explain the sources of this dispersion. The great majority of this research focuses on supply-side/production explanations, such as technology shocks, management skill, R & D, or investment patterns.² While these proposed explanations are undoubtedly important, I contend that demand-side (output market) conditions can also play an important role in explaining persistent productivity dispersion. I focus in this paper on the influence of one demand characteristic—product substitutability—on the equilibrium plant-level productivity distribution within an industry.

¹ Details regarding the data and methods used to construct plant-level productivity values will be discussed below.

² Bartelsman and Doms (2000) review much of this literature.

An obvious question arising from the above facts regards how such wide productivity dispersion can exist in equilibrium. One might expect a long-run tendency for industry output to be reallocated to more productive plants. They can produce output at lower cost than industry rivals and grab additional market share by undercutting their opponents' prices without sacrificing profits. If this process were to continue unabated, industry equilibrium would expectedly be characterized by a degenerate plant-level productivity distribution within the industry; all operating plants would share the same (highest possible) productivity level.

The above evidence suggests something impedes this reallocation process, at least partially. Imperfect product substitutability seems a likely candidate. It prevents industry customers from costlessly (in either a budgetary or utility sense) shifting purchases between industry producers. Thus more efficient (lower cost) plants cannot lure away all demand from their less efficient industry rivals simply with lower prices, and lower-productivity establishments are able to stay in business despite their cost disadvantage. As a result, the equilibrium productivity (cost) dispersion in an industry should be related to the extent of product substitutability. Industries with very segmented (in either geographic or product space) output markets can support large productivity differences, even in a long-run equilibrium. Conversely, high-substitutability industries should exhibit little dispersion. Further, since the productivity truncation only affects the low end of the distribution, greater substitutability implies higher central tendency in an industry's productivity distribution.³

Product substitution barriers are manifold. Transport costs prevent costless switching among suppliers even when industry products are otherwise identical. In the manufactured ice industry (SIC 2079), for example, it is unlikely that the physical characteristics of output vary much from plant to plant. However, the obvious transport barriers make manufactured ice in one

³ This raises questions about why even those producers protected from competition by imperfect substitutability would not seek to maximize efficiency. An implicit (and I believe reasonable) assumption underlying the intuitive premise of the paper is that improving productivity is not costless. The model below, as several models used in other contexts do, makes this stark by assuming that this cost is infinite; a producer's productivity draw is permanent and unchangeable. The assumption is not likely to be key to the results, however. One could introduce a costly productivity-improving technology and still obtain the same qualitative implications. In many theoretical frameworks, producers would have greater incentive to undertake productivity-enhancing investment when high product substitutability exposes them to intense competition.

locale an imperfect substitute for the same product in another. High-productivity plants would be unable to take market share from their less efficient industry cohorts given sufficient distance between them, supporting a range of productivity levels in equilibrium.

Physical product differentiation also limits substitutability. Idiosyncratic consumer preferences across attributes allow some producers to remain viable even if they are less physically efficient than their industry counterparts. Plants producing niche-market specialty products may often have higher per-unit costs than industry competitors who focus on mass production. However, niche producers can survive (and indeed thrive) if their product characteristics appeal to certain purchasers.

Branding and advertising can also lead to consumers perceiving physically identical products as being less than perfectly interchangeable. The classic example of name-brand bleach fetching a higher price than chemically identical generic alternatives is illustrative of this. Sufficient brand identity would allow a producer to operate even in the face an efficiency gap between itself and its industry competitors.

Real or perceived differences in services bundled with products, such as delivery speed, documentation, and product support, can also contribute to imperfect output substitutability. Finally, an array of intangible factors such as specific history-laden relationships between producers and their customers, interpersonal customer-manager interaction, and other assets of goodwill make costless substitution of another manufacturer's output impossible.

Alone or in combination, these factors allow productivity differences to persist among industry producers. Expectedly, as these substitutability factors vary across industries, certain moments of the productivity distribution should fluctuate in concert. Table 1 summarizes the substantial across-industry variance in plant-level productivity distribution moments. The between-industry standard deviation of within-industry interquartile productivity ranges (that is, the dispersion of productivity dispersion) is roughly one-third of the mean within-industry interquartile range. Similarly, the standard deviations of the wider intra-industry productivity ranges are roughly one-fourth of their means. Across-industry differences in plant productivity distributions are not restricted to second moments; within-industry median TFP levels have an across-industry coefficient of variation of 0.20.

The objective of this paper is to test if product substitutability differences are linked with the observed variation in these moments. Specifically, I test the notions forwarded above: that

greater product substitutability should be correlated with less productivity dispersion and higher central tendency in industries' plant-level productivity distributions.

This paper is a broadly-focused complement to Syverson (2003). That study explores the effect of exogenous differences in output substitutability *within* an industry (ready-mixed concrete) on the dispersion and central tendency of productivity distributions in local concrete markets. The results therein suggest that increased geographic barriers to substitution (lower demand and plant density in the concrete industry's case) lead to greater productivity dispersion and lower average productivity levels in the market. By exploiting substitutability variation within a given industry, that study holds constant many possible confounding influences on local productivity distributions. Its limited scope, however, makes generalizing the link between substitutability and productivity distribution moments (and thus the utility of drawing implications for aggregate production behavior) slightly more tenuous. The goal of the present paper is to test if this link holds more broadly across the economy.⁴

Differences in industry plant-level productivity distributions and their causes are of obvious interest to those interested in competition and production within industries. Moreover, as Basu and Fernald (1997) as well as others have pointed out, reallocation of production across industries with differing technological characteristics can be an important source of changes in aggregate production data. Therefore addressing issues regarding these across-industry differences could yield insights about aggregate productivity movements as well.

I show in the next section the link between the shape of the productivity distribution and product substitutability (as well as some technological parameters) in a simple model of entry and competition within an industry. I go on to test the predictions of the model and its extensions by collecting measures of product substitutability within industries and comparing variations in these factors to moments of industries' plant-level productivity distributions. These productivity moments are computed from data from roughly 200,000 establishments from the U.S. Census of Manufactures. To preview the results, I find that within-industry productivity dispersion (central tendency) does indeed tend to decrease (increase) when substitutability is high. These results hold even after controlling for several other plausible causes of

⁴ In a recent paper using data from 2300 firms in East Asia, Hallward-Driemeier, Iarossi, and Sokoloff (2002) find some evidence of a broad-based link between productivity dispersion and the level of market "integration," which

heterogeneity and appear robust to empirical modeling specifications. Furthermore, proxies for characteristics of industry technologies (fixed operating costs and sunk entry costs specifically) are also related to industries' plant-level productivity distribution moments in directions predicted by theory.

II. Theoretical Motivation

I formalize the above intuition using a theoretical framework where heterogeneous-productivity producers compete in an industry product market with (possibly imperfect) substitution across producers' outputs, which are varieties of the industry product. The model allows the equilibrium plant-level cost/productivity distribution to endogenously respond to variations in substitutability. Because I am concerned here with differences in productivity distributions across industries rather than intertemporal fluctuations within them, industry dynamics are not a primary concern. The equilibrium is a two-stage entry/production decision meant to model long-run differences in outcomes. For the sake of expositional clarity and to permit maximum transparency of the selection-driven mechanism, I assume a specific demand system. It is important to note, however, that similar qualitative implications can be obtained from other demand systems. While simple, the model shows in a fairly straightforward manner how differences in product demand and technology structures can create variation in industry productivity distribution moments.

A. Model

An industry is comprised of a continuum of producers of measure N . Each producer (indexed by i , where I is the set of industry producers) makes a distinct variety of the industry product. The representative industry consumer has preferences over these varieties given by

$$\begin{aligned}
 U &= y + \alpha \int_{i \in I} q_i di - \frac{1}{2} \eta \left(\int_{i \in I} q_i di \right)^2 - \frac{1}{2} \gamma \int_{i \in I} q_i^2 di \\
 &= y + \alpha \int_{i \in I} q_i di - \frac{1}{2} \left(\eta + \frac{\gamma}{N} \right) \left(\int_{i \in I} q_i di \right)^2 - \frac{1}{2} \gamma \int_{i \in I} (q_i - \bar{q})^2 di
 \end{aligned} \tag{1}$$

where y is the quantity of a numeraire good, q_i is the quantity of good i consumed, and

they define as consisting in part of factors such as transport costs and product differentiation.

$$\bar{q} \equiv \frac{1}{N} \int_{i \in I} q_i di, \quad \alpha > 0; \quad \eta > 0; \quad \gamma \geq 0.$$

Utility is a quadratic function in total consumption of the industry's output, minus a term that increases in the variance of consumption levels across varieties. This introduces an incentive to equate consumption levels of different products. The parameter γ embodies substitutability across varieties; an increase in γ imposes a greater utility loss from consuming idiosyncratically large or small quantities of particular q_i , therefore limiting consumer response to price differences across industry producers. As $\gamma \rightarrow 0$, substitutability becomes perfect: only the total quantity of industry varieties consumed—not its composition—affects utility. The parameters α and η shift demand for the industry's output relative to the numeraire. This is the utility function specified by Melitz and Ottaviano (2003) in their theoretical study of market size effects in trade. It is useful for my purposes because it embeds imperfect product substitutability in a parsimonious and tractable way.

Utility maximization by consumers implies that producers face the following demand function:

$$q_i = \frac{\alpha}{\eta N + \gamma} - \frac{1}{\gamma} p_i + \frac{\eta N}{\eta N + \gamma} \frac{1}{\gamma} \bar{p}, \quad (2)$$

where p_i is the price of good i and \bar{p} the average price among industry producers. Note that demand falls to zero at

$$p_{\max} = \frac{\gamma}{\eta N + \gamma} \alpha + \frac{\eta N}{\eta N + \gamma} \bar{p}, \quad (3)$$

so no producer will price above this level.

Industry producers operate at a constant marginal cost c_i that varies across producers, and productivity is defined to be some decreasing function of this cost. Thus productivity levels are idiosyncratic to industry producers. Producers must also pay a common fixed operating cost f if $q_i > 0$. Profits are therefore given by

$$\pi_i = \left(\frac{\alpha}{\eta N + \gamma} - \frac{1}{\gamma} p_i + \frac{\eta N}{\eta N + \gamma} \frac{1}{\gamma} \bar{p} \right) (p_i - c_i) - f. \quad (4)$$

Bertrand-Nash profit maximization implies the producer's optimal price (subject to $p_i \leq p_{\max}$) is

$$p_i = \frac{1}{2} \frac{\alpha\gamma}{\eta N + \gamma} + \frac{1}{2} \frac{\eta N}{\eta N + \gamma} \bar{p} + \frac{1}{2} c_i. \quad (5)$$

(The individual producer is too small relative to the industry to take into account the effect of its own pricing decision on the industry average.) Not surprisingly, the optimal price is increasing in the overall demand level (indexed by α), the average price charged by industry competitors, and the producer's cost level. Combining (5) with (2) gives the producer's quantity sold at the optimal price:

$$q_i = \frac{1}{2} \frac{\alpha}{\eta N + \gamma} + \frac{1}{2\gamma} \frac{\eta N}{\eta N + \gamma} \bar{p} - \frac{1}{2\gamma} c_i. \quad (6)$$

Then, using (5) and (6), maximized profits are

$$\pi_i = \frac{1}{4\gamma} \left(\frac{\alpha\gamma}{\eta N + \gamma} + \frac{\eta N}{\eta N + \gamma} \bar{p} - c_i \right)^2 - f. \quad (7)$$

These expressions imply a cost draw c^* such that operations are not profitable if $c_i > c^*$.⁵

Setting (7) equal to zero and solving for c^* gives this level explicitly:

$$c^* = \frac{\alpha\gamma}{\eta N + \gamma} + \frac{\eta N}{\eta N + \gamma} \bar{p} - 2\sqrt{\gamma f}. \quad (8)$$

Substituting this back into (7) yields maximized operating profits in terms of the cutoff cost level c^* and own costs:

$$\pi_i = \frac{1}{4\gamma} \left(2\sqrt{\gamma f} + c^* - c_i \right)^2 - f. \quad (9)$$

A large pool of ex-ante identical potential entrants decide whether enter the industry as follows. They first decide whether to pay a sunk entry cost s in order to receive a cost/productivity draw c_i from a known distribution with positive support and probability density function $g(c)$. If they pay s , they observe c_i and decide whether to begin production and earn the corresponding operating profits (9). Clearly, only those obtaining marginal cost draws yielding nonnegative operating profits (i.e., $c_i \leq c^*$) choose to produce in equilibrium. Others produce nothing, earn zero operating profits, and lose their sunk cost. Hence the expected gain from paying s is the expectation of (9) over $g(c)$, conditional upon drawing $c_i \leq c^*$. This expected

⁵ Note that due to the quadratic form of the profit function, while (7) implies positive profits for some $c_i > c^*$, such cost levels also imply that $p_i > p_{max}$ and therefore $q_i < 0$, which is impossible.

gain is obviously affected by the cutoff cost level c^* . Free entry pins down this value: c^* must set the net expected value of entry into the industry V^e equal to zero. Thus c^* satisfies

$$V^e = \int_0^{c^*} \left[\frac{1}{4\gamma} (2\sqrt{\gamma f} + c^* - c)^2 - f \right] g(c) dc - s = 0. \quad (10)$$

This expression summarizes the industry equilibrium. It combines the conditions that all producers make nonnegative profits from operations (net of fixed operating costs), and that entry occurs until the net expected value of taking a cost draw is zero.⁶

B. Comparative Statics

The primary comparative static that I seek to test empirically, the effect of within-industry product substitutability on the cost/productivity distribution, is presented here. I also derive secondary implications of the model regarding the effects of the technology parameters f and s . I control for these technological effects in some of the empirical specifications below.

When the model's parameters change, c^* adjusts to maintain equilibrium. Shifts in exogenous variables therefore affect the truncation point of the equilibrium cost/productivity distribution. So while the distribution of possible cost draws $g(c)$ is exogenous, the distribution among equilibrium producers—the truncation $g(c)/G(c^*)$, where $G(c^*)$ is the value of the cost cumulative distribution function at c^* —is endogenous and determined by the cutoff cost level. I test for this truncation using the moments of industries' plant-level cost/productivity distributions. Higher c^* results in higher (lower) average cost (productivity) levels in an industry and lower variation among producers' cost/productivity levels.⁷

Product Substitutability. From the implicit function theorem:

⁶ The equilibrium mass of producers N is determined by α , η , γ , f , \bar{c} , and c^* , and can be solved for by substituting the \bar{p} implied by (5) into (8).

⁷ The implication regarding within-industry productivity dispersion implicitly assumes some regularity conditions on $g(c)$, since for some distributions (even those with some arbitrary upper-bound cost level), it is possible that further truncation from below would actually increase dispersion moments rather than decrease them. However, for most common distributions, truncation implies a reduction in dispersion.

$$\frac{dc^*}{d\gamma} = \frac{-\partial V^e / \partial \gamma}{\partial V^e / \partial c^*}, \quad (11)$$

where from (10),

$$\frac{\partial V^e}{\partial \gamma} = \int_0^{c^*} \left[-\frac{1}{4\gamma^2} (2\sqrt{f} + c^* - c)^2 + \frac{1}{2\gamma} (2\sqrt{f} + c^* - c) \sqrt{\frac{f}{\gamma}} \right] g(c) dc. \quad (12a)$$

Some algebra shows:

$$\frac{\partial V^e}{\partial \gamma} = \int_0^{c^*} \left[-\frac{1}{4\gamma^2} (c^* - c)^2 - \frac{1}{2} \sqrt{\frac{f}{\gamma^3}} (c^* - c) \right] g(c) dc < 0, \quad (12b)$$

which is clearly negative since the terms in the brackets are less than zero throughout the region of integration. Further,

$$\frac{\partial V^e}{\partial c^*} = \left[\frac{1}{4\gamma} (2\sqrt{f} + c^* - c^*)^2 - f \right] g(c^*) + \int_0^{c^*} \left[\frac{1}{2\gamma} (2\sqrt{f} + c^* - c) \right] g(c) dc. \quad (13a)$$

The first term in this expression is equal to zero. (Intuitively, the marginal increase in V^e from letting in formally marginally unprofitable producer is zero.) Simplifying gives

$$\frac{\partial V^e}{\partial c^*} = \int_0^{c^*} \frac{1}{2\gamma} (2\sqrt{f} + c^* - c) g(c) dc > 0. \quad (13b)$$

(12b) and (13b) imply $dc^*/d\gamma > 0$; a decrease in substitutability (embodied in an increase in γ) leads to a higher cutoff cost level. This result is in accordance with the intuition in the introduction. When substitutability is low, relatively inefficient producers are protected from intense competition from their lower-cost competitors and can operate profitably in equilibrium. Given the inverse relationship between costs and productivity, this implies we should expect greater dispersion and lower central tendency in the productivity distribution in low-substitutability industries.

Fixed Operating Costs. Taking the derivative of (10) with respect to the fixed production cost and simplifying the result yields

$$\frac{\partial V^e}{\partial f} = \int_0^{c^*} \frac{1}{2\sqrt{f}} (c^* - c) g(c) dc > 0. \quad (14)$$

From above, $\partial V^e / \partial c^* > 0$, so the implicit function theorem implies $dc^*/df < 0$. Higher fixed

production costs lower the equilibrium cutoff cost level. High fixed costs, as with high product substitutability, make it more difficult for relatively inefficient producers to be profitable. Thus all else equal, high f industries should exhibit less dispersion and higher central tendency in their productivity distribution.

Sunk Entry Costs. The derivative of V^e with respect to the sunk entry cost s is -1 . This, combined with the results above, implies $dc^*/ds > 0$. Thus high sunk entry costs make it easier for relatively inefficient producers to survive in equilibrium. (Note that this is the opposite effect of high fixed operating costs.) To see this intuitively, suppose the number of equilibrium producers supported by the market size were fixed at some number n , and imagine the sunk cost approaching zero. With very low entry costs, it is extremely cheap for potential entrants to buy cost draws, so a large number end up doing so. The n lowest order statistics of these cost draws (i.e., those potential entrants that will produce in equilibrium) decrease when sunk costs fall. As a result, c^* falls with s —the cutoff cost level and sunk entry costs move in the same direction.⁸

C. Extending the Basic Model

Adding Transport Costs. One of the empirical product substitutability measures I use below is a proxy for average transport costs in an industry. Its purpose is to capture spatial differentiation differences across industries. While transport costs are not explicitly included in the model above, one could interpret the derivations with respect to the substitutability parameter γ as a reduced-form embodiment of spatial substitutability. This would imply that higher transport costs (lower spatial substitutability) increase the level of c^* . However, the same implication can be derived when the model is augmented to include a transport cost parameter directly. I do so in the appendix for interested readers.⁹

⁸ The implication that fixed operating and sunk entry costs move the cutoff cost/productivity level in opposite directions is a common feature of setups similar to the present model (see Asplund and Nocke (2003) and Melitz (2003), for example).

⁹ It can also be shown that c^* responds in the augmented model to changes in fixed operating and sunk entry costs in the same directions as above.

International Trade. Melitz (2003) examines theoretically the influence of international trade on the cutoff and average industry productivity levels when producers have heterogeneous productivity levels. He finds that increased trade exposure—either moving from autarky to trade or lowering trade barriers within a regime where trade already exists—drives low productivity domestic plants out of business and increases market shares of high-productivity domestic operators. These effects combine to raise the industry’s cutoff productivity level and average productivity.

The present model could be similarly modified to allow for the possibility of trade by incorporating several markets and additional costs incurred if a producer chooses to export. Given the common structures of the above model and the Melitz framework, equivalent effects of an industry’s foreign trade exposure are implied here. I attempt to control for trade exposure differences in some empirical specifications.

Entry and Exit. The two-stage game of entry and production above abstracts from continuous producer turnover in a dynamic setting. However, dynamics could be added to the model using a framework of the sort found in Hopenhayn (1992), Asplund and Nocke (2003), and Melitz (2003). These models share the common characteristic of allowing producer-level uncertainty while preserving deterministic industry aggregates, including the productivity distribution.

Both Asplund and Nocke (2003) and Melitz (2003) specify a productivity evolution process that induces a segment of industry producers to exit each period upon receipt of a sufficiently bad productivity innovation. Both papers assume that a fraction of producers ρ retain the same cost/productivity level from one period to the next.¹⁰ Since industry aggregates are constant, this share of producers remains in the industry. Melitz supposes the remaining fraction of producers receive a “killer” shock which forces exit. In Asplund and Nocke, the $1-\rho$ share of producers receive new productivity draws from a common ex-ante distribution (the same one from which entrants receive their draws). Thus, some producers receive new draws poor enough to require closure and liquidation, while others are able to remain in the industry—and may even be more productive than they were previously.

Both papers have the same implication regarding the effect of idiosyncratic productivity

¹⁰ To facilitate discussion, I have changed the notation from the original papers.

dynamics on the industry productivity distribution. As the persistence of the productivity process increases (i.e., ρ gets bigger), the cutoff productivity level also climbs, decreasing productivity dispersion and increasing the average productivity level. Intuitively, greater persistence implies a larger stream of discounted expected future profits for successful entrants because of the lower probability of receiving a negative productivity shock. Free entry requires that this be balanced by a lower probability of successful entry; i.e., a higher productivity cutoff value. I include a proxy for productivity persistence in one of the empirical tests below.

III. Empirical Method and Data

Testing the product substitutability–productivity distribution link implies a general empirical specification of the form

$$y_I = \beta_0 + X_{sI}B_s + X_{cI}B_c + \varepsilon_I.$$

Here, the plant-level productivity distribution moment y_I (a dispersion or central tendency measure) for industry I is a function of a constant, a vector of substitutability measures X_{sI} , a vector X_{cI} of other influences on the moments, and an industry-specific error term. I discuss the components of these vectors below.

The productivity distribution moments in this study are computed for 443 four-digit industries using plant-level data from the 1977 Census of Manufactures (CM).¹¹ The CM contains production data for every manufacturing establishment in the U.S., totaling roughly 300,000 plants. To lighten reporting burdens, particularly small plants (typically those with less than five employees—about one-third of plants), are classified as Administrative Record (AR) cases. Since all input data for these plants except the number of employees and total payroll is imputed, my sample includes only the roughly 200,000 non-AR plants in order to minimize productivity mismeasurement.

I compute industry productivity moments with four different productivity measures, estimating the model with moments from each of the corresponding distributions as a robustness check. Two are labor productivity measures: value added per employee, and value added per

¹¹ I am restricted to the 1977 cross section because the highly detailed Commodity Transportation Survey data used to measure geographic substitutability across industries is not available for later years. The CTS is, in effect, the binding data constraint in this study.

employee-hour. Value added is calculated as the difference between a plant's reported value of shipments and its expenditures on materials, parts, and energy. I use value added as an output measure because interplant differences in intermediate input intensity (primarily in materials expenditures) cause gross-output productivity measures to be quite noisy for some industries. Plant employee-hours are computed as reported production-worker hours plus nonproduction worker hours imputed according to the method of Davis and Haltiwanger (1991).¹²

In addition to these labor productivity measures, two total factor productivity (TFP) values are computed for each establishment. Both follow the typical form

$$tfp_i = y_i - \alpha_l l_i - \alpha_k k_i - \alpha_m m_i - \alpha_e e_i,$$

where the lower-case letters indicate logarithms of establishment-level TFP, gross output, labor hours, capital stock, materials, and energy inputs. The two TFP measures differ by the manner in which the factor elasticities $\alpha_j, j = \{l, k, m, e\}$ are computed. TFP1 uses input cost shares of the individual plants, while TFP2 uses the average cost shares across all industry plants for each plant in the corresponding four-digit industry. The plant-specific elasticities in TFP1 better account for within-industry technology differences manifested in input intensity differences, but are potentially vulnerable to measurement error because of the noisy nature of establishment-level data. Using the average input elasticities for all industry plants in TFP2 trades flexibility with regard to intraindustry technology differences for a reduction in spurious productivity dispersion due to measurement error. Reported wage bills, materials costs, and energy expenditures from the CM are used to compute input cost shares. Capital expenditures are computed by multiplying reported plant equipment and building stocks by their respective capital rental rates for each plant's corresponding two-digit industry.¹³

I measure industry productivity dispersion as the interquartile productivity difference divided by the industry's median productivity level.¹⁴ The dispersion measure is standardized to

¹² The plant's number of nonproduction workers is multiplied by the average annual hours for nonproduction workers in the corresponding two-digit industry (calculated from Current Population Survey data).

¹³ Capital rental rates are from unpublished data constructed and used by the Bureau of Labor Statistics for use in computing their Multifactor Productivity series. Formulas, related methodology, and data sources are described in U.S. Bureau of Labor Statistics (1983) and Harper, Berndt, and Wood (1989).

¹⁴ I check the results for robustness to other interquartile differences below.

prevent pure scale differences between industries—primarily a factor for labor productivity measures due to capital intensity variations—from causing productivity variation that is neither within the confines of the model nor very relevant to the paper’s hypothesis. Ordinal moments are used rather than the coefficient of variation because moments from plant-level data are especially vulnerable to influence of outliers.

For regressions with a central tendency measure as the dependent variable, I use median total factor productivity in the industry (both TFP1 and TFP2). Labor productivity levels are not included in the central tendency regressions because wide capital intensity variation yields average labor productivity differences between industries that are outside of the theoretical framework. (It is obviously not possible to remove scale effects in central tendency moments.) TFP is much less susceptible to such scale problems. Median TFP levels are used rather than the averages to counteract outlier effects.

A. Product Substitutability Factors

Ideally, one would regress these productivity moments on an average substitutability parameter implied by the industry’s demand system. It is unfortunately impossible to estimate each of these values given the number of industries and products. My strategy is to instead use a vector X_{st} of measurable proxies for substitution elasticities among the outputs of industry producers. To motivate my choices for variables included in X_{st} , I return to the earlier discussion on sources of substitutability barriers.

Geographic barriers to substitution arise when transport costs hinder producers from practicably selling their output beyond certain shipment distances. These distances, of course, depend on the magnitude of the transport costs. I compute two measures of transport costs for an industry; both use data from the 1977 Commodity Transport Survey (CTS). This CTS contains an enormous amount of information on manufacturers’ shipments at a detailed product level (most by five-digit product class). Included in this unusually rich survey are, for each product class, the average dollar value per pound of shipments and a decomposition of total tons shipped by distance category. This information is used to construct the two transport cost measures.

The first metric, VALUELB, is the natural logarithm of the weighted sum of the dollar-value-to-weight ratios of all product classes in a given four-digit industry. The weights are the

product classes' shares of total industry product tonnage shipped.¹⁵ There is an obvious relationship between the value of shipments per pound and product transportability. Goods valuable in relation to their weight are more economical to ship. Industries with high values of VALUELB expectedly have less geographically segmented output markets and greater product substitutability. (Summary statistics of the across-industry distribution of VALUELB and the other variables that will be discussed below are presented in Table 2.)

The second measure of geographic substitutability utilizes CTS product-class data on the tonnage shipped within each of seven distance-from-production-site categories. The measure, LOCAL, is a metric of the typical geographic size of industry producers' output markets. It is a weighted sum of the fraction of output (by ton) shipped within each distance category.¹⁶ Industries with plants that ship a large fraction of their output to nearby areas have high values of LOCAL and expectedly have less spatial product substitutability.¹⁷

Differences in the physical configuration of industry goods also create product

¹⁵ While there is close correlation between the CTS product categories and the corresponding four-digit SIC industries that contain them, they do not perfectly match. Using published descriptions of industry product types (from U.S. Office of Management and Budget, 1972), I was able to aggregate products into their corresponding industry for nearly every four-digit SIC. Shipment data for the ordnance industries (SICs 3282-4,9) were not available; these industries are not included in my sample. A concordance is available from the author upon request.

¹⁶ The categories are less than 100 miles, 100-199, 200-299, 300-499, 500-999, 1000-1499, and over 1500 miles. The weights are constructed as follows. The average shipment distance within each distance category is computed assuming uniformly distributed shipments within the category. (I use 1750 miles for shipments reported as over 1500 miles.) The sum of these seven distances is divided by each category's average distance, and these ratios are normalized so the weight of the under 100 miles category equals one. The resulting weights are: under 100 miles, 1; 100-199 miles, 0.333; 200-299, 0.2; 300-499, 0.133; 500-999, 0.067; 1000-1499, 0.04; and over 1500 miles, 0.029.

¹⁷ Use of LOCAL as a transport costs measure requires a caveat. If an industry's customers are geographically concentrated and industry plants choose to operate near their customers, it is possible that an industry could have high output substitutability despite a small average shipment radius. This is not a major issue for consumer goods industries, whose buyers are distributed throughout the country, but it may be for some industries which serve as suppliers of intermediate goods to specialized downstream buyers. I attempted to control for this possibility in specifications using LOCAL by including the measure of industry geographic concentration created by Ellison and Glaeser (1997). (Thanks to Glenn Ellison for providing the data at the four-digit SIC level.) Inclusion of this control did not noticeably change the coefficient on LOCAL.

substitutability barriers. Producers within textiles industries, for example, can have considerable variation in their product attributes. I measure this variation using the product differentiation index of Gollop and Monahan (1991). Their generalized Herfindahl-type product diversification index, DIVINDEX, takes into account not only the number of industry products (as defined by the SIC product classification system), but also how (un)equal the production shares of product lines are within the industry, as well as the dissimilarity of products as measured by the input shares of various intermediate products used to make them.

As a robustness check, I estimate specifications using a less sophisticated but more interpretable physical differentiation measure: the average primary product specialization ratio (PPSR) across industry establishments. PPSR is the fraction of plant revenue accounted for by products within the industry's primary product class (as assigned by the SIC product-level taxonomy). Industries with a lot of physical product variety are likely to be comprised of plants that produce a number of product types and therefore have low PPSR values. Conversely, a completely specialized industry will have an average PPSR equal to one.¹⁸

Product substitutability can also be shaped by advertising/branding effects. I account for such influences by using industry advertising expenditures data from the 1977 Benchmark Input/Output Tables. A measure of industry advertising intensity, ADV, is constructed as the ratio of advertising expenditures to total shipments.¹⁹

¹⁸ The across-industry correlation between DIVINDEX and PPSR is -0.813 , indicating a close (but naturally oppositely signed) correspondence. Both measures of physical product differentiation are limited by the SIC product classification system, of course. This leads to two vulnerabilities of DIVINDEX and PPSR as accurate gauges of product substitutability. Even the highly detailed seven-digit SIC product classification is only a blunt instrument for characterizing the enormous variety of manufactured products. A single SIC "product" may in truth encompass dozens, or even hundreds, of physically distinct products. The coarseness of the taxonomy will not cause empirical problems as long as product variety is undercounted at the same rate across industries, but this condition is unfortunately not testable. Second, product codes are somewhat inflexible over time, so new products that do not obviously fit into any of the categories in the existing system may be misclassified. As more new products are introduced, the original classification system matches the existing product space less completely. This drift should be minimized in this paper, because the SIC classification system underwent a major overhaul in 1972, not too long before my sample was taken. While neither vulnerability is a fatal flaw, one should keep them in mind when interpreting the results.

¹⁹ While the detailed BEA industry categorization used in the input-output tables roughly corresponds to the SIC

The industrial organization literature is divided on the question of the nature of the relationship between ADV and output substitutability.²⁰ One strand of research argues that advertising serves to create artificial product differentiation, largely along the lines of the branding motive discussed in the introduction. This view holds that industries with higher advertising intensities should exhibit more product differentiation. An opposing strand contends that advertising is informative and serves to educate consumers about superior products. Advertising expenditures under this view allow more productive firms to take market share away from less efficient competitors (by reducing search costs, say), increasing substitutability across industry producers. Of course, it is also possible that both effects act simultaneously; if their impacts have roughly equal magnitudes, estimates would show no overall influence of ADV on moments of industry productivity distributions.

The model above indicates that higher substitutability will be correlated with greater truncation of the plant-level productivity distribution, and therefore less dispersion and higher average productivity levels. Expressed in terms of the product substitutability measures used here, lower geographic segmentation (higher VALUELB or lower LOCAL) and less physical product differentiation (lower DIVINDEX or higher PPSR) correspond to greater product substitutability. The effect of higher advertising intensity (higher ADV) on substitutability is theoretically ambiguous.

B. Other Influences on the Productivity Distribution

The model and its extensions indicate factors besides product substitutability can shape industry productivity distributions. I include in X_{cl} controls for these other influences using variables constructed from several sources. It is not apparent beforehand whether excluding these other factors from the regressions would bias the substitutability coefficients, as the other factors may not be correlated with product substitutability. However, adding proxies for these other effects also allows further testing of the model's implications independent of any product substitutability effects, which is interesting in its own right.

four-digit system, data had to be pooled across some SIC industries to match more broadly defined BEA groups. Thus some four-digit industries have a common measured advertising-to-sales ratio.

²⁰ See, for example, the discussion in Tirole (1988) for a partial review of this literature.

As shown above, both sunk entry costs and fixed operating costs affect the critical productivity cutoff level, and therefore the moments of an industry productivity distribution. Controlling for these influences can be empirically difficult. It is not clear which fixed costs producers face are tied to entry and which are tied to operations. In the model, plants incur the sunk entry cost s before they learn their productivity level. If in reality plants must actually produce some output to learn their productivity levels, then any production-related overhead could be classified as either a fixed operating cost (f in the model) or an entry cost, at least in the first year of production. This scenario implies difficulties in separately measuring the influences of the two cost types, as they move the truncation point of the productivity distribution in opposite directions. Using observables that could be linked to either cost structure could yield inconclusive results. If, on the other hand, producers learn their productivity levels before starting operations, it is not immediately clear how to measure entry costs given that most data is collected after production is underway. I attempt to reconcile these confounding factors by assuming that sunk entry costs are related to post-production observables while using a fixed operating production cost measure that conceivably moves independently of entry costs.

I follow the method employed by Sutton (1991) to obtain a measure of sunk entry costs. This value, SUNKCOST, is the market share of an industry's median-sized plant multiplied by the capital-output ratio for the industry. The former element of this product is sometimes used as a measure of minimum efficient plant scale. Thus SUNKCOST is a proxy for the amount of capital (relative to the industry's total market size) required to build such a plant.²¹

Fixed operating costs are measured by the average ratio across industry establishments of non-production workers to total employment. This value, FIXEDCOST, proxies for the amount of overhead labor required by the industry technology. Because overhead labor is a fixed cost explicitly tied to production rather than entry, comovement between FIXEDCOST and entry costs should be related only through any inherent correlation between f and s in industry technologies, and not through erroneous measurement of entry costs. Both FIXEDCOST and SUNKCOST are constructed as proportions to remove scale effects across industries.

Some specifications include controls for international trade exposure. I use both import- and export-based metrics computed from the trade data discussed in Feenstra (1997). Industry

²¹ See Sutton (1991) for a thorough discussion of the advantages and limitations of this measure.

import penetration, $IMPPEN$, is the ratio of industry product imports to the sum of these imports plus the value of domestic production in the industry. Export intensity $EXPINT$ is the share of exports in total domestic output for the sector. Larger values of either variable should coincide with greater trade exposure. As mentioned above, Melitz (2003) shows that if these variables proxy for the extent of trade barriers in an industry (and therefore the extent of trade-driven truncation of the productivity distribution), they should expectedly have negative correlations with industry productivity dispersion and positive correspondence with the central tendency of industry productivity.

Differences in the ex-ante cost/productivity distribution $g(c)$ may also induce variation in industries' productivity distributions. To the extent that these are reflected in scale differences, normalizing dispersion moments to the median productivity level in the industry accounts for this influence. I also test for robustness of the dispersion results to the use of other interquantile differences, which allow determination of the effects across different subsets of the productivity distribution. While clearly not a flawless solution, these steps should remove a substantial amount of the influence of different ex-ante distributions across industries.

IV. Results

A. Benchmark Results

I first regress industry productivity distribution moments on each of the product substitutability measures. The results are presented in Table 3. Panel A shows the coefficients obtained by regressing each of the four productivity dispersion measures on the respective measures separately.²²

The results are consistent with the discussion above and the predictions of the model. Factors that plausibly increase industry product substitutability are negatively correlated with within-industry productivity dispersion. Regarding spatial substitutability, increases in $VALUELB$ (the average value per pound of an industry's output) and decreases in $LOCAL$ (an inverse measure of average shipment distance)—both of which correspond to greater substitutability—coincide with declines in productivity dispersion. Physical differentiation

²² Although the coefficients are listed in columns under the dispersion measures, the factor coefficients in this table are for single-variable (and a constant term) regressions.

factors play a similar role. Decreases in DIVINDEX (the Gollop-Monahan index) and increases in PPSR (the average fraction of plant revenue from the industry's primary product class), which indicate higher substitutability, are also negatively correlated with productivity dispersion. The coefficient on ADV (the industry ratio of advertising to sales) is positive in all regressions; industries with higher advertising intensities exhibit more measured productivity dispersion.

These findings are consistent across dispersion moments of all four productivity measures. For those substitutability measures with predicted directions of correlation with productivity dispersion, all of the estimates' signs are consistent the implications of the model. Further, 12 of these 16 are significant at the 5% level, and another one nearly so. These single-factor regressions indicate that productivity dispersion and product substitutability are undoubtedly correlated, and in directions consistent with theory.

Bivariate regressions with median industry TFP as the dependent variable are also largely consistent with expectations. As seen in Panel B of Table 3, industries with higher value-to-weight ratios and longer average shipment distances have higher median productivity levels on average. Greater physical product differentiation—a higher value of DIVINDEX or a lower PPSR—corresponds to a lower industry median TFP level. The corresponding coefficient estimates are consistently signed with the productions of the model and are significant at the 5% level in five of eight cases. The results from the bivariate regressions using ADV are more ambiguous: the coefficients are virtually zero, statistically speaking, and are oppositely signed.

I have two measures of both spatial substitutability (VALUELB and LOCAL) and physical product differentiation (DIVINDEX and PPSR). For the sake of brevity, and because the results in Table 3 suggest that all of these measures yield qualitatively similar results, from this point on I only report results for specifications using the measure of each that is less susceptible to measurement problems: VALUELB and DIVINDEX.²³ Estimates from regressions using the alternative measures, available from the author, largely match the findings presented below.

I next regress the moments on all output substitutability measures simultaneously. The outcomes are presented in Table 4. Again the consistency of the results is notable. In the

²³ Recall that LOCAL may confuse geographic clustering of an industry's customer base with low spatial differentiation. PPSR may indicate spuriously low product differentiation if industries spread many product types across a number of highly specialized plants.

productivity dispersion regressions (the first four columns), all estimates for VALUELB and DIVINDEX have the expected sign, and most are statistically significant at the 5% level (all are at 10%). As with the single-variable regressions, there is a positive and usually statistically significant correlation between advertising intensity and productivity dispersion. The estimated magnitudes of the responses to differences in the substitutability measures are nontrivial. A quadrupling of value density (which ranges from \$0.01 to \$150 per pound in my sample) corresponds to a decline in the labor productivity dispersion measures of roughly one-sixth of their standard deviation, and one-fifth of a standard deviation of the TFP dispersion measures. A one-standard-deviation increase in DIVINDEX coincides with a quarter-standard-deviation drop in labor productivity variability and one-eighth-standard-deviation decrease in TFP dispersion. An increase of one percentage point in advertising-to-sales ratio corresponds with labor productivity and TFP dispersion increases of roughly one-ninth to one-tenth of their respective standard deviations.

The product substitutability measures jointly explain seven to 12% of across-industry differences in productivity dispersion. Thus substantial productivity heterogeneity remains to be explained. This is not surprising given all of the across-industry variation in technological/supply-side influences that shape productivity distributions. Product substitutability is surely an economically relevant part of the story, however.²⁴ Furthermore, as argued in the introduction, there are also non-measurable product differentiation influences whose impacts obviously cannot be captured here. The present results do hint that these non-measurable factors affect industries' plant-level productivity distributions in the same manner as their measurable counterparts. It is possible that the combined effect of measurable and non-measurable substitutability differences is considerable.

The two rightmost columns of Table 4 show estimates from the median TFP regressions. Jointly estimating coefficients for all substitutability factors largely preserves the findings of the bivariate regressions. VALUELB is positive and highly significant in both TFP regressions. A quadrupling of value density across industries corresponds with an increase of about 10% in the

²⁴ Measurement error in plant productivity levels, doubtlessly present in establishment-level data sets, will create spurious productivity dispersion. Thus, the variation in true productivity dispersion moments explained by measurable product substitutability factors may be greater than the amount measured here.

median productivity level. DIVINDEX is negatively related to median productivity levels. The (insignificantly) estimated coefficients imply each standard deviation increase in the index corresponds to a drop in the median productivity level of 2 to 3%. Given the negative partial correlation seen between advertising intensity with productivity dispersion, the estimates of ADV from the central-tendency regressions are slightly puzzling. While the direction of the linkage between advertising intensity and substitutability cannot be pinned down theoretically, we should expect empirically that ADV has oppositely signed correlations with industry productivity dispersion and median productivity levels. These ADV estimates are much less precisely estimated than those in the dispersion regressions, however, so the positive coefficients here may be spurious. The product substitutability measures jointly explain roughly 13-14% of the variance in median productivity levels.

The empirical model is further enriched by adding SUNKCOST and FIXEDCOST as controls. The results are presented in Table 5. Apparently the correlations found above do not spuriously arise from comovement between the substitutability factors and features of the cost structure of industry technologies. In the productivity dispersion regressions, VALUELB and DIVINDEX retain their expected signs and are significantly estimated in every case. ADV still has a significant positive coefficient in the dispersion regressions, excepting a negative and insignificant coefficient in the model using the TFP measure computed with plant-specific input elasticities. The magnitudes of the substitutability coefficients are similar to those obtained without controlling for these cost measures.

Furthermore, the across-industry comovement between the sunk entry cost measure and productivity dispersion is as predicted by the model. The coefficient on SUNKCOST is positive and statistically distinguishable from zero in every specification. The implications of the FIXEDCOST coefficients are more ambiguous. With regard to labor productivity dispersion, fixed production costs are found to have their expected negative correspondence. However, the TFP dispersion results are not in tune with the model's predictions: both FIXEDCOST coefficients are positively signed, and one of these is precisely estimated. Adding these two controls for sunk entry and fixed production costs improve the model's explanatory power slightly, with a typical increase in R^2 of about 0.03.

The results of the median industry productivity level regressions with SUNKCOST and FIXEDCOST echo the findings of the dispersion regressions. Value-to-weight ratios and the

product diversity index have nonzero coefficients with the expected signs. Unlike in the productivity level regressions without fixed cost controls, the ADV coefficients in this specification are negatively signed (albeit insignificantly estimated). This is consistent with the estimated positive correlation between ADV and productivity dispersion. All SUNKCOST and FIXEDCOST estimated coefficients have the expected signs. In a turnabout from the dispersion regressions, here it is the FIXEDCOST estimates that are statistically significant while the SUNKCOST coefficients are less precise. The explanatory power of the model for across-industry differences in median productivity levels ranges from 15 to 18%.²⁵

These findings largely agree both with intuitive priors and the model presented above. Measurable product substitutability factors have expected correlations with moments of industries' plant-level productivity distributions. More spatially localized industries have plant-level productivity distributions with greater dispersion and lower central tendencies. Industries with greater physical product differentiation have higher dispersion and lower median productivity levels. Advertising intensity is positively correlated with productivity dispersion, and at least when technology controls are added, is weakly associated with declines in median productivity levels. Moreover, these results are found in a number of empirical specifications. Substitutability factors are correctly signed and significantly estimated in models ranging from simple bivariate correlations to those including other substitutability measures and controls for sunk entry and fixed operating costs. The results also hold across several productivity measures. The estimated effects of sunk and fixed costs, while somewhat weaker than those for the product substitutability measures because of the difficulty finding measurable proxies, are also on

²⁵ SUNKCOST combines measures of the median establishment-level market share and the capital-to-output ratio in an industry. Arguably, either of these could be related to industries' cost structures separately through other channels. While its interpretation is not critical to the key results here on product substitutability, the results indicate that the measure may be related to the shape of within-industry productivity distributions in directions predicted by theory. To see if the two components of SUNKCOST have separable effects, I ran a specification where the components entered separately. Median market share was significantly and positively correlated with productivity dispersion and negatively correlated with the median productivity level in the industry. These correlations are the same as those implied for SUNKCOST. On the other hand, increases in the industry capital-to-output ratio were associated (statistically significantly) with less productivity dispersion and higher median productivity levels. Entering the components separately had no substantive effect on the product substitutability measure coefficients.

balance consistent with the theory.

B. Robustness Checks

To see if the results discussed above hold in more generalized frameworks, I have conducted several robustness checks. These are described here.

Foreign Trade. I estimate a specification that includes in X_{cl} industry-level measures of import penetration and export intensity (IMPPEN and EXPINT) to see if international trade affects the measured relationship between product substitutability and productivity distribution moments. The results are presented in Table 6. Importantly, the qualitative and quantitative features of the substitutability factor and sunk and fixed costs estimates are unaffected when the extra controls are included. It does not appear that the results obtained above arise due to omitted variable bias regarding this other influence.

As for the coefficients on the additional controls, the weight of the evidence suggests (although a greater proportion of their coefficients are statistically insignificant) that industries with greater exposure to international trade (higher IMPPEN and EXPINT) have more productivity variability. This correspondence is counter to the implications of Melitz (2003). The benefits of exposure to foreign markets enjoyed by the more productive domestic firms should drive the least efficient domestic producers out of business, thereby decreasing productivity dispersion. Perhaps the positive comovement seen here is explained in part by reverse causation, if foreign producers deliberately target industries with wide productivity distributions to better their relative competitive position. It could also be that foreign trade serves to increase product differentiation in an industry, counteracting the competitive effect of trade productivity dispersion.

The regressions using median industry productivity as the dependent variable are also largely unaffected by the inclusion of the additional controls. Product substitutability factors are still correlated with productivity distributions' central tendencies in the predicted manner. Advertising intensity again has a negative but weak comovement with the median. Three of the four sunk and fixed cost coefficients have the expected sign, but only the FIXEDCOST coefficients are statistically significant. The coefficients for the trade exposure controls yield results that, because they have oppositely signed correlations with the median productivity, are

consistent with their measured correlations with productivity dispersion. All are insignificant, however.

Productivity Dispersion Measure. I check if the results are sensitive to the productivity dispersion measure (the interquartile range divided by the median) by performing the industry productivity dispersion regressions using the difference between the 90th and 10th productivity percentiles divided by the median as the dispersion measure. These results are shown in Table 7. The results are qualitatively consistent with those presented earlier. The FIXEDCOST coefficients gain statistical significance in the labor productivity regressions, while DIVINDX and SUNKCOST lose significance in the TFP dispersion regressions. Excepting these differences, the choice of the range over which to measure productivity dispersion does not seem to greatly influence the results.²⁶

Dynamics. The theoretical framework abstracts from dynamic evolution of producers' productivity levels. Asplund and Nocke (2003) and Melitz (2003) show that changes in the persistence of the (exogenous) productivity process affect the cutoff productivity level; more persistence implies a higher cutoff. To control for the possible influence of across-industry difference in producer-level productivity persistence, I estimate a specification that includes in X_{ct} the fraction of industry plants in the 1972 Census of Manufactures—the census that most immediately precedes my data—that still operate (in the same industry) in the 1977 CM.²⁷ This measure, SURVRT, is meant to proxy for the persistence of the industry's plant-level productivity evolution. A higher SURVRT value, all else being equal, implies a higher probability that operating plants receive updated productivity draws above the threshold (i.e., a higher ρ). Increases in SURVRT should expectedly be correlated with lower productivity dispersion and higher median TFP levels.²⁸

²⁶ The magnitudes of the coefficients are changed because the dependent variable is now scaled differently. I also estimated the model using the 95th-5th percentile range and obtained similar results.

²⁷ I do include Administrative Record plants when computing survival rates, because here I only need to know of their existence, not their production specifics.

²⁸ Controlling for differences in industries' productivity evolution processes with survival rates is at best an

The results of this exercise are shown in Table 8. The results for both the productivity dispersion and the central tendency regressions closely match those in Table 6. Furthermore, the coefficients on SURVRT have the predicted signs. While the sensitivity of survival rates to changes in industry equilibria means that SURVRT may capture more than exogenous influences on the evolution of plant productivity levels, it does suggest that any mismeasurement in this regard is not correlated with the other regressors in a way that would affect the benchmark results.

Capital Measurement. As discussed above, I have excluded Administrative Record plants from my sample because much of their production data is imputed. The remaining establishments report virtually all production data directly. The exceptions to this are establishments not in the current Annual Survey of Manufactures (ASM) panel. (Roughly 35% of the 198,000 establishments in my sample are in the ASM panel.) These plants have imputed capital stocks, and these imputations can then in turn lead to TFP mismeasurement. To ensure that productivity mismeasurement arising from capital stock imputation is not driving the results, I re-estimate the TFP specifications using only ASM plants to compute within-industry productivity moments.

Another form of capital mismeasurement arises when plants vary their capital services inputs by changing the intensity with which they utilize their capital stock. When this is the case, the value of a plant's capital stock does not accurately reflect capital's contribution to production. Systematic differences in capital utilization patterns across industries could potentially affect the above findings for those specifications using TFP measures.

To check for the influence of variable capital utilization, I estimate a specification where I compute plant TFP levels according to the suggestion of Basu and Kimball (1997). They show that under the assumptions of cost minimization by plants and a production function that is Leontief in capital services and materials (i.e., one cannot substitute extra materials for capital in production) that TFP can be defined as

imperfect solution. Measured survival rates are likely to confound any underlying dynamics in the producer-level productivity process with the industry equilibrium effects of changes in product market and technological parameters that also affect exit rates. (Indeed, exploring these effects is an interesting avenue for future research.) Hence I only report results including SURVRT as a robustness check rather than incorporating this control into the main specification.

$$tfp_i = y - \alpha_l l_i - (\alpha_k + \alpha_m) m_i - \alpha_e e_i - \delta \alpha_l h_i,$$

where variables are defined as in the standard TFP measure above, h_i is the log of hours per employee, and δ is a parameter which Basu and Kimball estimate as having a value of 1.06. Intuitively, including the product of labor's cost share and logged hours per worker and adding capital's cost share to material's cost share in the TFP measure controls for increases in capital utilization intensity because materials use is proportional to capital services flows. Including hours per worker captures variations in the intensity of labor utilization.

The results of these exercises can be found in Table 9.²⁹ It does not appear that capital—and therefore TFP—mismeasurement due to either through capital stock imputation or variable utilization is driving the results obtained above. While some precision is lost in the ASM-plant-only results, the product substitutability estimates, both in the productivity dispersion and central tendency regressions, are qualitatively (and to a lesser extent, quantitatively) comparable to those discussed above. The notable difference is that the (still imprecisely estimated) coefficients on ADV in the central tendency regressions are now positive.

V. Caveats

The key results linking output substitutability and the moments of the industry productivity distribution seem to be robust across a number of empirical specifications, with the possible exception of a weak connection between industry advertising intensity and the central tendency of the productivity distribution. As with nearly any study, however, the findings come with some caveats. Several potential concerns are discussed below, along with mitigating factors that may minimize their influence on the key results.

It is important to note that the empirical tests above were performed using moments of *measured* productivity distributions. The introductory discussion and the model make the common conceptualization that productivity is the efficiency of input use with respect to production of output, and as such is related to production costs. Empirical productivity measures, however, are not so cleanly obtained. One particular difficulty is that producer output is measured in terms of revenue rather than more appropriate units, due to a lack of

²⁹ For the specification adjusting for variable utilization, I again use the full sample of plants, including those not in the ASM.

comprehensive physical output data or plant-specific deflators. Plant-level price variation enters into output measures and can create measured productivity variation independent of efficiency differences.

Melitz (2000) points out one effect of this is an under-measurement bias in between-plant productivity differences, where the size of this bias is larger when the elasticity of output substitution is low. Notice, however, that this effect works against my empirical results. Intra-industry productivity dispersion should be most underestimated when substitutability is lowest. The results above, despite this possible influence, show that productivity dispersion in an industry *grows* with the amount of product substitution. The true dispersion might be even greater than measured in low-substitutability industries, and the true correlation with substitutability greater.

A similar byproduct of revenue-based output measures is that any across-industry differences markups would create variation in median revenue-based TFP through price effects. If high fixed operating costs in an industry support higher average markups, for instance, this would appear as a higher median TFP. It is possible that this is in part driving the related results in the central tendency regressions. However, the influence of markups on median TFP differences may also tend to work against some of the results above rather than spuriously create them. If higher markups are sustainable in low substitutability industries, this would induce a negative correlation between measured TFP and substitutability: revenue-based TFP would tend to overstate (understate) true productivity in low- (high-) substitutability industries. The empirical results indicate that despite this possibility, average TFP levels tend to instead be positively correlated with substitutability.

The measurement problems inherent to quantifying sunk and fixed costs are additional empirical hurdles. While I took care in finding proxies for these influences on productivity distributions, the resulting controls are at best approximate. However, this concern is balanced by the facts that the product substitutability results are qualitatively invariant to inclusion of technological controls, and that the sunk entry and fixed operating cost proxies are usually observed to be correlated with productivity moments in the expected directions. Additionally, even if these proxies were not able to separately identify the influence of sunk and fixed costs, their inclusion still can remove their joint influence on industry productivity distributions.

The model also assumes that product substitutability varies exogenously across industries. This is not necessarily true in reality. While across-industry differences in spatial substitutability (caused by the inherent physical characteristics of products) might be presumed to be out of the hands of producers, advertising intensity certainly is not, and perhaps a great deal of physical differentiation is endogenous as well. Productivity dispersion itself may cause producers to increase differentiation by physically altering their product or through their advertising behavior. Less efficient plants might have the incentive to put greater distance in product space between themselves and their more efficient competitors. This would reverse the direction of causation implied by the present theory. While an exploration of such producer efforts is a worthy research topic, it is beyond the scope of this paper. My focus here is determining whether there is indeed a systematic relationship between product substitutability and moments of the productivity distribution, not characterizing the causal connection. Since endogenous substitutability would preserve the direction of the correlations implied by the model, the correspondences found here are consistent with the presence of both exogenous and endogenous substitutability differences.

Finally, exploiting across-industry differences to empirically test the implications of the model raises issues of inter-industry heterogeneity affecting the results. Unfortunately, data limitations prevent the use of within-industry substitutability changes from being measured at a disaggregate level. However, I reference the aforementioned within-industry case study of similar concept in Syverson (2003). The results there within an industry are consistent with those found here across industries, suggesting that the links between substitutability and productivity moments found above are legitimate.

VI. Conclusion

The evidence presented in the paper suggests that product substitutability—a characteristic of industry demand—is systematically related to the shape of the industry's equilibrium plant-level productivity distribution. Measurable factors likely correlated with high substitutability, such as low transport costs and less physical product differentiation, are shown to be negatively related with productivity dispersion and positively with median productivity in an industry. These findings are robust; they are found both in simple bivariate correlations as well as when controls for other influences on industries' productivity distributions are included

in empirical specifications. Additionally, the empirical results suggest that across-industry differences in these other influences on the productivity distribution, such as the size of sunk entry and fixed operating costs, are correlated with variability in productivity distribution moments in the expected direction. The exception to this is an industry's trade exposure, which seems to correspond with productivity moments in directions opposite that predicted in Melitz (2003).

These results suggest that, while the technological/supply-side factors that have been the focus of the related literature doubtlessly play a role in creating productivity dispersion, demand-side influences are also important. Measurable substitutability factors explain a nontrivial fraction of the total interindustry variation in productivity moments. Further, additional unmeasured (or unmeasurable) types of substitutability barriers may explain some of the remaining variation. Exploring the specific output market mechanisms driving these results may be a fruitful path for further research.

The findings offer help in understanding why productivity differences exist within industries and what factors affect their magnitude, a puzzle discussed in the introduction. On a broader scale, they also lend insight into how aggregate productivity dynamics might be affected, either by shifts in output shares across industries with different shapes of their productivity distribution, or by shifts over time of the product substitutability factors within industries.

Appendix: Incorporating Transport Costs into the Model

Assume industries operate in two markets, each identical to the one in the base model. I assume, as in Melitz and Ottaviano (2003), that producers can produce not only for their home market (as above), but for the outside market if they so choose. Selling to the outside market, however, involves paying transport costs to ship output. I assume these are of the standard “iceberg” variety, where $\tau > 1$ units must be shipped for one unit to arrive. Thus the marginal cost of producing for the outside market is τc_i , while still only c_i for the home market. A large set of potential entrants in each market considers the entry decision specified above.

Producers can charge different prices in each market. Therefore the optimal price in the home market is still that given by (5) above, but now the optimal price and resulting demand in the outside market are given by:

$$p_{i,outside} = \frac{1}{2} \frac{\alpha\gamma}{\eta N + \gamma} + \frac{1}{2} \frac{\eta N}{\eta N + \gamma} \bar{p} + \frac{\tau}{2} c_i \quad (\text{A.1})$$

and

$$q_{i,outside} = \frac{1}{2} \frac{\alpha}{\eta N + \gamma} + \frac{1}{2\gamma} \frac{\eta N}{\eta N + \gamma} \bar{p} - \frac{\tau}{2\gamma} c_i. \quad (\text{A.2})$$

Variable profits from the outside market are then:

$$\pi_{i,outside} = \frac{1}{4\gamma} \left(\frac{\alpha\gamma}{\eta N + \gamma} + \frac{\eta N}{\eta N + \gamma} \bar{p} - \tau c_i \right)^2. \quad (\text{A.3})$$

Note that because of the symmetric markets assumption, the average price \bar{p} (which includes prices of both home and exporting outside-market producers) is the same in both markets.

From inspection of (6) and (A.2), it is clear that there will be a range of productivity levels for which home-market sales are positive but will be zero for the outside market. Define c_h^* and c_o^* and as the cost levels where $q_{i,home} = 0$ and $q_{i,outside} = 0$, respectively. From (6) and (A.2), these levels are

$$c_h^* = \frac{\alpha\gamma}{\eta N + \gamma} + \frac{\eta N}{\eta N + \gamma} \bar{p} \quad \text{and} \quad c_o^* = \frac{1}{\tau} \left(\frac{\alpha\gamma}{\eta N + \gamma} + \frac{\eta N}{\eta N + \gamma} \bar{p} \right) \quad (\text{A.4})$$

Clearly, any potential entrants drawing $c_i > c_h^*$ will choose not to operate, since they would have zero sales in both markets. However, because of the fixed operating costs, there will also be a set of producers with $c_i < c_h^*$ that will find production unprofitable. If fixed costs are not too large, there will be a cost level $c^* \in [c_o^*, c_h^*]$ where a producer with $c_i = c^*$ will not sell to the outside market and will be just indifferent to operating in the home market, since the variable profits from doing so are just enough to cover the fixed operating costs.³⁰ This level is given by

³⁰ The fixed cost must satisfy $c^* \geq c_o^*$ in equilibrium; i.e.,

$$f \leq \frac{1}{4\gamma} \left[\frac{\tau - 1}{\tau} \left(\frac{\alpha\gamma}{\eta N + \gamma} + \frac{\eta N}{\eta N + \gamma} \bar{p} \right) \right]^2.$$

If fixed operating costs are larger than this, but not so large as to make any entry unprofitable, only those producers who can sell in both markets will choose to operate, since the extra sales are necessary to recoup the high fixed

(8), and the relationship between c_o^* and c^* is

$$c_o^* = \frac{1}{\tau} \left(c^* + 2\sqrt{\gamma f} \right). \quad (\text{A.5})$$

Expression (A.4) can be substituted into (A.3) to obtain

$$\pi_{i, \text{outside}} = \frac{\tau^2}{4\gamma} (c_o^* - c_i)^2, \quad (\text{A.6})$$

and the value of entry is now

$$V^e = \int_0^{c_o^*} \left[\frac{1}{4\gamma} (2\sqrt{\gamma f} + c^* - c)^2 - f \right] g(c) dc + \int_0^{c_o^*(c^*)} \left[\frac{\tau^2}{4\gamma} (c_o^* - c)^2 \right] g(c) dc - s = 0. \quad (\text{A.7})$$

I have explicitly noted in the above expression that c_o^* is a function of c^* , with $\partial c_o^*/\partial c^* > 0$. The comparative static of interest is $dc^*/d\tau$. The relevant components of the implicit function theorem are as follows (note that c_o^* is also a function of τ):

$$\frac{\partial V^e}{\partial \tau} = -\frac{c^* + 2\sqrt{\gamma f}}{4\gamma} (c_o^* - c_o^*)^2 g(c_o^*) + \int_0^{c_o^*(c^*)} \frac{\tau}{2\gamma} (c_o^* - c) \left[(c_o^* - c) + \tau \frac{\partial c_o^*}{\partial \tau} \right] g(c) dc \quad (\text{A.8a})$$

Using the fact that $\partial c_o^*/\partial \tau = -c_o^*/\tau$, this simplifies to

$$\frac{\partial V^e}{\partial \tau} = \int_0^{c_o^*(c^*)} -\frac{\tau}{2\gamma} c (c_o^* - c) g(c) dc < 0, \quad (\text{A.8b})$$

which is negative since the integral is over $c \leq c_o^*$. Further using the fact that $\partial c_o^*/\partial c^* = 1/\tau$,

$$\frac{\partial V^e}{\partial c^*} = \int_0^{c_o^*} \frac{1}{2\gamma} (2\sqrt{\gamma f} + c^* - c) g(c) dc + \int_0^{c_o^*(c^*)} \left[\frac{\tau}{2\gamma} (c_o^* - c) \right] g(c) dc > 0. \quad (\text{A.9})$$

Therefore $dc^*/d\tau > 0$; higher transport costs are akin to lower product substitutability. They support efficiency gaps between industry competitors because they act as a barrier keeping certain consumers from shifting purchases to more productive producers.

costs. Here, I consider the case where there are both home-only and home-and-outside market producers. Melitz and Ottaviano (2003) obtain similar qualitative implications in the exporters-only case with no fixed operating costs.

References

- Asplund, Marcus and Volker Nocke, "Firm Turnover in Imperfectly Competitive Markets," Penn Institute for Economic Research Working Paper 03-010, 2003.
- Bartelsman, Eric J. and Mark Doms, "Understanding Productivity: Lessons from Longitudinal Microdata," *Journal of Economic Literature* 38:3 (2000), 569-595.
- Basu, Susanto and John G. Fernald, "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy* 105:2 (1997), 249-283.
- Basu, Susanto and Miles S. Kimball. "Cyclical Productivity with Unobserved Input Variation." NBER Working Paper no. 5915, February 1997.
- Davis, Steven J. and John Haltiwanger, "Wage Dispersion between and within U.S. Manufacturing Plants, 1963-1986," *Brookings Papers on Economic Activity, Microeconomics* (1991), 115-180.
- Ellison, Glenn and Edward L. Glaeser, "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach," *Journal of Political Economy* 105:5 (1997), 889-927.
- Feenstra, Robert C., "NBER Trade Database, Disk 3: U.S. Exports, 1972-1994, with State Exports and Other U.S. Data," National Bureau of Economic Research Working Paper no. 5990 (1997).
- Gollop, Frank M. and James L. Monahan, "A Generalized Index of Diversification: Trends in U.S. Manufacturing," *Review of Economics and Statistics* 73:2 (1991), 318-330.
- Hallward-Driemeier, Mary, Giuseppe Iarossi, and Kenneth L. Sokoloff, "Exports and

- Manufacturing Productivity in East Asia: A Comparative Analysis of Firm-Level Data,” National Bureau of Economic Research Working Paper no. 8894 (2002).
- Harper, Michael J., Ernst R. Berndt, and David O. Wood, “Rates of Return and Capital Aggregation Using Alternative Rental Prices,” in D.W. Jorgenson and R. London (eds.), *Technology and Capital Formation* (Cambridge, MA: MIT Press, 1989).
- Hopenhayn, Hugo A., “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica* 60:5 (1992), 1127-1150.
- Melitz, Marc J., “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica* 71:6 (2003), 1695-1725.
- Melitz, Marc J., “Estimating Firm-Level Productivity in Differentiated Product Industries,” Mimeo, Harvard University, (2000).
- Melitz, Marc J. and Gianmarco I. P. Ottaviano, “Estimating Firm-Level Productivity in Differentiated Product Industries,” Mimeo, Harvard University (2003).
- Sutton, John, *Sunk Costs and Market Structure* (Cambridge, MA: MIT Press, 1991).
- Syverson, Chad, “Market Structure and Productivity: A Concrete Example,” Mimeo, University of Chicago (2003).
- Tirole, Jean, *The Theory of Industrial Organization* (Cambridge, MA: MIT Press, 1988).
- U.S. Bureau of the Census, *1977 Commodity Transport Survey* (Washington, DC: U.S. Government Printing Office, 1980).

U.S. Bureau of Labor Statistics. *Trends in Multifactor Productivity, 1948-81*, Bulletin 2178
(Washington, DC: U.S. Government Printing Office, 1983).

Table 1: Dispersion Across Industries of Within-Industry Productivity Distribution Moments

Productivity Measure	Within-Industry Productivity Moment	Mean	Std Dev.	IQ Range
Labor Productivity log(Value Added / Employees)	Median	3.174	0.407	0.449
	IQ Range	0.662	0.208	0.213
	90-10 Percentile Range	1.417	0.388	0.407
	95-5 Percentile Range	2.014	0.568	0.565
Labor Productivity log(Value Added / Hours)	Median	2.521	0.376	0.428
	IQ Range	0.653	0.216	0.242
	90-10 Percentile Range	1.391	0.389	0.391
	95-5 Percentile Range	1.969	0.553	0.570
Total Factor Prod. 1 (Plant-Specific Input Elasticities)	Median	1.642	0.370	0.474
	IQ Range	0.447	0.146	0.153
	90-10 Percentile Range	0.986	0.238	0.276
	95-5 Percentile Range	1.356	0.291	0.329
Total Factor Prod. 2 (Industry Average Input Elasticities)	Median	1.790	0.342	0.430
	IQ Range	0.290	0.087	0.102
	90-10 Percentile Range	0.651	0.173	0.196
	95-5 Percentile Range	0.935	0.233	0.296

Note: This table summarizes plant-level productivity distribution moments across 443 (four-digit SIC) manufacturing industries. Rows correspond to moments of within-industry producer productivity distributions; columns show the across-industry mean and dispersion of these moments. “IQ range” is the interquartile range. Due to data disclosure restrictions, calculated medians are actually the average of 49th and 51st percentile values.

Table 2: Summary Statistics of Regression Variables

Variable	Mean	Std. Dev.
Value Added per Employee Dispersion	0.209	0.058
Value Added per Employee-Hour Dispersion	0.259	0.073
TFP1 Dispersion	0.293	0.145
TFP2 Dispersion	0.165	0.047
VALUELB (dollar value per lb. of shipments)	0.352	1.641
LOCAL (index of average shipment distance)	0.412	0.163
DIVINDEX (physical product differentiation index)	0.151	0.096
PPSR (primary product specialization ratio)	0.896	0.068
ADV (advertising expenditures per dollar of revenue)	0.010	0.019
FIXEDCOST (fixed operating costs measure)	0.232	0.093
SUNKCOST (sunk entry costs measure)	2.66E-3	0.012
IMPPEN (import penetration rate)	0.076	0.104
EXPINT (export intensity)	0.059	0.069

Note: This table presents summary statistics of variables used in the empirical tests. See the text for detailed descriptions of the variables. Reported moments are across 443 (four-digit SIC) manufacturing industries.

Table 3: Regression Results—Bivariates on Product Substitutability Factors

A. Dispersion Regressions (Interquartile Range ÷ Median Productivity as Dependent Variable)

Substitutability Factor	Productivity Measure			
	Labor (Output/Employee)	Labor (Output/Hour)	TFP1	TFP2
VALUELB	-6.29E-3* (1.69E-3)	-7.21E-3* (2.12E-3)	-0.022* (0.004)	-6.01E-3* (1.52E-3)
LOCAL	0.048* (0.170)	0.045* (0.021)	0.082* (0.041)	0.023 (0.014)
DIVINDEX	0.062* (0.301)	0.076 (0.039)	0.293* (0.088)	0.089* (0.030)
PPSR	-0.064 (0.041)	-0.084 (0.055)	-0.553* (0.127)	-0.126* (0.043)
ADV	0.703* (0.151)	0.784* (0.170)	0.238 (0.244)	0.543* (0.098)

B. Central Tendency Regressions (Median Productivity as Dependent Variable)

Substitutability Factor	Productivity Measure	
	TFP1	TFP2
VALUELB	0.079* (0.012)	0.079* (0.011)
LOCAL	-0.286* (0.115)	-0.126 (0.103)
DIVINDEX	-0.151 (0.240)	-0.239 (0.226)
PPSR	0.695* (0.345)	0.848* (0.321)
ADV	0.031 (0.727)	-0.250 (0.714)

Note: The table reports the coefficients obtained when within-industry productivity moments are regressed on substitutability measures. The upper panel shows results for productivity dispersion moments, the bottom for median productivity levels. Note that the reported coefficients are from bivariate regressions—the productivity moment is regressed on each substitutability measure separately (all regressions include a constant). TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Table 4: Regression Results—All Substitutability Factors

Subst. Factor	Productivity Dispersion Regressions				Central Tendency Regressions	
	Labor (emp)	Labor (hours)	TFP1	TFP2	TFP1	TFP2
VALUELB	-6.23E-3* (1.66E-3)	-7.15E-3* (2.09E-3)	-0.023* (0.004)	-6.07E-3* (1.49E-3)	0.079* (0.012)	0.079* (0.011)
DIVINDEX	0.052 (0.029)	0.065 (0.038)	0.308* (0.084)	0.083* (0.028)	-0.212 (0.217)	-0.295 (0.201)
ADV	0.663* (0.143)	0.735* (0.163)	0.028 (0.256)	0.487* (0.088)	0.303 (0.597)	0.067 (0.611)
R ²	0.094	0.074	0.106	0.118	0.125	0.149

Note: The table reports the coefficients obtained when within-industry productivity moments are regressed on the substitutability measures simultaneously. TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Table 5: Regression Results—Model with Sunk and Fixed Costs

	Productivity Dispersion Regressions				Central Tendency Regressions	
	Labor (emp)	Labor (hours)	TFP1	TFP2	TFP1	TFP2
VALUELB	-5.56E-3* (1.70E-3)	-5.94E-3* (2.13E-3)	-0.022* (0.004)	-5.41E-3* (1.50E-3)	0.077* (0.012)	0.076* (0.011)
DIVINDEX	0.067* (0.031)	0.081* (0.039)	0.254* (0.089)	0.068* (0.030)	-0.416 (0.216)	-0.482* (0.200)
ADV	0.705* (0.156)	0.796* (0.178)	-0.017 (0.259)	0.483* (0.088)	-0.098 (0.641)	-0.330 (0.621)
SUNKCOST	0.418* (0.143)	0.850* (0.182)	1.726* (0.578)	0.679* (0.123)	-0.008 (0.624)	-0.849 (0.502)
FIXEDCOST	-0.059 (0.031)	-0.075* (0.037)	0.156* (0.072)	0.038 (0.027)	0.726* (0.163)	0.683* (0.167)
R ²	0.110	0.104	0.136	0.153	0.155	0.181

Note: The table reports the results obtained when industry-level sunk entry cost and fixed operating cost measures are added. TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Table 6: Regression Results—Model with Trade Exposure Measures

	Productivity Dispersion Regressions				Central Tendency Regressions	
	Labor (emp)	Labor (hours)	TFP1	TFP2	TFP1	TFP2
VALUELB	-6.60E-3* (1.87E-3)	-7.93E-3* (2.44E-3)	-0.023* 0.004	-7.20E-3* (1.70E-3)	0.080* (0.012)	0.080* (0.011)
DIVINDX	0.071* (0.031)	0.087* (0.040)	0.249* (0.089)	0.068* (0.029)	-0.422* (0.215)	-0.481* (0.199)
ADV	0.689* (0.158)	0.780* (0.179)	0.025 (0.262)	0.496* (0.086)	-0.088 (0.645)	-0.365 (0.625)
SUNKCOST	0.284 (0.162)	0.625* (0.212)	1.668* (0.591)	0.538* (0.141)	0.341 (0.707)	-0.596 (0.578)
FIXEDCOST	-0.044 (0.031)	-0.056 (0.037)	0.139 (0.081)	0.038 (0.028)	0.703* (0.190)	0.687* (0.192)
IMPPEN	0.074* (0.027)	0.122* (0.033)	0.024 (0.080)	0.072 (0.024)	-0.187 (0.170)	-0.129 (0.134)
EXPINT	-0.002 (0.055)	0.025 (0.072)	0.122 (0.135)	0.075 (0.056)	-0.077 (0.288)	-0.159 (0.245)
R ²	0.126	0.131	0.139	0.188	0.158	0.183

Note: The table reports the results obtained when industry-level trade exposure measures are added. TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Table 7: Regression Results—Alternative Dispersion Measure

Productivity Dispersion Regressions				
	Labor (emp)	Labor (hours)	TFP1	TFP2
VALUELB	-9.37E-3* (3.68E-3)	-0.011* (0.005)	-0.050* 0.010	-0.013* (0.003)
DIVINDEX	0.063* (0.061)	0.089* (0.080)	0.424 (0.217)	0.087 (0.066)
ADV	0.963* (0.243)	1.037* (0.302)	0.075 (0.472)	0.897* (0.319)
SUNKCOST	0.746 (0.755)	0.948* (0.949)	0.420 (1.136)	0.267 (0.334)
FIXEDCOST	-0.127* (0.061)	-0.194* (0.075)	0.169 (0.145)	-0.027 (0.059)
IMPPEN	0.177* (0.059)	0.249* (0.073)	-0.005 (0.149)	0.107 (0.061)
EXPINT	-0.030 (0.107)	-0.003 (0.129)	0.335 (0.238)	0.122 (0.103)
R ²	0.098	0.099	0.104	0.081

Note: The table reports the results obtained when using an alternative productivity dispersion measure (based on the range between the 10th and 90th percentile plants in the industry productivity distribution rather than the interquartile range). TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Table 8: Regression Results—Model Controlling for Survival Rate

	Productivity Dispersion Regressions				Central Tendency Regressions	
	Labor (emp)	Labor (hours)	TFP1	TFP2	TFP1	TFP2
VALUELB	-7.96E-3* (1.80E-3)	-9.26E-3* (2.39E-3)	-0.022* 0.004	-7.60E-3* (1.71E-3)	0.084* (0.013)	0.079* (0.012)
DIVINDX	0.097* (0.033)	0.112* (0.042)	0.235* (0.090)	0.076* (0.031)	-0.486* (0.221)	-0.477* (0.205)
ADV	0.650* (0.154)	0.743* (0.176)	0.046 (0.261)	0.485* (0.085)	0.007 (0.635)	-0.370 (0.633)
SUNKCOST	0.345* (0.133)	0.684* (0.173)	1.635* (0.580)	0.555* (0.134)	0.189 (0.712)	-0.587 (0.564)
FIXEDCOST	-0.046 (0.030)	-0.057 (0.036)	0.139 (0.080)	0.038 (0.028)	0.706* (0.191)	0.687* (0.192)
IMPPEN	0.062* (0.027)	0.110* (0.034)	0.030 (0.080)	0.069* (0.024)	-0.156 (0.170)	-0.131 (0.136)
EXPINT	0.011 (0.055)	0.037 (0.072)	0.115 (0.134)	0.078 (0.056)	-0.108 (0.289)	-0.157 (0.244)
SURVRT	-0.094* (0.026)	-0.093* (0.034)	0.051 (0.058)	-0.027 (0.021)	0.232 (0.167)	-0.014 (0.156)
R ²	0.154	0.147	0.140	0.192	0.162	0.183

Note: The table reports the results obtained when the industry's plant survival rate (the fraction of industry plants surviving a five-year span) is included in the regression. TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.

Table 9: Regression Results—Capital Measurement Robustness Checks

	Productivity Dispersion Regressions				Central Tendency Regressions			
	ASM Plants Only		Variable Utilization		ASM Plants Only		Variable Utilization	
	TFP1	TFP2	TFP1	TFP2	TFP1	TFP2	TFP1	TFP2
VALUELB	-0.021*	-7.87E-3*	-0.020*	-4.63E-3*	0.084*	0.088*	0.047*	0.049*
	0.005	(2.05E-3)	(0.004)	(2.13E-3)	(0.012)	(0.012)	(0.011)	(0.010)
DIVINDX	0.140	0.059	0.188*	0.047	-0.490*	-0.527*	-0.297	-0.354*
	(0.098)	(0.037)	(0.090)	(0.037)	(0.216)	(0.209)	(0.191)	(0.179)
ADV	0.143	0.595*	0.063	0.486*	-0.038	-0.639	0.251	0.003
	(0.244)	(0.142)	(0.235)	(0.120)	(0.650)	(0.625)	(0.522)	(0.524)
SUNKCOST	1.133	0.291	1.628*	0.338	0.540	-0.257	0.662	0.058
	(0.663)	(0.201)	(0.610)	(0.205)	(0.709)	(0.620)	(0.704)	(0.649)
FIXEDCOST	0.112	0.008	0.109	0.013	0.798*	0.749*	0.718*	0.701*
	(0.081)	(0.035)	(0.079)	(0.036)	(0.194)	(0.199)	(0.163)	(0.172)
IMPPEN	0.059	0.104*	0.034	0.086*	-0.269	-0.134	-0.108	-0.043
	(0.090)	(0.030)	(0.075)	(0.026)	(0.177)	(0.147)	(0.132)	(0.108)
EXPINT	0.201	0.065	0.115	0.057	-0.043	-0.076	0.011	-0.094
	(0.163)	(0.066)	(0.151)	(0.066)	(0.286)	(0.272)	(0.244)	(0.207)
R ²	0.081	0.127	0.107	0.089	0.176	0.199	0.113	0.128

Note: The table reports the results obtained when plant-level TFP measures are adjusted to account for possible capital measurement error. One set of results computes productivity moments using only plants in the Annual Survey of Manufactures panel (for which capital stocks are not imputed). The other implements a correction for variable capital utilization rates across plants within the industry. TFP1 is computed using plant-specific input elasticities; TFP2 uses industry average elasticities. Heteroskedasticity-robust standard errors are in parentheses. An asterisk indicates significance at the 5% level.