Web Appendix for The Micro Elasticity of Substitution and Non-Neutral Technology

Devesh Raval

Federal Trade Commission

devesh.raval@gmail.com

March 22, 2015

1 Data Notes

In this section, I go over how the data was constructed for this paper, including the details of the construction of the main dataset for the paper, the local area wage series, various instruments for the local area wage, and measures of perpetual inventory capital. This section also details how I construct quantity produced for a set of homogeneous product industries.

1.1 Main Dataset

The main datasets that I use in this paper are the 1987 and 1997 Census of Manufactures. I remove all Administrative Record plants because these plants do not have data on output or capital. I also eliminate a set of outliers and missing values from the dataset. I first remove all plants born in the given Census year, as well as a small set of plants with missing age data. I then remove plants with zero or missing data on the following variables: average revenue product of capital, average revenue product of labor, capital share, capital labor ratio, and plant level wage. I also remove plants above the 99.5th percentile or below the 0.5th percentile of their 4 digit SIC industry on these variables to remove plants with potential data problems. I have examined how robust my results are to these outlier corrections, and have found similar estimates of the elasticity of substitution when I include the omitted plants in the dataset.

The most important variable in this study is the factor cost ratio, which is the ratio of capital costs to labor costs. I construct both capital costs and labor costs in nominal terms for the given Census year. For labor costs, I use the total salaries and wages paid by the plant.

For capital costs, I multiply capital stock measures by rental rates of capital. In the 1987 Census, the Census asked plants to report the book value of structures capital separately from equipment capital. Thus, I construct the capital stock for structures capital separately from equipment capital for 1987. Because the book value reported in the Census is a historical gross cost measure (although it accounts for capital retirements), I multiply the book value of capital by a current net cost to historical gross cost deflator based upon estimates of the current net value of capital and historic gross value of capital constructed by the Bureau of Economic Analysis at the 2 digit SIC level. Because this deflator is not base 1987, I then use investment deflators to convert each capital stock to 1987 dollars. Finally, I use a set of unpublished rental rates of capital created by the Bureau of Labor Statistics to measure total equipment costs and structures costs for the plant. These rental rates are based upon the ratio of capital income to capital stock.

In the 1997 Census, the Census only asked plants to report the total value of capital. Here, I construct a capital deflator and rental rate for both structures and equipment capital as in 1987, although I use the investment deflator to convert capital to 1997 dollars. I then average both the capital deflator and rental rate of capital for structures and equipment capital, weighting each type of capital by its share of overall capital based upon data on structures and equipment capital for the plant's 4 digit SIC industry from the NBER Productivity Database.

1.2 Local Wages

I construct measures of the local wage in order to estimate the elasticity of substitution across plants, using two different datasets to measure the local area wage. The first dataset that I use in the Census 5% samples of Americans. The Population Censuses have data on both wages and MSA geographic location for a large sample of workers.

To obtain the local wage, I first calculate the individual wage for prime age men (with age between 25 and 55) who are employed in the private sector as workers earning a wage or salary. I calculate the wage as an hourly wage, which I define as total yearly wage and salary income divided by total hours worked. I measure total hours worked as weeks worked per year multiplied by hours worked per week. I remove all individuals with zero or missing income or zero total hours worked. For 1990, incomes above the Census top code of $140,000 are set to the state median of wage and salary income above the top code. For 2000, incomes above the Census top code of $175,000 are set to the state mean of wage and salary income above the top code.

Before calculating local area wages, I adjust measures of local wages for differences in worker characteristics through regressions with the individual log wage as a dependent variable. I include education through a set of dummy variables based upon the worker's maximum educational attainment, which include four categories: college, some college, high school degree, and high school dropouts. I define experience as the individual's age minus an initial age of working that depends upon their education status, and include a quartic in experience in the regression. I also have data on the race of workers and so include three race categories of white, black, and other. I include six occupational categories: Managerial and Professional; Technical, Sales, and Administrative; Service, Farming, Forestry, and

Fishing; Precision Production, Craft, and Repairers; and Operatives and Laborers. Finallly, I include thirteen industrial categories: Agriculture, Forestry, and Fisheries; Mining; Construction; Manufacturing; Transportation, Communications and Other Public Utilities; Wholesale Trade; Retail Trade; Finance, Insurance, and Real Estate; Business and Retail Services; Personal Services; Entertainment and Recreation Services; Professional and Related Services; and Public Administration. I then calculate the local area wage as the MSA average of residual wages from a regression that includes all of these characteristics, where I allow all regression coefficients to vary by year. Because the Economic Census is conducted in different years from the Population Censuses, I match the 1987 Census of Manufactures to wages from the 1990 Population Census, and the 1997 Census of Manufactures to wages from the 2000 Population Census.

The second dataset that I use is the Longitudinal Business Database, which contains data on payroll and employment for all US establishments. I construct the establishment wage as total payroll divided by total employment. I measure the local wage as the mean log wage at the county level, though I have examined robustness to using the median log wage instead. I match the 1987 Longitudinal Business Database to the 1987 Census of Manufactures and the 1997 Longitudinal Business Database to the 1997 Census of Manufactures. Because these wages are constructed from establishment data, I cannot make adjustments for differences in workers within or across establishments.

1.3 Instruments

I use two different sets of instruments for the local wage for robustness checks on my estimates of the elasticity of substitution. The first instrument is an instrument for labor demand based upon the differential impact of national level shocks to industry employment across locations. Positive national shocks to an industry should increase labor demand, and so wages, more in areas with high concentrations of that industry. Formally, the predicted growth rate in employment for a given location is the sum across industries of the local employment share of this industry multiplied by the 10 year change in national level employment for that industry.

The implicit assumption here is that changes in industry shares at the national level are independent of local manufacturing plant productivity. To help ensure that this assumption holds, I exclude manufacturing industries from the labor demand instrument. I calculate the instrument defining locations by MSAs and industries at the SIC 4 digit level. I then match each instrument to the establishment based wage for that year in my instrumental variable regressions. For 1987, I use the 1976 to 1986 instrument because the SIC 4 digit industry definitions change significantly in 1987.

The second set of instruments I use are local housing prices and rents. Local housing costs are a simple proxy for employee cost of living, and the nominal wages that plants pay their workers should increase with local cost of living to keep workers indifferent across locations, because real wages remain constant. I use the Census 5 percent samples to construct measures of housing rents and housing prices across MSAs. Housing rent is measured as monthly gross rent for households that rent, and housing price as the current value of the house for households that own their house.

I then run a set of regressions on the log housing rent or log housing price that strip out a set of relevant housing characteristics to capture local area differences in housing costs separately from differences in housing quality. I drop all observations where the household moved in to the house more than 10 years ago, as information on the current value of the house may be inaccurate for households that moved more than ten years ago. I drop all observations where the structure age, number of rooms, or number of bathrooms is not recorded. I then include a set of indicators for the number/type of units in the structure, indicators for the presence of a kitchen and for plumbing, an indicator on whether the household is black, an indicator on whether the housing unit is a condo, the number of people in the household per room, number of rooms, number of bedrooms, a set of indicators for the year the structure was built, and an indicator on whether the structure has any commercial use. All variables are interacted with the period since the household moved in to the house, where the period is defined as 0-1 years ago, 2-5 years ago, or 6-10 years ago. Coefficients on all variables are allowed to vary by Census year. I then construct the local housing price or housing rent as the average of the residual housing price or rent within the MSA, and match these to the corresponding local wages for the same Census year.

1.4 Perpetual Inventory Capital

The Annual Survey of Manufactures tracks about 50,000 plants over five year panel rotations that are more heavily weighted towards large plants. The ASM has data on plant investment over time as well as book values of the stock of capital, which I use to construct perpetual inventory measures of capital.

I also take into account retirements of the capital stock, as data on retirements of capital stock are available from 1977-1987 excepting 1986. For 1973-1976 and 1986 I can calculate an imputed value for retirements as end of year capital subtracted from beginning year capital and yearly investment; I lower investment if this value is negative. Plants retire their capital stock at a rate of about 4 percent a year, which is concentrated in a few plants retiring a lot of capital stock. Since firms retiring capital deduct the retirement values from their book value, the book value incorporates depreciation from retirements.

I calculate perpetual inventory measures of capital through the following capital accumulation equation, as in Caballero et al. (1995):

$$K_t = (1 - \delta^a)K_{t-1} + I_t - R_t$$

where K_t is period t capital stock, I_t is period t investment, R_t is period t retirements, and δ^a is the in use depreciation rate. I build separate capital stocks for structures and equipment capital. To calculate the in use depreciation rate δ^a, I first calculate δ^r the average yearly rate of capital retirements (total retired capital stock divided by beginning gross capital stock) across plants from 1977 to 1985 by 2 digit SIC industry. I then initially define the in use depreciation rate as:

$$\delta^a = \delta - \delta^r$$

where δ is the overall 2 digit SIC depreciation rate calculated by the BLS minus this yearly retirement rate.

I account for retirements by building a set of capital vintages for each year that the plant exists in the dataset. Retirements are taken out of the gross capital stock of the earliest vintages of capital, as I assume FIFO retirement of capital. I initialize capital stock by the initial sample year book value, so for the first year that the plant exists in the dataset, capital is set to book value of capital. I deflate this book value by a net current cost to gross historical cost deflator. In subsequent years, each vintage is investment deflated through the investment deflator. Real investment is added to capital, and in use depreciation subtracted from capital. After this process, I recalculate the retirement depreciation rate as capital retired net of in use depreciation divided by net overall capital stock, and then recalculate all of the capital vintages to construct an overall capital measure.

The ASM plant samples also have data on the value of non monetary compensation given to employees, such as health care or retirement benefits, which I use to better measure payments to labor.

1.5 Homogenous Product Industries

I follow a similar process to Foster et al. (2008) in constructing data on homogenous product industries. I use eleven homogenous products: Boxes, Bread, Carbon Black, Coffee, Concrete, Flooring, Gasoline, Block Ice, Processed Ice, Plywood, and Sugar. All of the products are defined as in Foster et al. (2008). I use data from 1987-1997 as capital data was imputed before 1987 for non-ASM plants. I do not use data for 1992 for Processed Ice because of data errors, 1987 for Boxes because of a product definition change, and 1997 for Concrete because quantity data was not recorded. I remove Census balancing codes imputed by the Census to make product level data add up to overall revenue data in cases where I can identify them, receipts for contract work, miscellaneous receipts, and resales of products, and products with negative values.

I then remove all plants for which the product's share of plant revenue, measured after removing the balancing codes and other items mentioned above, is less than 50 percent. For each product, I have measures of both total quantity produced and revenue, which allows me to calculate product price as revenue over quantity. I delete all plants for which the ratio of product price to median product price is between .999 and 1.001, as these plants likely have quantity data imputed by the Census. I also remove plants with prices greater than ten times the median price or less than one-tenth the median price as potential mismeasured outliers.

2 Alternative Explanations

In this section, I examine a number of alternative models, including a gross output instead of value added production framework, a non competitive market for labor through unions, and measurement errors in capital.

2.1 Additional Factors of Production

This paper has assumed a value added production function with capital and labor as factors of production. However, Basu and Fernald (1997) point out that a value added production function requires either perfect competition or that materials are Leontief with capital and labor. The results of this paper rely upon the separability of materials from the capital-labor aggregate. For example, let $F(AK, BL)$ represent the production function for the capital-labor aggregate. A separable gross output production function is then:

$$Y = G(F(AK, BL), h(M))$$

The quantity or price of materials affects only the levels of capital and labor but not the factor cost ratio. Thus, the cost minimization conditions imply the same estimation procedure for the elasticity of substitution, and the same formula for the ratio of labor augmenting productivity to capital augmenting productivity B/A. Given expressions for the

G and h functions, we can invert the above function for $F(AK, BL)$ and so solve for capital and labor augmenting productivity. For example, a Cobb Douglas production function between the capital-labor aggregate $F(AK, BL)$ and materials means that:

$$\begin{aligned} Y &= (F(AK, BL))^{1-\alpha_m} M^{\alpha_m} \\ F(AK, BL) &= (\frac{Y}{M^{\alpha_m}})^{\frac{1}{1-\alpha_m}} \end{aligned}$$

I have examined the productivity correlations with size adjusting productivity using the above gross output Cobb Douglas production function and found similar results.

If the production function is not separable between materials and the capital-labor aggregate, the price of materials could affect the ratio of capital costs to labor costs. If plants face different materials prices, the factor cost ratio will be correlated with the level of materials as plants adjust the levels of all factors.

I thus control for the plant's share of materials in total costs in some of my main regressions. Table I reports estimates of the elasticity of substitution after controlling for materials intensity. The estimates of the elasticity of substitution are not affected by the materials controls.

Table I Further Robustness Checks for Plant Capital-Labor Substitution Elasticity

	Materials Cost Share	State FE	MSA Union Intensity
1987	0.64 *(0.02)*	0.48 *(0.05)*	0.52 *(0.04)*
1997	0.66 *(0.01)*	0.46 *(0.04)*	0.41 *(0.04)*

Note: All regressions include industry dummies, age fixed effects, and a multiunit status indicator and have standard errors clustered at the two digit industry-area level. Wages are based on establishment data in the first column, and on worker data in the second and third columns, and are defined in the text.

2.2 Unionization

So far, I have assumed that manufacturing plants face competitive input markets for capital and labor. Unions violate this assumption by bargaining collectively with management over worker pay, benefits, and duties. In the years that I study, union strength in manufacturing had already declined considerably from its peak. 25 percent of manufacturing workers were covered by a union in 1987, far below the 37 percent covered just ten years earlier in 1977. By 1997, only 17 percent of manufacturing workers were covered by a union. Since only a minority of manufacturing workers are covered by unions, union biases may be relatively unimportant.

A union affects the simple cost minimization conditions in a couple of different ways. First, a powerful union could force the plant to pay workers a premium wage over the local area wage. The union premium u_p appears in the cost minimization conditions as follows:

$$\log(rK/wL) \quad = \quad -(1-\sigma)\log(w/r) - (1-\sigma)\log(u_p) + (1-\sigma)\log B/A + \sigma \log \frac{\alpha}{1-\alpha}$$

A union premium only affects the elasticity of substitution if plant unionization varies with the wage of a plant's location. If high wage areas have more union plants, the estimate of the elasticity of substitution is biased downwards. Unions may also affect the level of productivity that the plant chooses to have by restricting management's powers to introduce new labor augmenting technologies such as automation.

Evidence in the labor and IO literature on the importance of unions is mixed. Many studies, including Hirsch (2008), have found substantial union wage premia from worker data (in Hirsch's case, from the Current Population Survey (CPS)). Schmitz (2005) and Dunne et al. (2010) also provide evidence from the cement industry and mining industry that union power lowered productivity by forcing management to adopt less efficient work practices. On the other hand, DiNardo and Lee (2004) examine manufacturing plants where a union narrowly won or lost an election and find no differences in wages or productivity between union and non union plants.

One simple check on the prevalence of unions is to examine the elasticity of substitution using within state variation in wages. Within state variation in wages removes any differences in union intensity from state level regulations such as right to work laws.[1] The within state estimates reported in the second column of Table I are only slightly different than the estimates allowing across state variation; the elasticity of substitution for 1987 falls slightly from 0.52 to 0.48.

[1] Holmes (1998), for example, shows that plants do indeed respond to right to work laws, as industrial activity is higher than average in areas in right to work states adjacent to non right to work states.

Ideally, I would also control for plant level union status in the elasticity of substitution regressions. Since I do not have data on plant level union status, I control instead for the MSA union intensity calculated by Hirsch and Macpherson (2003) from CPS data in the third column of Table I. The estimates of the elasticity of substitution are 0.52 in 1987 and 0.41 in 1997 controlling for MSA level union intensity. The 1987 estimates are unchanged from before and the 1997 estimates fall. Thus, correcting for unionization does not substantially affect the conclusions of this paper.

2.3 Measurement Errors in Capital

An alternative model is that persistent measurement errors in capital cause the wide variation in factor cost ratios across plants. In all of my empirical work, I keep the factor cost ratio or productivity as the dependent variable to prevent biases from measurement error.

I examine the salience of measurement error using data from the Annual Survey of Manufactures, which tracks about 50,000 plants over five year panel rotations. These plants generally have more accurate data, both because they have participated in the plant surveys for multiple years and because they have the investment history required to construct perpetual inventory measures of capital. I create perpetual inventory measures of capital for 1987 by initializing capital stock by the initial sample year book value, adding investment and subtracting both capital retirements and in use depreciation over time in a process similar to Caballero et al. (1995). I detail this procedure in Section 1.4. The ASM plant samples also have data on the value of non monetary compensation given to employees, such as health care or retirement benefits, which I use to better measure payments to labor.

The factor cost ratio using the perpetual inventory capital stock is highly correlated with the book value measure, with a correlation of 0.85 after taking out industry fixed effects. If I switch my measure of capital to the perpetual inventory measure of capital, I estimate an elasticity of substitution of 0.47 using quality adjusted worker wages, close to the 1987 estimate of 0.52 using book value measures of capital for the full Census.

3 Stylized Facts

The main paper estimated capital augmenting and labor augmenting productivity and found evidence for labor augmenting productivity. This section provides an additional set of stylized facts against neutral technology; the plant ratio of capital costs to labor costs, or factor cost ratio, should not vary with productivity when technical differences are neutral.

3.1 Persistent Within Industry Variation

Figure 1 depicts the smoothed density for the capital share for the ready mixed concrete industry in 1987. The mode of the capital share distribution for ready mixed concrete is slightly above 0.2. However, many plants have capital shares below 0.1 or above 0.3, and a long tail of plants have even higher capital shares. Neutral technological differences should

Figure 1 Capital Share Dispersion for Ready Mixed Concrete in 1987

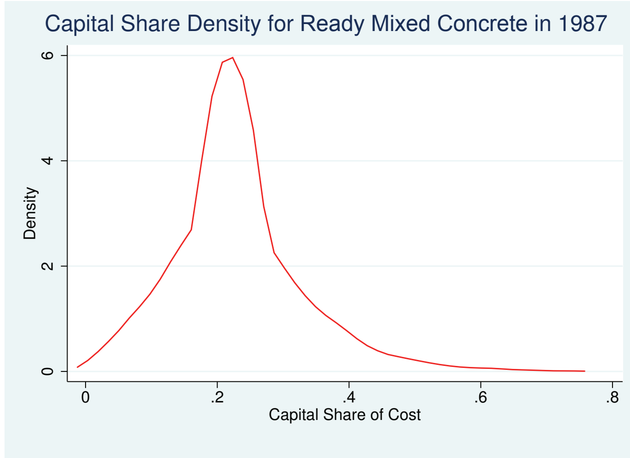

not cause any differences in capital shares across plants.

This within industry dispersion in plant capital shares exists across manufacturing industries. I measure the magnitude of this dispersion by the 75/25 and 90/10 ratios of the capital share distribution. Table II reports these statistics for the median, 25th percentile, and 75th percentile industry across all 459 manufacturing industries in the 1987 Census. For the median industry, the capital share for the 75th percentile plant is almost double that of the 25th percentile plant; the 90th percentile plant has a capital share almost four times that of the 10th percentile plant. Moreover, the 75/25 ratio and 90/10 ratios of the capital share vary only slightly between the 25th percentile industry and 75th percentile industry. For example, the 75/25 ratio for the capital share is 1.6 for the 25th percentile industry, 1.8 for the median industry, and 2.1 for the 75th percentile industry. This variation is similar for the factor cost ratio; from now on I only report statistics for the factor cost ratio.

Within industry differences in the factor cost ratio are persistent across time. I examine

Table II Dispersion in the Capital Share and Factor Cost Ratio in 1987

		Industry Percentile		
Variable	Statistic	Median	25th	75th
Capital Share	75/25 Ratio	1.8	1.6	2.1
	90/10 Ratio	3.9	3.2	4.6
Factor Cost Ratio	75/25 Ratio	2.1	1.9	2.4
	90/10 Ratio	5.4	4.6	6.7

Note: The table contains the 75/25 ratio and 90/10 ratio of each variable for the median industry, 25th percentile industry, and 75th percentile industry in the 1987 Census.

persistence to account for factors that cause temporary variation in capital shares, including idiosyncratic measurement errors and factor adjustment costs.[2] Table III contains estimates of the 10 year autocorrelation coefficient for the factor cost ratio between the 1987 and 1997 Census after controlling for industry fixed effects. The factor cost ratio is substantially autocorrelated over time with a coefficient of 0.32 over ten years. The implied one year autocorrelation coefficient given an AR(1) model of persistence is 0.89.[3]

I also examine the same 10 year autocorrelation using value added weights to measure the autocorrelation of the largest manufacturing plants. The factor cost ratio is even more persistent for the largest manufacturing plants, with a ten year correlation of 0.37 and an implied AR(1) one year autocorrelation of 0.91. The factor cost ratio has the same order of magnitude of persistence as revenue TFP, which is well known to be highly persistent (Bartlesman and Doms (2000)). Neutral productivity differences would imply that the persistence of productivity and the factor cost ratio are due to independent processes.

[2]While factor adjustment costs would lead to temporary persistence, the level of persistence observed is not consistent with standard models of adjustment costs.

[3]Under an AR(1) model the one year coefficient is the ten year coefficient to the power $\frac{1}{10}$.

Table III Persistence in Factor Cost Ratio between 1987 and 1997

	Ten Year Persistence	
Log(Factor Cost Ratio)	0.32 *(0.004)*	0.37 *(0.003)*
Log(Revenue TFP)	0.27 *(0.003)*	0.39 *(0.003)*
Weights	No	Value Added

Note: All regressions control for four digit SIC industry. Log TFP is measured as log value added minus log capital and log number of employees, each weighted by four digit industry level cost shares.

3.2 Correlation with Size

Large manufacturing plants have higher capital shares than the norm for their industry. The first graph in Figure 2 displays the nonparametric relationship between the plant factor cost ratio and value added in 1987. Each variable is calculated relative to its industry mean in order to control for industry effects. The largest plants of the industry have almost a 50 percent higher factor cost ratio than the smallest plants. The factor cost ratio is positively correlated with value added across the value added distribution, except for a slight dip for the smallest plants in the industry. This dip may be due to mismeasurement of capital utilization for small plants. Accounting for utilization would lower the factor cost ratio for low output firms and raise the factor cost ratio for high output firms, and so bolster my findings of a correlation between the factor cost ratio and value added.

The positive correlation between the factor cost ratio and value added implies that the average products of capital and labor do not move together with plant value added. The second graph in Figure 2 depicts the nonparametric relationship of each average revenue product with value added. Labor is measured through the wage bill. The average revenue

Figure 2 Factor Cost Ratio and Average Factor Products by Value Added for 1987

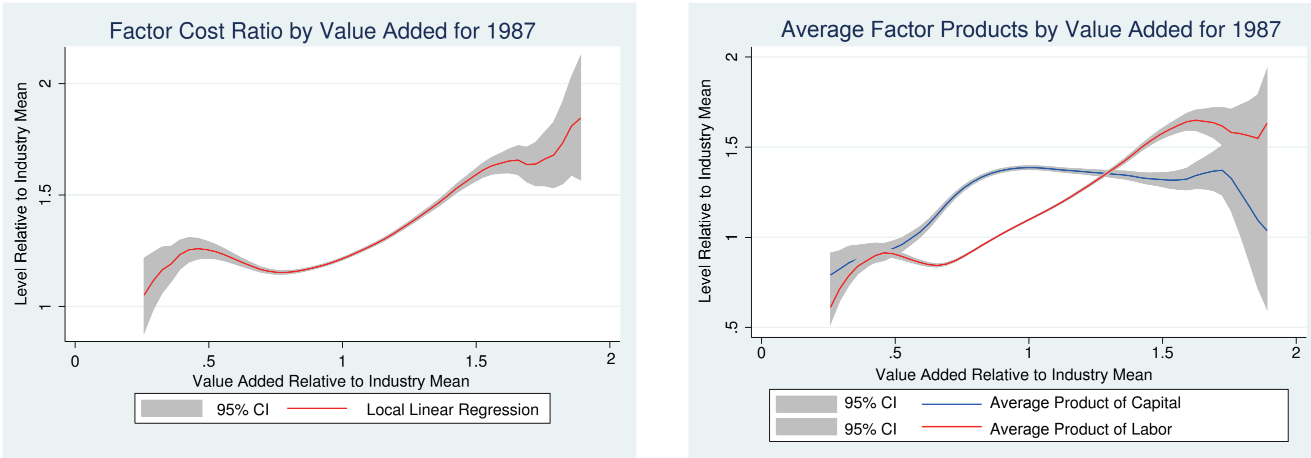

Note: Each graph depicts a local polynomial regression of either the factor cost ratio or average product of capital and average product of labor on plant value added, after adjusting all variables for industry effects by dividing by the industry mean.

product of capital increases by about 40 percent but levels off after the smallest plants in the industry. The average revenue product of labor increases by about 100 percent and continues to rise after the average product of capital flattens. Neutral productivity differences cannot explain the lack of co-movement between these average products.

The same basic relationships with plant value added hold with controls for plant level age through a set of dummy variables, plant single establishment status, and the state in which the plant is located. Table IV reports the coefficient on log value added for regressions with this extensive set of controls; the dependent variables are the log factor cost ratio or the log average factor product. The factor cost ratio increases by an average of 6 percent in 1987 and 2 percent in 1997 with a 100 percent increase in value added. The correlation between the factor cost ratio and value added is significantly higher for the largest plants in manufacturing. After weighting for value added, the factor cost ratio increases by 9 percent in 1987 and 10 percent in 1997 with a 100 percent increase in value added. The average product of labor also always rises faster with value added than the average revenue product of capital.

Table IV Correlations with Size for Factor Cost Ratio and Average Factor Products

	1987		1997	
Factor Cost Ratio	0.06 *(0.001)*	0.09 *(0.005)*	0.02 *(0.001)*	0.10 *(0.013)*
Average Product of Capital	0.07 *(0.002)*	0.03 *(0.006)*	0.10 *(0.001)*	0.07 *(0.02)*
Average Product of Labor	0.13 *(0.001)*	0.14 *(0.001)*	0.12 *(0.004)*	0.17 *(0.009)*
Weights	No	Value Added	No	Value Added

Note: Each cell contains the coefficient from a regression with log value added as the independent variable and the left hand side variable as the dependent variable, and includes controls for dummy variables for age and state, single establishment status and four digit SIC industry. Reported standard errors are robust to arbitrary degrees of heteroskedasticity.

3.3 Alternative Production Functions

The main paper assumes a CES production function; given my estimates of an elasticity below one, plants with higher labor augmenting productivity would have a higher factor cost ratio. In this section, I discuss alternative models for the stylized facts discussed in this section.

One alternative framework for production is that plants have Cobb Douglas production functions with different factor elasticities. Different factor elasticities can explain the persistent variation in the factor cost ratio across plants within US industries. However, a model with different factor elasticities does not imply the systematic relationship between size and the factor cost ratio that I find. Second, an elasticity of substitution of one implies no relationship between the factor cost ratio and the local wage. Cobb Douglas production functions with different factor elasticities cannot rationalize the stylized facts of the US micro data.

Another alternative is to introduce heterotheticity, either through a fixed cost of labor or another production function such as the translog. Heterotheticity would allow the relative share of factors to vary with size without introducing non-neutral productivity. For example, the overhead labor model used in Bartlesman et al. (2009) has labor is used both as a fixed operating cost and as a variable input for production. In this model, the average revenue product of labor increases with size because a larger plant has a higher fraction of labor devoted to production.

An overhead labor model further implies that the slope of the average product of labor - size relationship declines with plant size as the fraction of labor used in production increases towards one. As shown in Figure 2, the slope of the relationship between the average product of labor on value added does not fall with size. The average revenue product of labor increases even faster in regressions that reflect the behavior of large plants through value added weights than for the overall sample. These facts are inconsistent with a heterothetic model with overhead labor.

References

Bartlesman, Eric and Mark Doms, "Understanding Productivity: Lessons from Longitudinal Microdata," *Journal of Economic Literature*, 2000, *38* (3).

_ **, John Haltiwanger, and Stefano Scarpetta**, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," Technical Report, NBER Working Paper No. 1549 2009.

Basu, Susanto and John Fernald, "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 1997, *105* (2).

Caballero, Ricardo, Eduardo Engel, and John Haltiwanger, "Plant Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, 1995, *2.*

DiNardo, John and David Lee, "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001," *Quarterly Journal of Economics*, 2004, *119* (4).

Dunne, Tim, Shawn Klimek, and James Schmitz, "Does Competition Spur Productivity? Evidence from U.S. Cement Manufacturing Industry in the Post WWII Period," 2010. Mimeo.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, March 2008, *98* (1), 394–425.

Hirsch, Barry T., "Sluggish Institutions in a Dynamic World: Can Unions and Industrial Competition Coexist?," *Journal of Economic Perspectives*, 2008, *22* (1).

_ **and David A. Macpherson**, "Union Membership and Coverage Database from the Current Population Survey: Note," *Industrial and Labor Relations Review*, 2003, *56* (2).

Holmes, Thomas, "The Effects of State Policies on the Location of Industry: Evidence from State Borders," *Journal of Political Economy*, August 1998, *106* (4), 667–705.

Schmitz, James, "What Determines Productivity? Lessons from the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries Following Their Early 1980s Crisis," *Journal of Political Economy*, 2005, *113* (3).