Estimation of a CES Production Function with Factor Augmenting Technology*

Devesh Raval

July 11, 2011

Abstract

Both the recent literature on production function identification and a considerable body of other empirical work assume a Cobb-Douglas production function. Using a comprehensive dataset of US manufacturing plants, I provide evidence against this assumption and present an alternative production function that better fits the data. A Cobb Douglas production function has two empirical implications that I show do not hold in the data: a constant cost share of capital and neutral productivity. First, the capital share falls as local area wages rise. Second, the capital share rises with plant revenue. Labor productivity rises with plant revenue, but capital productivity does not. I show that a CES production function with labor augmenting productivity differences and an elasticity of substitution between labor and capital less than one can account for these facts. I identify the elasticity of substitution with local area differences in wages and find an elasticity of substitution of .52 in 1987, significantly different than one. I then find that labor augmenting productivity is persistent over time and strongly correlated with plant size. The form of the production function and productivity is important for macroeconomics and growth theory. Specifying the correct form of the production function is more generally important for empirical work, as I demonstrate by applying my methodology to address questions of misallocation of capital.

*I would especially like to thank my committee members Ali Hortacsu, Sam Kortum, and Chad Syverson for all of their support and guidance on this paper. I have also benefited from conversations with Fernando Alvarez, Costas Arkolakis, Alejo Costa, Chang-Tai Hsieh, Erik Hurst, Matthias Kehrig, Steven Levitt, Asier Mariscal, Benni Moll, Emi Nakamura, Ezra Oberfield, Adi Rom, Andy Zuppann, as well as participants at the UChicago IO Lunch, Micro Lunch, Labor Working Group and Trade Working Group. I would also like to thank Frank Limehouse for all of his help at the Chicago Census Research Data Center and Randy Becker for assistance with deflators for the micro data. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

The recent approach to production function estimation and the empirical work using it has generally assumed a Cobb Douglas production function. (Olley and Pakes (1996), Levinsohn and Petrin (2003)) The Cobb-Douglas assumption places two strong restrictions on the data. First, all productivity differences are Hicks neutral, which implies that labor productivity and capital productivity (revenue per unit of capital) must move together. Second, a unitary elasticity of substitution implies that cost minimizing firms will keep a constant capital share of cost.

Using micro data on more than 180,000 manufacturing plants from the US Census of Manufacturers, I find empirical evidence against both of these implications. I first examine the dispersion in capital costs relative to labor costs in the 1987 Census, documented in Table 1. For the 75th percentile manufacturing plant, capital costs relative to labor costs are about double that of the 25th percentile plant for the median industry. In the tails of the distribution, the 90th percentile plant has capital costs relative to labor costs that are more than five times that of the 10th percentile plant. These differences in the factor cost ratio (the ratio of capital costs to labor costs) are fairly persistent over time with a within industry 10 year autocorrelation coefficient of .32. This persistence rules out measurement error or adjustment costs as the main explanation for differences in the capital share across plants.

One explanation of these facts on capital share dispersion is that manufacturing plants produce with plant specific Cobb Douglas production functions, which would allow factor shares to vary within industry. However, I document stylized facts in the data that do not match plant specific Cobb Douglas production functions either. First, the factor cost ratio falls with local area wages. Plants within the same industry in high wage areas have higher labor costs relative to capital costs than those in low wage areas. If plants just had different Cobb-Douglas share parameters, selection of plants into local areas would lead to high capital share plants operating in high wage areas, the opposite of what I find.

Second, the factor cost ratio is strongly correlated with plant revenue. In 1987, the factor cost ratio rises by 6% for a 100% rise in value added, so the largest plants have about a 45%

higher factor cost ratio than the smallest plants. Labor productivity and capital productivity move separately with plant revenue as well. Labor productivity rises with plant revenue while capital productivity is flat with plant revenue after the smallest plants. Neither of these facts is consistent with a simple Cobb Douglas story either. Given these data facts, what is the nature of production technology and technical differences?

A CES production function with labor augmenting productivity differences and an elasticity of substitution less than one can better explain these facts. I first assume that firms minimize cost given competitive factor markets. When the elasticity of substitution is less than one (so labor and capital are complements), firms substitute less than one for one towards capital when wages rise. The factor cost ratio then falls when wages rise as in the data.

Labor augmenting productivity is akin to having more effective labor for the same number of workers. Since labor and capital are complementary, firms with labor augmenting productivity have a higher marginal product of capital relative to labor when factor proportions are kept fixed. Cost minimizing firms set relative marginal products equal to the relative factor prices they face, so firms with more labor augmenting productivity increase capital relative to labor. Profit maximizing firms facing a downward sloping elastic demand curve expand output and revenue when labor augmenting productivity rises. Thus, labor augmenting technical differences will lead larger firms to have higher factor cost ratios, and have higher labor productivity relative to capital productivity, as in the US micro data.

The factor bias of productivity critically depends upon the main parameter in the production function: the long run elasticity of substitution between labor and capital. I estimate the elasticity using differences in wages across local areas. Because these wage differences are highly persistent over time, with a 10 year autocorrelation coefficient for MSA wages of .9, they can identify the long run elasticity of substitution. Plants will adjust their levels of factors to match persistent differences in the wages they face. I first estimate the elasticity using an entire manufacturing cross-section and assume that the elasticity is the same across industries. I control for variation in capital intensity across four digit SIC industries through industry fixed effects. For overall manufacturing, MSA

level wage differences imply an elasticity significantly less than one: .52 in 1987 and .46 in 1997.

I then allow the elasticity of substitution to vary across industries by estimating the elasticity of substitution separately for 2 digit manufacturing industries. I can reject an elasticity of one for 15 out of 19 two digit industries using MSA level wages in 1987. All of my estimates are less than one and most of them range between .4 and .7. I also estimate the elasticity of substitution for ten large four digit industries that have plants in most geographic areas in the US, to avoid potential selection problems of plants into local areas. The prototypical example here is ready mixed concrete: every local area with construction activity has concrete plants. I reject Cobb-Douglas for 10 out of 10 industries with county wages. The elasticity of substitution for concrete is .77 for the 1987 county level estimates.

These estimates of the elasticity of substitution then allow me to identify measures of productivity. Cost minimization implies that labor augmenting productivity is the residual from the regression of the factor cost ratio on the local area wage. Without data on prices, I cannot separately identify neutral productivity from demand shocks. I can identify a combination of the two by backing out output from revenue through the demand function. Both measures of productivity are fairly persistent over time. Labor augmenting productivity is correlated with plant revenue, though the correlation of labor augmenting productivity with revenue is somewhat lower than the sum of neutral productivity and demand shocks. However, Foster, Haltiwanger, and Syverson (2008) point out that demand shocks are an important driver of firm profitability. When I use data on homogenous products where I can separate demand shocks and neutral productivity, neutral productivity has a low or negative correlation with plant sales and quantity produced, unlike labor augmenting productivity. I also find that labor augmenting productivity is correlated with total exports conditional on exporting, unlike the neutral productivity and demand shock measure.

Since the production function is one of the basic building blocks of economic theory, my results on the production function have implications for a number of economic questions. First, labor augmenting productivity differences across firms can inform a growing literature on why productivity varies across firms. For example, Bloom and Reenen (2007) find that manufacturing plants

with better management techniques have higher productivity. Many of the specific management practices they look at, including improvements in management that monitor workers' performance and provide incentives for work, are labor augmenting, as they mean workers get more done in the hours they work.

Many results in macroeconomics and growth theory also depend upon the form of the production function. A number of recent papers in growth theory have firms operating with a production function with an elasticity of substitution less than one and labor augmenting technologies. Acemoglu (forthcoming) shows that such a production function satisfies the conditions under which labor scarcity will encourage innovation as technological improvements will be strongly labor saving. Jones (2005) builds microfoundations for the standard growth model in which firms face a menu of production functions that have a low elasticity of substitution and both capital and labor augmenting productivities. If these productivities are distributed Pareto, the economy will have a unitary elasticity of substitution and all technical growth is labor augmenting.

Unlike micro Cobb-Douglas production units, micro CES production units with non neutral productivity do not easily aggregate up to an aggregate production function. Oberfield and Raval (2011) show that the macro elasticity of substitution can be characterized from the intensive margin can be characterized as a convex combination of the micro elasticity of substitution and the elasticity of demand. The weights of this convex combination depend upon the distribution of capital shares. An additional margin of substitution across firms means that the macro elasticity will be higher than the micro elasticity. A larger elasticity of demand or more disperse distribution of non neutral productivity allows more substitution across firms and so increases the macro elasticity of substitution.

The type of production technology I find holds in the US data affects questions of misallocation as well. A number of papers have shown that output or capital frictions can cause large losses in aggregate productivity, and so explain productivity differences across countries. (Banerjee and Duflo (2005), Restuccia and Rogerson (2008)) Hsieh and Klenow (2009), for example, find eliminating frictions could increase aggregate TFP over 30% for the US and over 100% for India and

China. However, these studies have assumed micro production units produce using identical Cobb Douglas production functions.

Changing this assumption affects results on misallocation in two ways. First, a low elasticity of substitution makes it more difficult for the firm to substitute away from a factor facing frictions and so increases the cost of these frictions. Second, labor augmenting productivity differences lead to differences in cost and revenue shares of factors within an industry. A Cobb Douglas production function would imply that all of these differences are misallocation frictions. Thus, we have to consider differences in labor augmenting productivity to correctly estimate the contribution of misallocation to aggregate productivity.

Section 2 goes over stylized facts on capital share differences, Section 3 a model of the firm's production problem, Section 4 the estimation of the elasticity of substitution and productivity and Section 5 some departures from the assumptions of the model. Section 6 contains an application to misallocation and Section 7 concludes.

1.1 Literature Review

My work is related to both a literature on productivity estimation and a literature estimating the labor capital elasticity of substitution. The recent IO literature on productivity estimation since Olley and Pakes (1996) has focused on estimators for the Cobb-Douglas production function that avoid endogeneity problems. Unobserved productivity shocks will affect both output and factor levels since firms choose their levels of output and inputs jointly.

Gandhi, Navarro, and Rivers (2009) is the only paper to apply the methodology stemming from Olley and Pakes to a wider array of production functions. They use revenue share equations from the first order conditions of the production function to separate measurement error from Hicks neutral productivity as neutral productivity does not enter their share equations. They then estimate the production function parameters with timing moments that assume that innovations to productivity are uncorrelated with prior levels of inputs. The revenue share equations allow them

to estimate more broad production functions than Cobb Douglas, such as the CES or translog. In their CES estimation case, they find an elasticity considerably above 1 which is very different from my results.

However, their assumption of neutral productivity implies that any non-neutral productivity would become measurement error. Non-neutral productivity removes the separation between measurement error and productivity that they exploit and leads to severe biases in their approach.

A broader literature in macroeconomics and labor economics focuses on estimating the elasticity of substitution between labor and capital, using a variety of different techniques. Most of these papers estimate the macro elasticity: the elasticity of substitution for the aggregate production functions. The macro elasticity of substitution is not the same parameter as the micro elasticity of substitution. Since at the macro level substitution is possible both within firms and across firms, the macro elasticity of substitution will be higher than the micro elasticity of substitution.

The standard approach to estimating the elasticity of substitution has been to use changes over time in factor prices at the macro level or panel changes in the user cost of capital at the micro level. However, this time series source of variation in subject to two major biases. First, a number of papers have recently shown that labor augmenting technical change can bias estimates of the elasticity of substitution towards one. Both Antras (2004) and Klump, McAdam, and Willman (2007) control for labor augmenting technical change, either through time trends that imply exponential growth in labor augmenting productivity or other parametric functional forms. They find that estimates of the macro elasticity fall from one to .8 and .6 for the US, respectively. Because I use cross-sectional variation to identify the elasticity of substitution, my approach is not subject to the same bias from non neutral technological progress.

Second, studies using shocks to the rental rate of capital may not estimate the long run micro elasticity of substitution because firms face adjustment costs. Since adjustment costs mean that capital cannot adjust instantaneously, the short run elasticity of substitution can be much lower than the long run elasticity of substitution. **?** finds that adjustments for adjustment costs through lag values of variables can have a large impact on estimated elasticities. Firms will also not respond

much to transitory changes in the rental rate of capital. In this paper I identify the elasticity with differences in local area wages that are very persistent over time, which eliminates these problems.

Because of the above biases to time series or panel estimation, estimates of the elasticity of substitution vary considerably, depending upon the panel data set and type of user cost variation used. Chirinko (2008) and León-Ledesma, McAdam, and Willman (2010) both provide surveys of the previous literature. Depending upon the type of technical change and country studied, researchers have found the macro elasticity to be below, equal, or greater than one. Looking across 2 digit SIC industries,**?** find estimates of the micro elasticity ranging from 0 to 2. Chirinko, Fazzari, and Meyer (2004) attempt to avoid these problems by using only long run changes in the user cost of capital. They compare changes in the firm's average capital-output ratio over two long time intervals to changes in the average user cost and control for biased technical change through time effects. They find estimates of the micro elasticity of .44 that are similar to my baseline estimates using MSA level wages.

2 Stylized Facts

The Cobb Douglas production function is

$$F(K,L) = AK^{\alpha}L^{1-\alpha}$$

This assumption places a number of restrictions on the data. First, all productivity is neutral. To show this, I introduce a labor augmenting productivity measure B to the production function. However, this labor augmenting productivity B can always be rewritten as a new Hicks neutral shifter $\tilde{A}$:

$$Y = AK^{\alpha}(BL)^{1-\alpha} = AB^{1-\alpha}K^{\alpha}L^{1-\alpha} = \tilde{A}K^{\alpha}L^{1-\alpha}$$

Thus, firm productivity does not affect the relative proportions of factors that are used by the firm and so labor productivity and capital productivity should move together. A firm facing competitive inputs and minimizing cost sets the relative marginal products of capital and labor equal to relative factor prices. For the Cobb Douglas production function, this implies that the capital share of cost is constant:

$$\begin{aligned} \frac{rK}{rK+wL} &= \alpha \\ \frac{rK}{wL} &= \frac{\alpha}{1-\alpha} \end{aligned}$$

In this section, I use the US data on manufacturing plants to test this strong implication. The ratio of capital costs to labor costs, which I call the factor cost ratio, is very disperse across plants within 4 digit industries, persistent over time, and correlated with measures of plant size. Labor productivity has a much stronger correlation with plant size than capital productivity. These stylized facts then motivate a production function where productivity is not neutral, which I describe in Section 3 and then estimate in Section 4.

2.1 Data

I use US data on manufacturing plants from the Census of Manufactures and Annual Survey of Manufactures (ASM). The Census of Manufactures is a census of all manufacturing plants taken every five years. I then drop manufacturing plants with missing or outlier data. A typical Census sample has more than 180,000 plants and considerable variation within ages, geographic areas, and industries. I always look within industry to avoid including within industry differences in capital intensity in my results.

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

The main variables I look at in this study are capital costs, labor costs, and value added at the plant level. I measure value added by undeflated value added. Labor costs are the total salaries and wages at the plant level. For the ASM plant samples, I also have data on the value of non monetary compensation given to employees, such as health care or retirement benefits. For these samples I use this benefits data to better measure payments for labor.

The primary data constraint for this study is data on capital. Before 1987 the Census did not ask questions on capital stocks for non ASM plants. Thus, I primarily use the samples after the Census began to measure capital data.

I measure capital by the end year book value of capital, deflated using a current cost to historic cost deflator. The 1987 Census has book values for equipment capital and structures capital separately, so I construct capital stocks for each and then combine them. For capital costs, I multiply these measures of capital stocks by the appropriate rental rate. For the rental rate, I use unpublished two digit BLS rental rates calculated by capital income to capital stock measures. However, both the capital deflator and rental rate end up captured in the industry fixed effect as they do not vary within industry. This rental rate thus assumes that the rental rate of capital is constant between plants within the same industry.

For the general Census of Manufactures I cannot use perpetual inventory methods, because investment is not recorded in non-Census years. However, for the ASM samples, I create perpetual inventory measures of capital by initializing capital stock by the initial sample year book value, adding investment and substracting capital retirements and in use depreciation over time in a process similar to **?**.

2.2 Persistent Within Industry Variation

Even within a 4 digit SIC industry, manufacturing plants have large differences in their capital shares. Figure 2.1 shows the smoothed kernel density for the capital share for the ready mixed concrete industry in 1987. The capital share for ready mixed concrete is quite disperse. While the mass of the distribution is concentrated at .2, a substantial proportion of plants have capital shares

above .3 and below .1.

Figure 2.1: Capital Share Dispersion for Ready Mixed Concrete in 1987

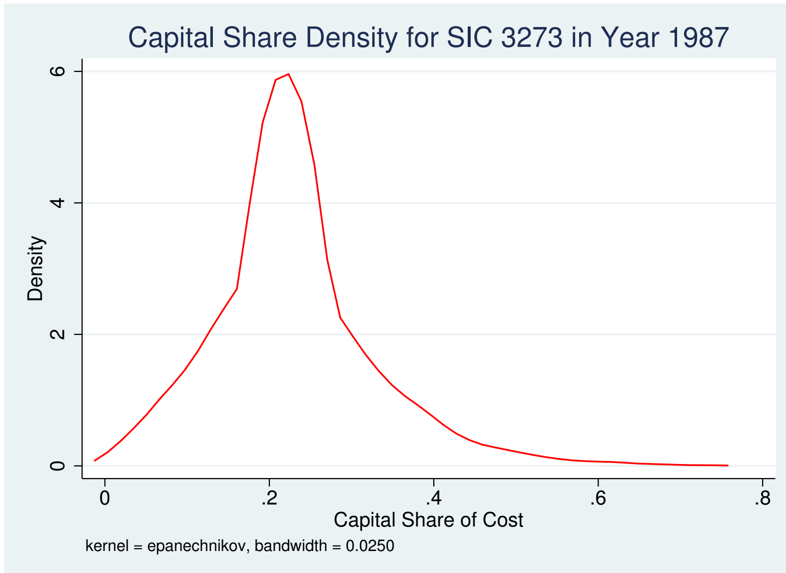

The X axis is the plant capital share. The Y axis is the density of the capital share. The graph was generated using kernel density estimation.

Most US industries have considerable variation in capital shares across manufacturing plants. Table 1 reports the 75/25 ratio and 90/10 ratio of the capital share distribution for the median, 25th percentile, and 75th percentile industry across all 459 manufacturing industries. For the median industry in the 1987 Census, 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 for the median industry. These large within industry differences exist across manufacturing industries. The 75/25 ratio and 90/10 ratios of the capital share are not too different between the 25th percentile industry and 75th percentile industry. For example, the 75/25 ratio is 1.6 for the 25th percentile industry, 1.8 for the median industry, and 2.1 for the 75th percentile industry.

I have also reported the factor cost ratio, which is the ratio of capital costs to labor costs for the plant. Since the capital share is a monotone function of the factor cost ratio, all the stylized facts for one hold for the other. From now on I only report statistics for the factor cost ratio.

Table 1: Dispersion in Capital Share and Factor Cost Ratio within 4 digit Industries for the 1987 Census of Manufactures

		Median	25%	75%
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

For each industry, I calculated the 75/25 ratio and 90/10 ratio for each variable. I have then reported the median, 25%, and 75% of these ratios across industries.

This within industry dispersion in the factor cost ratio is also persistent across time. Measurement errors in capital or adjustment costs will cause temporary differences in capital shares. Idiosyncratic measurement error will cause temporary dispersion in capital shares. Firms that face adjustment costs will not always adjust their capital share to match the static first order conditions. Instead, they will only adjust capital infrequently when their capital stock gets too far away from the optimal capital stock. However, most firms will have readjusted their capital stock over a ten year time horizon, so adjustment costs would not cause long run persistence in the factor cost ratio. Adjustment costs of capital would imply that firms with high capital shares today would be lower in the future as they readjusted their capital share downwards.

Figure 2.2 displays the local polynomial regression of the 1997 factor cost ratio relative to its industry mean against the 1987 value. Plants whose capital share is high in 1987 also have a high capital share in 1997. Only for plants with an extremely high level of capital costs to labor costs, more than four times the industry mean, does this relationship level off.

Figure 2.2: Persistence of Factor Cost Ratio between 1987 and 1997

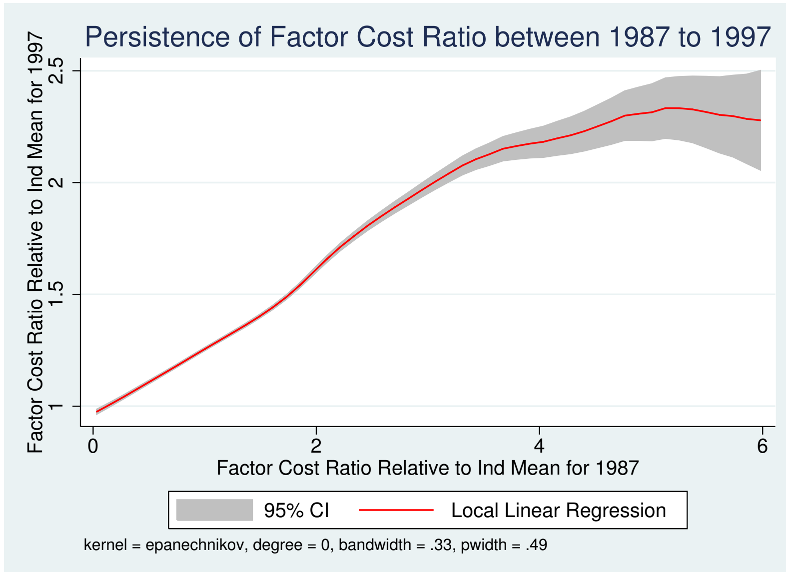

The X axis is a plant's factor cost ratio in 1987 relative to its industry mean. The Y axis is the plant's factor cost ratio in 1997 relative to its industry mean. The graph was generated using local polynomial regression.

I also compute the 10 year autocorrelation in the factor cost ratio using the 1987 and 1997 Censuses after controlling for industry fixed effects. I compare these values with the autocorrelation values for conventionally measured log TFP, as TFP is well known to be autocorrelated over time. Table 2 contains the estimates.

The factor cost ratio is substantially autocorrelated over time. A 10 year coefficient of .32 implies a one year autocorrelation of .89. I also run weighted regressions with value added weights which measure the autocorrelation of the biggest plants in the industry. In these weighed regressions, a ten year correlation of .37 implies an one year autocorrelation of .91. These levels of autocorrelation are quite similar to TFP, which has a ten year correlation of .27 in the unweighted regressions and .39 in the weighted regressions.

Another way to look at persistence is to look at transitions between percentiles in the capital share distribution. I look at a transition table between quantiles of the capital share distribution. Within industry year, I assign plants to quartiles based on their factor cost ratio and examine how much movement there is between quartiles over ten years. Table 3 reports this transition table. For

Table 2: Persistence in Factor Cost Ratio between 1987 and 1997

	Ten Year	Implied One Year	Ten Year	Implied One Year
Log(Factor Cost Ratio)	.32 *(.004)*	.89 *(.001)*	.37 *(.003)*	.91 *(.001)*
TFP	.27 *(.003)*	.88 *(.001)*	.39 *(.003)*	.91 *(.001)*
Weights	No	No	Value Added	Value Added

All regressions contain 4 digit SIC industry dummies. TFP is measured by log(Value Added) minus log capital and log labor (no of employees) weighted by 4 digit industry level cost shares. The implied one year coefficient is the ten year coefficient to the 1/10 power.

the factor cost ratio, 47.2% of the largest quartile plants in 1987 are in the largest quartile of plants in 1997. Similarly, 39% of the lowest quartile plants in 1987 were also in the lowest quartile in 1997. Thus, the differences in capital share within industry are also persistent over time.

2.3 Correlation with Size

I also find that plants with higher value added have higher capital shares. I first examine the nonparametric regression of the plant's factor cost ratio on the plant's value added in Figure 2.3. To control for industry effects, I calculate each relative to its industry mean. Here, the largest plants of the industry have a much higher factor cost ratio than the smallest plants, about 45% higher for 1987. However, the factor cost ratio does dip slightly for plants producing around half the industry average. Figure 2.4 depicts another way to see this relationship, using the within industry percentile of value added instead of the level of value added. First, the increase in the factor cost ratio is not just a difference between small and large plants in the industry. Even among the largest 20% of plants in a given industry we see a substantial increase in capital costs relative to labor costs for larger plants.

Table 3: Transition Table between Quartiles of Factor Cost Ratio from 1987 and 1997

1987 Quartile	1997 Quartile			
	Q1	Q2	Q3	Q4
Q1	**39.3%**	27%	19.4%	14.3%
Q2	22.5%	**30.3%**	23.9%	23.4%
Q3	18.9%	26.1%	**26.8%**	28.2%
Q4	12%	19.3%	21.6%	**47.2%**

Here Quartile 1 has the smallest 25% of plants by within 4 digit industry capital share and Quartile 4 the largest 25% of plants by within 4 digit industry capital share.

Figure 2.3: Factor Cost Ratio by Value Added for Year 1987

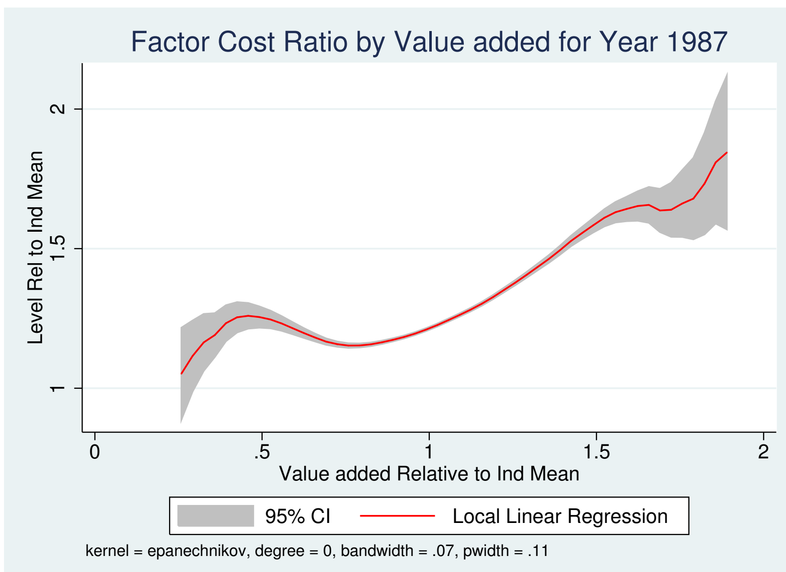

The X axis is a plant's value added relative to its industry mean. The Y axis is the factor cost ratio relative to its industry mean. The graph was generated using local polynomial regression.

Figure 2.4: Factor Cost Ratio by Value Added Percentile for Year 1987

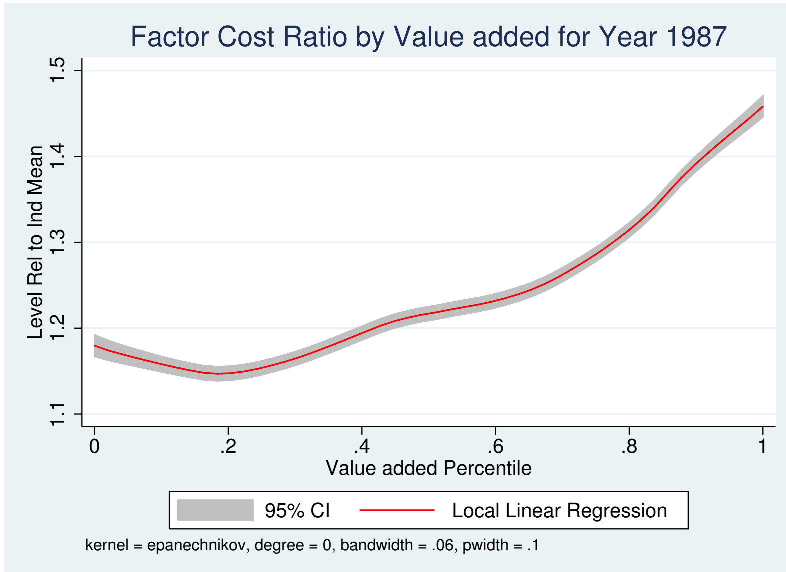

The X axis is a plant's value added percentile within the industry. The Y axis is the factor cost ratio relative to its industry mean. The graph was generated using local polynomial regression.

Second, the factor cost ratio again dips, but for only the smallest 20% of plants. This dip is slightly larger in 1997 but vanishes in 2002. The rise in the factor cost ratio after the smallest 20% of plants is always present, however. I can only measure the amount of capital owned by firms, not the amount of capital used by the firms. If a few firms shut down or produce much less than usual, they are not using much of their capital stock. These firms will have very low output but a high factor cost ratio relative to the industry. This example illustrates a general problem: capital utilization rates are not observed. Differences in utilization lower the true factor cost ratio for low output firms and raise the true factor cost ratio for high output firms, and so should only bolster my findings of rising capital costs relative to labor costs with value added.

Another way to see these facts that relates more directly to productivity is to examine the average revenue product of capital and average revenue product of labor. Hicks neutral productivity differences should affect the average revenue products of capital and labor symmetrically. I measure labor as the wage bill- labor productivity rises considerably faster if I measure labor by the number of employees. Larger plants have higher wages which may reflect better quality workers.

I then nonparametrically regress the factor cost ratio on value added in Figure 2.5, where both values are relative to the industry mean values. The average revenue product of capital increases by about 40% but levels off after the smallest plants. The average revenue product of labor increases by about 100% and continues to rise after the average product of capital is constant or slightly declining. The average product of capital and average product of labor do not move together as plant value added increase. These relationships look similar in 1997.

Figure 2.5: APK and APL by Value added for Year 1987

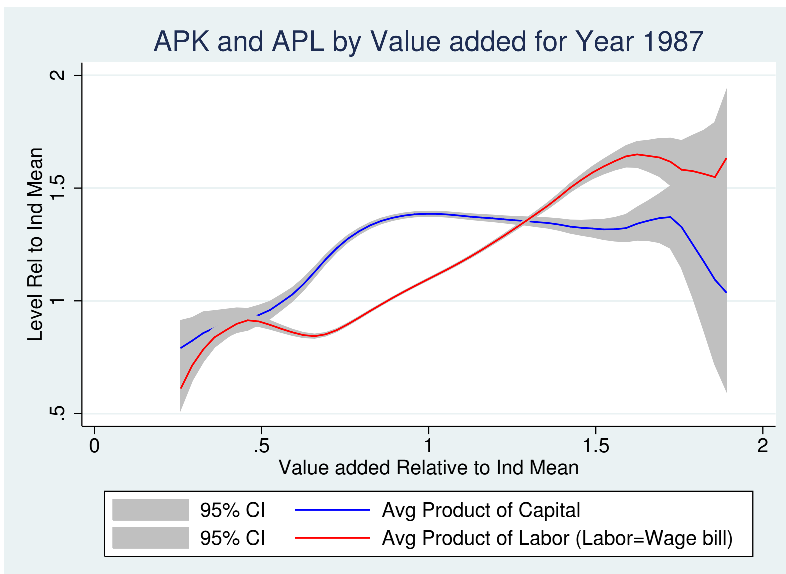

The X axis is a plant's percentile of value added relative to its industry. The Y axis is the log average product of capital or average product of labor after taking out industry averages. Here labor is defined as the plant's wage bill. The graph was generated using local polynomial regression.

The previous results only controlled for industry effects. I now run regressions that depict the same facts but control for plant level age through a set of dummy variables, the plant's single establishment status, and state, using the 1987 and 1997 Censuses. Table 4 contains regressions where each cell represents a different regression with log factor cost ratio, the log average product of capital and the log average product of labor as the dependent variables and log value added as the independent variable, as well as the controls. I also report these regressions weighting for value added, which reflect the same relationships for the largest plants in the Census. The factor cost ratio increases faster with value added in the weighted regressions than the unweighted regressions,

Table 4: Correlations with Size for Factor Cost Ratio, APK, APL

	1987		1997	
Log(Factor Cost Ratio)	.06	.09	.02	.10
	(.001)	*(.005)*	*(.001)*	*(.013)*
Log(APK)	.07	.03	.10	.07
	(.002)	*(.006)*	*(.001)*	*(.02)*
Log(APL)	.13	.14	.12	.17
	(.001)	*(.001)*	*(.004)*	*(.009)*
Weights	No	Value Added	No	Value Added

All of these coefficients are from regressions with the LHS variable as the dependent variable and log of value added as the independent variable. Controls include dummy variables for age and state, single establishment status and 4 digit SIC industry. I use robust standard errors.

consistent with the graphs above. A 100% increase in value added increases the factor cost ratio by 6% in the unweighted 1987 regressions and 9% in the weighted 1987 regressions. This increase is weaker in 1997 for the unweighted regressions, but not for the weighted regressions. A 100% increase in value added increases the factor cost ratio by 2% in the unweighted 1997 regressions and 10% in the weighted 1997 regressions. The average product of labor always rises significantly faster than the average revenue product of capital. This difference is substantially higher in the weighted regressions.

2.3.1 Capital Robustness

The above results use book values of capital for the capital stock. For the ASM sample I can calculate a perpetual inventory measure of the capital stock. Perpetual inventory measurement depreciates each vintage of capital by its age and deflates each vintage by its investment year's investment deflator. I also take into account retirements of the capital stock. Plants retire their capital stock at a rate of about 4% 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

Table 5: Capital Robustness Checks: Correlations with Size for Factor Cost Ratio, for ASM plants only

	1987	
Book Value of Capital	.06 *(.007)*	.10 *(.006)*
Perpetual Inventory Measure of Capital	.03 *(.007)*	.08 *(.006)*
Weights	ASM Weight	Value Added * ASM Weight

All of these coefficients are from regressions with the LHS variable as the dependent variable and log of value added as the independent variable. Controls include dummy variables for age and state, single establishment status and 4 digit SIC industry. I use robust standard errors. I construct the perpetual inventory stock using data on book values of capital, investment, and retirements.

book value measures incorporate the depreciation from retirements. Data on retirements of capital stock are available for all years up to 1987, after which retirements are recorded only in Census years. The factor cost ratio using the perpetual inventory capital stock is highly correlated with the book value measure, with a raw correlation of .93 and a correlation of .85 after taking out industry fixed effects. Table 5 shows the correlation between value added and the factor cost ratio using the perpetual inventory measure of capital. Here I still find the same patterns as before, though the correlations are slightly weaker for the perpetual inventory measures of capital. In the unweighted regressions, the factor cost ratio increases by 6% for a 100% increase in value added using a book value measure and by 3% using a perpetual inventory measure. In the weighted regressions, the factor cost ratio increases by 10% for a 100% increase in value added using a book value measure and by 8% using a perpetual inventory measure.

2.3.2 Firm Specific Rental Rates

I have also assumed that the rental rate of capital is constant across plants within an industry. Another potential reason for my findings is differences in rental rates of capital across plants, as

large plants may face low rental rates of capital and so have higher capital shares. However, the rental rate of capital should vary at the firm level, so I control for it using firm level measures. I have to restrict the sample to multi unit plants to control for firm level effects. I control for the rental rate of capital in two ways: using the firm's total size and using firm fixed effects. Firm size will control for the rental rate of capital if the rental rate is correlated with firm size, while firm fixed effects control for all firm level differences. These measures will control for firm level differences in productivity as well as differences in the rental rate of capital.

For firm size, I use total firm employment as a measure of size either by including deciles of the firm size distribution for manufacturing plants, or a quartic in log firm employment. I have to use total employment for firm size as I only have data on employment and payroll at the firm level. I have also varied firm size measures using payroll instead of employment and found similar results. As Table 6 shows, I still find that the factor cost ratio is correlated with plant value added. A 100% increase in value added increases the factor cost ratio by 4% in 1987 and 6% in 1997 in the unweighted regressions and by 7% in 1987 and 8% in 1997 in the weighted regressions. The last row in Table 6 contains the fixed effect results. Here the weights in the weighted regressions are at the firm level and so are the average value added for a manufacturing plant of the firm. These regressions with firm fixed effects have correlations similar to the regressions with firm size controls in the unweighted regressions and slighly smaller in the weighted regressions. A 100% increase in value added increases the factor cost ratio by 4% in 1987 and 6% in 1997 in the unweighted regressions and by 5% in 1987 and 7% in 1997 in the weighted regressions.

3 Theoretical Setup

I have shown that there are persistent within industry differences in the capital share across plants. The plant capital share is also correlated with plant value added. These facts are incompatible with a Cobb Douglas production function and neutral technology. In this section, I introduce a CES production function with non neutral technology to help explain these facts and explore

Table 6: Robustness Checks: Firm Level Controls

	1987		1997	
Firm Size=Quartic in Log Firm Employment	.04 *(.002)*	.07 *(.007)*	.06 *(.003)*	.08 *(.014)*
Firm Size= Deciles of Log Firm Employment	.04 *(.002)*	.07 *(.007)*	.06 *(.003)*	.08 *(.015)*
Firm Fixed Effects	.04 *(.003)*	.05 *(.007)*	.06 *(.005)*	.07 *(.009)*
Weights	No	Value Added	No	Value Added

All of these coefficients are from regressions with the LHS variable as the dependent variable and log of value added as the independent variable. Controls include dummy variables for age, state and 4 digit SIC industry. These regressions only include multiunit firms.

the implications of cost minimization and profit maximization. The direction of technological differences across plants critically depends upon the value of the elasticity of substitution between labor and capital.

I assume that the production function has a constant elasticity of substitution σ.[1] Here I allow both neutral and labor augmenting technologies. If productivity is Hicks neutral, improvements in productivity affect labor and capital symmetrically. If productivity is labor augmenting (Harrod neutral), an increase in productivity is equivalent to having more labor. While here I have introduced productivity as neutral or labor augmenting, I could always rewrite the production function in terms of capital augmenting and labor augmenting technologies. The CES Production Function with constant returns to scale is then:

$$Y = A(\alpha K^{\frac{\sigma-1}{\sigma}} + (1-\alpha)(BL)^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}} \tag{3.1}$$

A is Hicks neutral productivity and B is labor augmenting productivity. Y is the quantity output

[1] The elasticity of substitution depends upon the curvature of the production function, as in the definition of Hicks (1932): $\sigma = \frac{F_k F_l}{F_{kl} F}$

of the firm, not revenue. σ is the elasticity of substitution: labor and capital are complements if $\sigma < 1$ and substitutes if $\sigma > 1$. The distribution parameter α governs how much capital contributes to output relative to labor. α is not separately identified from A and B when σ is not equal to one.[2]

The elasticity of substitution determines how the ratio of marginal products depends upon productivity. Taking derivatives of the CES production function, we have that:

$$\frac{MPK}{MPL} = (B)^{\frac{1-\sigma}{\sigma}}(\frac{K}{L})^{-\frac{1}{\sigma}}\frac{\alpha}{1-\alpha} \tag{3.2}$$

Neutral productivity changes do not affect the ratio of marginal products because Hicks neutral productivity increases the marginal productivity of capital and labor by the same percentage. The elasticity of substitution determines how the ratio of marginal products responds to changes in labor augmenting productivity B. Keeping factor proportions constan, the marginal product of capital rises relative to that of labor when labor augmenting productivity increases if capital and labor are complements and falls if capital and labor are substitutes. Only for the Cobb-Douglas case does productivity not affect the ratio of marginal products.

3.1 Cost Minimization

So far I have kept factor proportions constant. A cost minimizing firm sets marginal products equal to factor prices by adjusting the levels of its factors. Cost minimization thus implies that:

$$\frac{r}{w} = \frac{MPK}{MPL} = (B)^{\frac{1-\sigma}{\sigma}}(\frac{K}{L})^{-\frac{1}{\sigma}}\frac{\alpha}{1-\alpha} \tag{3.3}$$

I invert the above equation to examine the factor cost ratio.

$$\frac{rK}{wL} = B^{1-\sigma}(\frac{r}{w})^{1-\sigma}(\frac{\alpha}{1-\alpha})^{\sigma} \tag{3.4}$$

[2] One useful property of the CES production function is that it nests a number of famous simple cases. When the elasticity of substitution σ converges to 0, we have the Leontief production function, $Y = A\min(K/\alpha, BL/(1-\alpha))$. When the elasticity of substitution σ is 1, we have the Cobb-Douglas production function $Y = AK^{\alpha}(BL)^{1-\alpha}$. When the elasticity of substitution σ converges to infinity, we have the linear production function $Y = A(\alpha K + (1-\alpha)BL)$.

First, the elasticity of substitution σ controls how the factor cost ratio reacts to changes in factor prices.[3] Wage increases reduce the factor cost ratio when $\sigma < 1$. A Cobb-Douglas production function has a constant capital share of cost, as $\frac{rK}{wL} = \frac{\alpha}{1-\alpha}$.

I can also use equation 3.4 to characterize what happens to the factor cost ratio when the level of labor augmenting productivity changes. If $\sigma < 1$, firms respond to increases in labor augmenting productivity by raising their factor cost ratio until marginal products are equal to factor prices. Thus, increases in labor augmenting productivity are also *labor saving* if $\sigma < 1$. If $\sigma > 1$, firms decrease their capital-labor ratio when labor augmenting productivity rises.

Increases in labor augmenting productivity B are equivalent to increases in labor for firms. If labor and capital are complements, firms want to increase the amount of capital they hold. The marginal product of capital rises until firms increase capital relative to labor. Since the elasticity of the factor cost ratio to changes in labor augmenting productivity B is $1-\sigma$, the factor cost ratio increases less than proportionately with B unless the production function is Leontief.

I summarize this section in the following proposition:

Proposition 1. *If the firm production function is CES with labor and capital as inputs, and firms cost minimize facing competitive factor markets, then when $\sigma < 1$:*

1. The factor cost ratio $\frac{rK}{wL}$ will decrease with increases in w/r with an elasticity of $\sigma - 1$.
2. The factor cost ratio increases with labor augmenting productivity B with an elasticity of $1-\sigma$.

3.2 Profit Maximization

So far I have not placed any assumptions on the level of output. I now introduce a demand side where each firm produces a differentiated product and has a downward sloping demand curve for

[3] Cost minimization thus implies the Robinson (1933) definition of the elasticity of substitution, $\sigma = -\frac{d\log(K/L)}{d\log(r/w)}$

their product. The demand curve is isoelastic:

$$Y = \frac{D^{\varepsilon-1}}{P^{\varepsilon}} \tag{3.5}$$

D is a demand shifter: firms with higher D can sell more of the product at the same price. ε is the elasticity of demand for the firm's product. I assume that the elasticity of demand is greater than one as is required for firms to set pricwa. I can rewrite expression 3.5 in terms of revenue, so firm revenue depends upon price in the following way:

$$PY = (D/P)^{\varepsilon-1} \tag{3.6}$$

Isoelastic demand implies that the optimal price for the firm is a constant markup over marginal cost:

$$P = \frac{\varepsilon}{\varepsilon-1}C \tag{3.7}$$

C is the marginal cost of the firm's product. Since the production function has constant returns to scale, the marginal cost of production does not depend upon the amount produced. Cost minimization implies that the firm's marginal cost is:[4]

$$C = \frac{1}{A}(\alpha^{\sigma}r^{1-\sigma} + (1-\alpha)^{\sigma}(\frac{w}{B})^{1-\sigma})^{\frac{1}{1-\sigma}} \tag{3.8}$$

Both measures of productivity reduce the firm's marginal cost. The capital distribution parameter α and labor augmenting productivity B govern how much labor contributes to the marginal cost relative to capital. I can then solve for the price substituting the marginal cost from equation 3.8

[4] The marginal cost is the Lagrange multiplier on the production function in the cost minimization problem. To obtain the marginal cost, substitute in the first order conditions for labor and capital into the production function and then solve for the Lagrange multiplier.

into the markup from equation 3.7. Equation 3.6 implies that the firm's revenue is:

$$PY = (AD)^{\varepsilon-1}(\frac{\varepsilon-1}{\varepsilon})^{\varepsilon-1}(\alpha^{\sigma}r^{1-\sigma}+(1-\alpha)^{\sigma}(\frac{w}{B})^{1-\sigma})^{-\frac{\varepsilon-1}{1-\sigma}} \tag{3.9}$$

Both the Hicks neutral productivity A and demand shifter D increase revenue symmetrically. With heterogeneous goods, A and D are isomorphic: one can not tell whether a given product is produced more efficiently with higher A or just has higher demand D.

Firms with higher labor augmenting productivity B also have higher revenue. Cost minimization then implies that firms with higher revenue should have a higher factor cost ratio. Since revenue increases proportionately with cost through the constant markup, the average revenue products of capital and labor depend upon B:

$$\frac{PY}{L} = \frac{\varepsilon}{\varepsilon-1}w((\frac{\alpha}{1-\alpha})^{\sigma}(\frac{Br}{w})^{1-\sigma}+1) \tag{3.10}$$

$$\frac{PY}{K} = \frac{\varepsilon}{\varepsilon-1}r((\frac{1-\alpha}{\alpha})^{\sigma}(\frac{w}{Br})^{1-\sigma}+1) \tag{3.11}$$

Hicks neutral productivity does not affect the average revenue products. Improvements in neutral productivity induce the firm to produce more until the marginal return of factors meets factor prices. This increase in production pushes the firm down its demand curve until its price falls and average revenue products remain constant. Labor augmenting productivity, by contrast, shifts the average revenue products of labor and capital in opposite directions. A firm with high labor augmenting productivity B increases its capital-labor ratio, depressing its average revenue product of capital and pushing up its average revenue product of labor. I summarize this section with the following proposition:

Proposition 2. *If the firm production function is constant returns to scale and CES with labor and capital as inputs, and firms profit maximize facing competitive factor markets and an isoelastic demand function, then when* $\sigma < 1$*:*

1. Revenue will increase with labor augmenting productivity B.

2. The average revenue product of labor $\frac{PY}{L}$ increases with labor augmenting productivity B and the average revenue product of capital $\frac{PY}{K}$ decreases with B.

4 Productivity Estimation

In this section I estimate the elasticity of substitution, the main parameter in the CES production function, and then obtain measures of neutral and labor augmenting productivity. I use the conditions from cost minimization and local area wage variation to identify the elasticity of substitution. With the elasticity of substitution in hand, I can form measures of productivity as residuals in the data. I describe my identification strategy in Subsection 4.1, the estimation of the elasticity of substitution in Subsection 4.2, and both the estimation of productivity and some results on the direction of productivity in Subsection 4.3.

4.1 Identification Strategy

The first order cost minimization conditions for capital and labor of the CES production function imply that:

$$\log(rk/wl) \quad = \quad -(1-\sigma)\log(w/r)+(1-\sigma)\log B+\sigma\log\frac{\alpha}{1-\alpha} \tag{4.1}$$

The elasticity of substitution determines how the firm changes its factor cost ratio when the wage to rental rate ratio changes. To identify the elasticity of substitution, I need variation in factor prices. I use variation in the wage at the local area level for identification in this study. I control for industry differences in capital intensity, α in the equation above, through 4 digit SIC industry fixed effects. This identification strategy assumes that the local area wage is the wage that the plant faces when it chooses its levels of factors. Plants in the same local area that have a higher wage bill relative to their number of workers (the plant level wage) are then assumed to have higher skilled workers.

This identification strategy uses crosssectional differences in local area wages and not differences in wages over time. Panel differences in wages can have severe biases if wage change over time is correlated with biased productivity change, as pointed out in Antras (2004) and León-Ledesma, McAdam, and Willman (2010). This identification strategy still assumes that the local area wage is orthogonal to the plant's rental rate for capital and level of labor augmenting productivity. Under these conditions, local area wage variation provides a consistent estimate of the elasticity of substitution. In Subsection 4.2.2, I employ fixed effect and instrument approaches to control for both rental rate of capital differences and for local area productivity differences affecting local area wages.

My identification strategy uses the factor cost ratio and not the capital-labor ratio to obtain the elasticity for two reasons. First, any measurement error in wages will bias the elasticity towards one with the factor cost ratio, and towards zero with the capital-labor ratio. Since I am testing whether the elasticity is one, I prefer this bias to be towards one. Second, a capital labor ratio implies that the error term will include differences in worker quality across plants. I later use the above equation to extract a measure of labor augmenting productivity. I do not want this measure of productivity to include worker quality differences across plants, which are really just variation in the level of worker input.

4.1.1 Wage Data

I measure local area wages from data on workers and data on employers. The first source of wage variation is from the Census 5% sample of Americans. The Population Censuses have the largest survey of workers and contain data on both the MSA of workers. I calculate the wage as wage and salary income divided by the total hours worked for prime age private sector men. The Population Censuses also allow me to control for differences in worker quality across areas in my measure of the local area wage. I then construct the average log wage for each MSA after controlling for education, experience, industry, occupation, and race of workers. Since the Population Census is only conducted every ten years, I match the 1990 Population Census to the 1987 Census

of Manufactures and the 2000 Population Census to the 1997 Census of Manufactures.

The second source of local area wages is from establishment data in the Longitudinal Business Database (LBD). The LBD contains employment and payroll data for all US establishments. I define the wage as payroll divided by employment. I then construct average log wages for each county in the United States. A major advantage of this data is that I can include plants in counties that are not in MSAs. Since the LBD is yearly, I match the Manufacturing Censuses to wages from the appropriate year of the LBD. For the LBD wages, though, I can not adjust for differences in worker quality across areas as I can with worker data.

4.1.2 Persistent Wage Variation

These wages have a substantial amount of variation across local areas. Figure 4.1 depicts the differences in the MSA log wage for 1990 across selected MSAs. The zero point in the figure is the US average in 1990. In 1990, the New York MSA has the highest wage among these MSAs with a wage about 20% higher than the national average. The San Antonio MSA, however, has wages about 16% lower than the national average. The difference between the New York MSA and San Antonio MSA is a log differential of .35. This wage variation exists within state boundaries as well. For example, San Francisco has much higher wages than San Diego.

Figure 4.1: MSA Level Adjusted Log Wages in 1990 for Selected MSAs

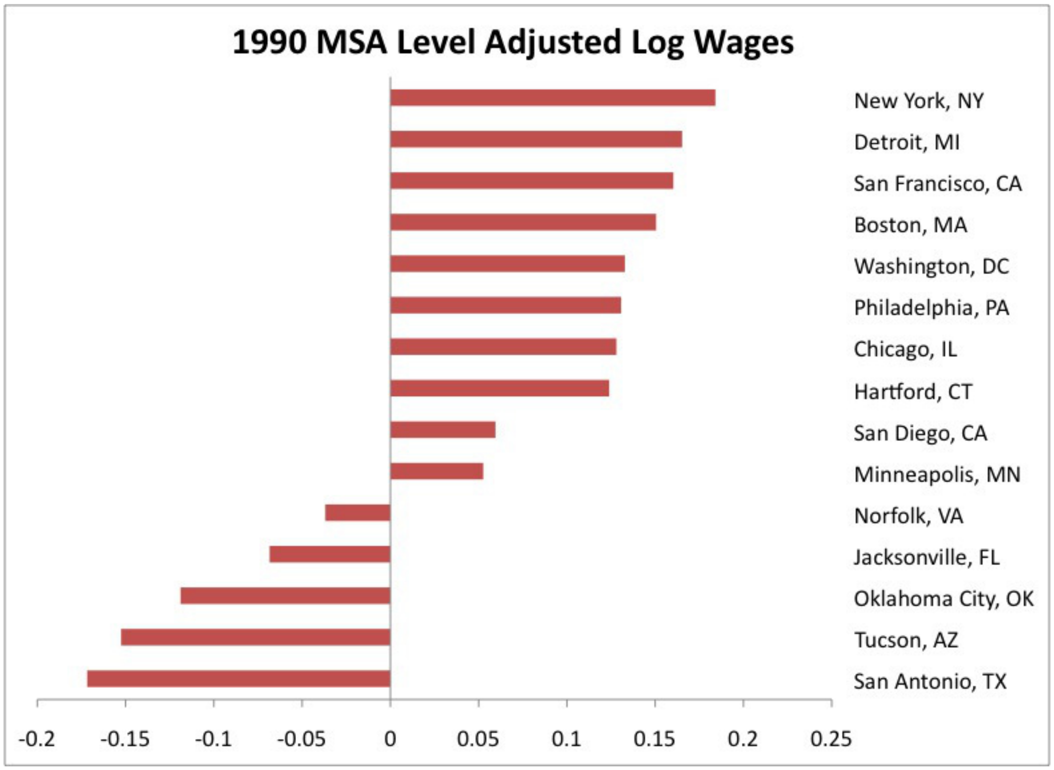

The X axis is the average log wage at the MSA level after controlling for education, experience, race, industry, and occupation, where the wage is calculated as the wage and salary income over total number of hours worked for a worker in the Census 5% sample data.

Moreover, this variation is persistent over time. Wage persistence is important as plants with adjustment costs will react much less to temporary differences in wages than permanent differences in wages. Persistent differences in wages will identify the correct theoretical parameter for this study, the long run elasticity of substitution. Figure 4.2 displays the differences in the MSA log wage for 1990 and 2000 across the same MSAs as before, where the red line is 1990 and blue line 2000. The high wage areas in 1990, such as New York, Detroit, and San Francisco, also have high wages in 2000, while low wage areas in 1990 have low wages in 2000. A few MSAs do see big changes between 1990 and 2000. Wages fall in Norfolk and San Diego, both big defense spending areas, with the end of Cold War defense spending. Wages rise in San Francisco with the tech boom of the late 1990s. Still, the correlation for all MSAs in the log wage between 1990 and 2000 is

very high at .90.

Figure 4.2: MSA Level Adjusted Log Wages in 1990 and 2000 for Selected MSAs

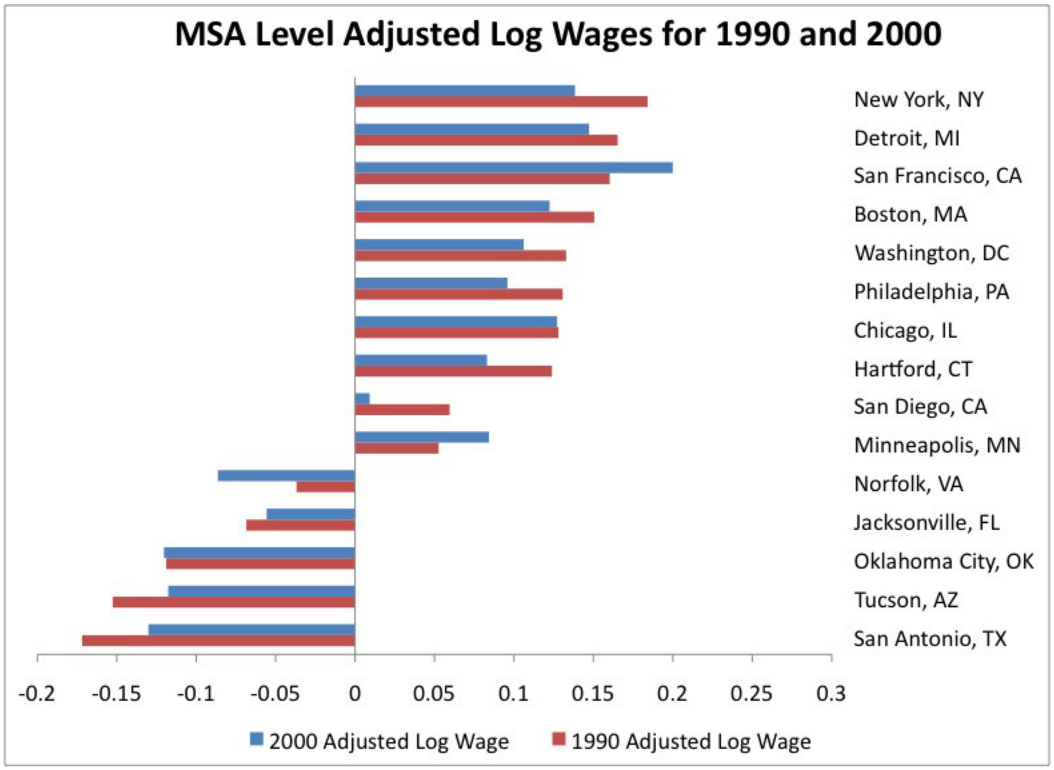

The X axis is the average log wage at the MSA level after controlling for education, experience, race, industry, and occupation, where the wage is calculated as the wage and salary income over total number of hours worked for a worker in the Census 5% sample data.

4.2 Estimates of the Elasticity of Substitution

Figures 4.3 and 4.4 nonparametrically plot the industry demeaned factor cost ratio against the MSA level wage and the county level wage in 1987. In both figures, the factor cost ratio is strongly decreasing in the local area wage. The log wage increases by about 40% as the factor cost ratio falls by 20% in Figure 4.3, indicating an elasticity of substitution close to a half.

Figure 4.3: Factor Cost Ratio by MSA Level Wage for 1987

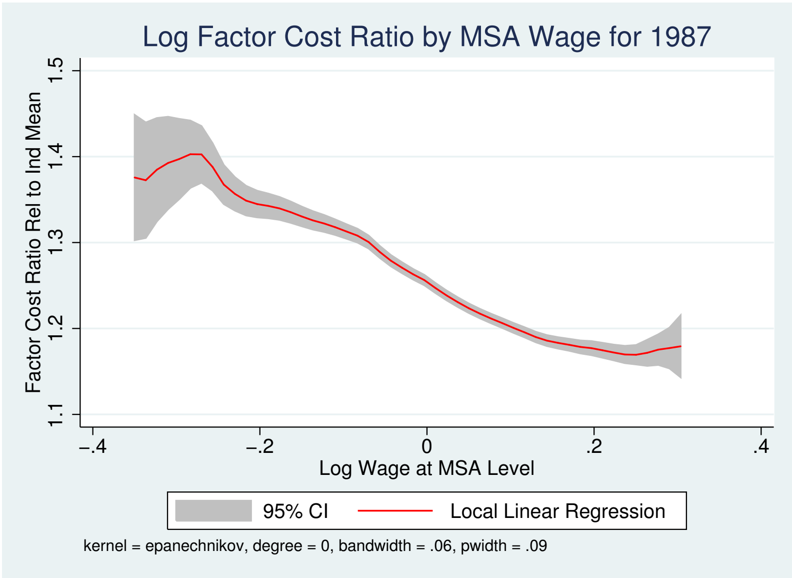

The X axis is the average log wage at the MSA level after controlling for education, experience, race, industry, and occupation, where the wage is calculated as the wage and salary income over total number of hours worked for a worker in the Census 5% sample data. The Y axis is the log factor cost share after taking out industry averages.

Figure 4.4: Factor Cost Ratio by County Level Wage for 1987

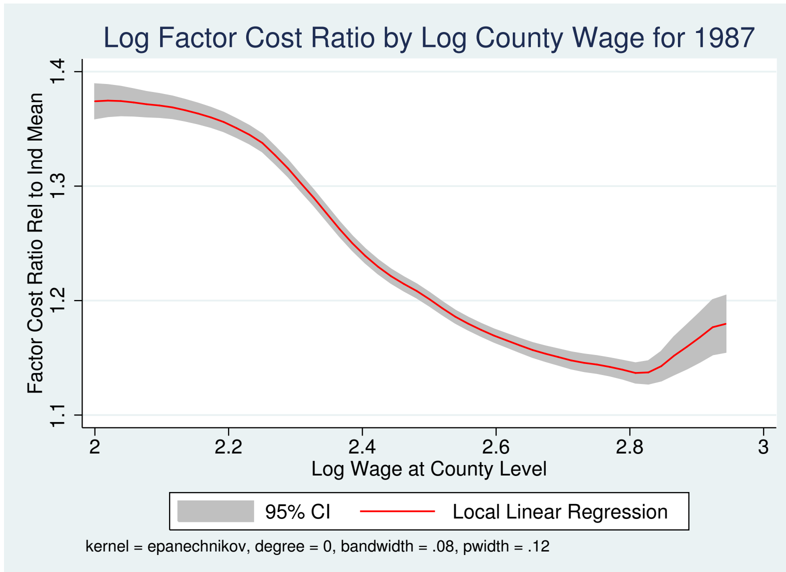

The X axis is the average log wage at the county level, where the wage is computed as payroll/number of employees at the establishment level for establishments in the Longitudinal Business Database. The Y axis is the log factor cost share after taking out industry averages.

A constant elasticity of substitution implies a linear relationship between the factor cost ratio and local area wage. Table 7 displays estimates of the elasticity of substitution across all manufacturing. I have used both MSA and county sources of variation in these regressions. I cluster standard errors at the two digit SIC- local area level where the local area is based on the source of the wage variation. This level of clustering adjusts the standard errors for correlated shocks to the factor cost ratio in local areas for plants in the same broad industry.

Table 7: Elasticities of Substitution between Labor and Capital for All Manufacturing

	MSA Level			County Level		
1987	.43	.52	.48	.59	.65	.69
	(.04)	*(.04)*	*(.05)*	*(.03)*	*(.03)*	*(.03)*
1997	.35	.46	.46	.59	.67	.69
	(.04)	*(.03)*	*(.04)*	*(.01)*	*(.01)*	*(.01)*
Source of Wage Data	Worker Data			Employer Data		
Controls	No	Yes	Yes	No	Yes	Yes
State Dummies	No	No	Yes	No	No	Yes
N	~125,000			~180,000		

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for MSA-level regressions, two digit SIC-MSA, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages and the wage is wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry. Controls are age fixed effects and a multi-unit status indicator.

The estimates of the elasticity of substitution from the MSA wage regressions are .43 in 1987 and .35 in 1997. When I add controls for age through fixed effects and for plant multiunit status these estimates rise slightly to .52 in 1987 and .46 in 1997. Using the county wages from the LBD, estimates of the elasticity of substitution are .59 without controls and .65 in 1987 and .67 in 1997 with controls. In all of these cases, I can easily reject that the production function is Cobb-Douglas.

I also run wage regressions with state level fixed effects which restricts all of the wage variation to be across local areas within a state. The previous regressions could have problems if state-level regulations affect both wages and the factor cost ratio.[5] Using within state MSA level wage variation, I find that the elasticity of substitution is .48 for 1987 and .46 for 1997. Using within state county level wage variation, I find that the elasticity of substitution is .69 for 1987 and 1997.

[5] For example, right to work laws make it more difficult for firms to unionize, which could both lower wages and make it easier for firms to automate and change their factor cost ratio. Holmes (1998) shows that plants do indeed respond to right to work laws, as industrial activity is higher than average in right to work states adjacent to non right to work states.

These estimates are slightly higher than those without state fixed effects, but I can still easily reject an elasticity of one.

In Table 7, the elasticities of substitution from the MSA regressions are all lower than those from the county regressions. The main reason for this difference is the adjustment for worker quality differences across areas. In Table 8, I show how the estimate of the elasticity of substitution changes after I adjust the wage for quality differences across areas. Unadjusted, the estimate of the elasticity of substitution in .62 in 1987 and 1997, quite similar to the (unadjusted) county level wage results. Adjusting the wage for quality differences lowers the amount of wage variation. For example, the New York MSA is 26% above the national average for the unadjusted wage and 20% above the national average for the quality adjusted wage in 1990. The San Antonio MSA is 20% below the national average for the unadjusted wage in 1990 and 16% below for the quality adjusted wage in 1990. Since the variation in the factor cost ratio remains constant as the wage variation falls, the elasticity of substitution falls. If I have not fully controlled for worker quality differences across areas, the elasticity of substitution will be even lower than my estimates. I also measure a manufacturing specific wage by only using manufacturing workers in a given local area but still controlling for worker quality otherwise. With the manufacturing worker only wage, I estimate the elasticity of substitution to be .48 in 1987 and .42 in 1997, similar to the estimates using all workers.

Table 8: Elasticities of Substitution for All Manufacturing, for Different Wage Specifications

	MSA Level		
1987	.62	.43	.48
	(.03)	*(.04)*	*(.04)*
1997	.62	.35	.42
	(.03)	*(.04)*	*(.03)*
Source of Wage Data	Worker Data		
Wage Data is:	Unadjusted, All Industries	Adjusted, All Industries	Adjusted, Manufacturing Only
Controls	No		
N	∼125,000		

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for MSA-level regressions, two digit SIC-MSA, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry in the second and third specifications. Controls are age fixed effects and a multi-unit status indicator.

4.2.1 Industry Estimates

So far, I have assumed that the elasticity of substitution is constant across all industries. I now allow the elasticity of substitution to vary at the two digit industry level. Two digit SIC industries are major industry groupings within manufacturing. For example, Textiles or Primary Metals are two digit SIC industries, while Carpets and Rugs (SIC 2273) and Steel Blast Furnaces (SIC 3312) are four digit SIC industries within these broader two digit SIC industries. Table 9displays estimates of the elasticity of substitution for each two digit SIC industry for 1987 and 1997. I exclude the tobacco industry because it is much smaller than the other two digit industries.

Table 9: Elasticities of Substitution between Labor and Capital for Two Digit SIC Industries

SIC Two Digit Industry:	Level of Wage Variation				N for 1987
	MSA Level, 1987	MSA Level, 1997	County Level, 1987	County Level, 1997	
20: Food Products	.67 *(.10)*	.87 *(.11)*	.71 *(.04)*	.81 *(.04)*	∼10,000
22: Textiles	.70 *(.16)*	.30 *(.24)*	.70 *(.08)*	.50 *(.07)*	∼3,500
23: Apparel	.82 *(.11)*	.40 *(.09)*	1.03 *(.04)*	.83 *(.03)*	∼12,000
24: Lumber and Wood	.23 *(.12)*	.48 *(.11)*	.43 *(.04)*	.53 *(.04)*	∼15,000
25: Furniture	.42 *(.14)*	.18 *(.17)*	.47 *(.05)*	.49 *(.06)*	∼6,000
26: Paper	.20 *(.16)*	.20 *(.15)*	.46 *(.06)*	.61 *(.05)*	∼4,000
27: Printing and Publishing	.57 *(.05)*	.50 *(.08)*	.69 *(.03)*	.66 *(.03)*	∼26,000
28: Chemicals	.41 *(.15)*	.51 *(.21)*	.52 *(.08)*	.54 *(.08)*	∼6,500
29: Petroleum Refining	.70 *(.23)*	.53 *(.28)*	.80 *(.11)*	.81 *(.11)*	∼1,500
30: Rubber	.64 *(.13)*	.42 *(.14)*	.60 *(.05)*	.59 *(.04)*	∼8,500
31: Leather	.43 *(.28)*	.46 *(.36)*	.88 *(.12)*	.99 *(.12)*	∼1,000
32: Stone, Clay, Glass, Concrete	.47 *(.11)*	.80 *(.16)*	.64 *(.04)*	.58 *(.05)*	∼9,000
33: Primary Metal	.42 *(.17)*	.26 *(.19)*	.69 *(.06)*	.67 *(.07)*	∼4,000
34: Fabricated Metal	.33 *(.09)*	.25 *(.09)*	.56 *(.04)*	.59 *(.04)*	∼20,000

Table 1.3: (contd)

SIC Two Digit Industry:	Level of Wage Variation				N for 1987
	MSA Level, 1987	MSA Level, 1997	County Level, 1987	County Level, 1997	
35: Machinery	.54 *(.08)*	.52 *(.11)*	.68 *(.02)*	.70 *(.03)*	~25,000
36: Electrical Machinery	.48 *(.12)*	.51 *(.12)*	.65 *(.07)*	.70 *(.05)*	~8,000
37: Transportation Equip	.65 *(.16)*	.77 *(.16)*	.70 *(.06)*	.75 *(.06)*	~5,000
38: Instruments	.74 *(.10)*	.71 *(.13)*	.67 *(.07)*	.74 *(.07)*	~4,500
39: Misc	.43 *(.13)*	.38 *(.12)*	.54 *(.04)*	.53 *(.05)*	~6,500
Source of Wage Data	Census 5% individual samples		Longitudinal Business Database		
Controls	Yes				

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages and the wage is wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry.

First, almost all of the point estimates are below one. In only one case, the Apparel industry estimate in 1997 using county level wage differences, do I find a point estimate of the elasticity that is above one. I can reject that the elasticity of substitution is one for 15 of 19 industries using MSA level wages and 16 of 19 industries using county level wages in 1987. Most of the estimates are concentrated in a narrow band between .40 and .70. 15 out of the 19 estimates in 1987 are in this band, both using the MSA wages and the county wages. The across industry differences in

elasticity also accord with intuition. Industries that primarily use chemical processes to produce output, and so would have less scope for labor capital substitution, have elasticities on the low side. The paper industry has elasticities of .2 in the MSA estimates, and the chemicals industry between .4 and .5. The transportation industry, with assembly line manufacturing that allows for greater labor capital substitution, has elasticities between .65 and .80. In general, however, the SIC two digit level estimates are similar to those using the entire manufacturing sector.

I also examine a set of large, geographically diverse industries at the narrow 4 digit level to control for selection effects. Selection of manufacturing plants into local areas based on factor costs should bias the elasticity of substitution towards one. Plants with high labor shares would prefer to locate in low wage areas, giving low wage areas a higher factor cost ratio. This selection effect explains, for example, why Cobb Douglas production functions with different capital shares can not explain the data.

To see how this force of selection affects my results, I look at 10 four digit SIC industries with substantial geographic variation. I select all the industries that are located in at least 300 MSAs or at least 250 MSAs and 48 states, dropping industries that have less than 1,000 plants or are miscellaneous industries. Newspaper Publishing, Commercial Lithographic Printing, and Ready Mixed Concrete are among the biggest of these industries. Ready Mixed Concrete is perhaps the best test case; since ready mixed concrete cannot be shipped very far, every location that has construction activity must have concrete plants. Thus, concrete plants can not select into where they locate based on factor costs. Table 10 displays the estimates of the elasticity of substitution using MSA and county level wages for these industries. I can reject Cobb-Douglas for eight out of ten industries with MSA level wages and ten out of ten industries with county level wages in 1987. For ready mixed concrete, I find an elasticity of substitution of .36 in the MSA-level regressions and .77 in the county level regressions in 1987, and .98 in the MSA level regressions and .62 in the county level regressions in 1997. Only the 1997 MSA estimate is close to one.

Table 10: Elasticities of Substitution between Labor and Capital for Selected 4 Digit Industries

SIC Four Digit Industry	Level of Wage Variation				N
	MSA Level, 1987	MSA Level, 1997	County Level, 1987	County Level, 1997	
2711: Newspaper Publishing	.43 *(.20)*	NA	.82 *(.07)*	NA	~4,000
2752: Commercial Printing, Lithographic	.69 *(.07)*	.52 *(.09)*	.72 *(.03)*	.67 *(.04)*	~12,000
3272: Concrete Products, Except Block and Brick	.41 *(.20)*	.87 *(.39)*	.49 *(.10)*	.49 *(.10)*	~2,000
3273: Ready Mixed Concrete	.36 *(.17)*	.98 *(.24)*	.77 *(.07)*	.62 *(.07)*	~4,000
3441: Fabricated Structural Metal	.35 *(.18)*	.58 *(.26)*	.52 *(.10)*	.75 *(.10)*	~1,500
3444: Sheet Metal Work	.65 *(.19)*	.28 *(.20)*	.55 *(.09)*	.58 *(.07)*	~3,000
2051: Bread and other Bakery Products, except Crackers	.91 *(.24)*	1.17 *(.24)*	.71 *(.12)*	.99 *(.13)*	~1,000
2421: Sawmills and Planing Mills	.13 *(.32)*	.11 *(.29)*	.74 *(.08)*	.80 *(.09)*	~3,000
2431: Millwork	.04 *(.24)*	.49 *(.23)*	.11 *(.09)*	.51 *(.09)*	~1,500
2434: Wood Kitchen Cabinets	.26 *(.23)*	.66 *(.24)*	.31 *(.11)*	.60 *(.09)*	~2,000
Source of Wage Data	Census 5% individual samples		Longitudinal Business Database		
Controls	Yes				

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages and the wage is wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry.

4.2.2 Endogeneity Problems

My identification strategy relies upon differences in wages across local areas in the US. The obvious concern with this strategy is that the local area wage is endogenous to factors that also affect plant factor costs. The two potential problems I focus on here are differences in rental rates across local areas in the US and differences in productivity across areas causing wage differences.

The first concern is that the rental rate of capital also varies with the local area wage. For equipment capital, a national market for equipment means that the rental price is the same across local areas. For structures capital, however, local area wages and prices of materials can affect the price of capital. One way to control for rental rate differences is to examine equipment capital separately. I thus look at the elasticity of substitution between equipment capital and labor in Table 11 for 1987. Using the quality adjusted MSA wages, the elasticity between labor and equipment capital is .45 in the regressions without controls and .53 in the regressions with controls. Using the non quality adjusted county wages, the elasticity between labor and equipment capital is .61 in the regressions without controls and .67 in the regressions with controls. These values are very similar to the estimates using the full capital stock.

Table 11: Elasticities of Substitution between Labor and Equipment Capital for All Manufacturing

	MSA Level		County Level	
1987	.45	.53	.61	.67
	(.04)	*(.03)*	*(.03)*	*(.03)*
Source of Wage Data	Census 5% individual samples		Longitudinal Business Database	
Controls	No	Yes	No	Yes
N	$\sim$125,000		$\sim$180,000	

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages and the wage is wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry. Controls are age fixed effects and a multiunit indicator.

I also include firm fixed effects in the regression in the sample of plants that belong to multiunit firms. Firm fixed effects control for firm level differences in the cost of capital as well as firm level productivity differences. Table 12 reports these estimates. The estimates of the elasticity of substitution with firm fixed effects are .49 in 1987 and .48 in 1997 and are quite similar to my baseline estimates.

Another concern is that differences in wages across local areas are caused by productivity differences in these areas. The fixed effect regression mentioned above already controlled for firm level differences in productivity. Another way to examine this issue is to look at local areas where manufacturing is small relative to the local economy. For these local areas, productivity shocks to manufacturing will not affect local wages much. Table 12 also contains estimates where I restrict the sample to only counties where the share of manufacturing employment is below the national median. With this restricted sample, I find estimates of the elasticity of .65 in 1987 and .66 in 1997 using county level wages, on par with the full sample estimates.

Table 12: Robustness Checks for Elasticity of Substitution between Labor and Capital

Robustness Check:	Low Manufacturing Employment Counties Only	Controls for Firm Fixed Effects
1987	.65 *(.02)*	.49 *(.05)*
1997	.66 *(.01)*	.48 *(.08s)*
Wage Variation:	County Level, Employer Data	MSA Level, Worker Data
Controls	Yes	
N	~80,000	~45,000

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 (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages . For the Census 5% sample data, I calculate the residual wage after controlling for composition differences, where the wage is wage and salary income over total number of hours worked. I report results only for plants in counties where the manufacturing percentage of employment is below the national median for that year. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry.

I also examine the salience of endogeneity problems through various instrumental variables for the local area wage. Each instrument selects out a different slice of the variation in wages. The first instrument is an instrument for labor demand based on the fact that national level industry shocks affect local areas differently based upon the initial industrial composition of the area. Formally, this instrument is the interaction between initial local area employment shares of industries and the 10 year national employment growth rate of these industries as used in **?**. I restrict the industries to non-manufacturing industries to avoid potential bias problems. Table 13 contains the instrument estimates. Using changes in industries at the 2 digit SIC level, I obtain estimates close to zero but with high standard errors. I can only still an elasticity of one, however. For 1997, I can use industry changes at the much more narrow 4 digit SIC level and obtain a much more precise estimate of

.33.[6] This labor demand instrument based estimate is much lower than the OLS estimate with county wages of .67. This difference may be due to the endogeneity problems discussed earlier. Another interpretation of the lower estimate is that this source of wage variation is more temporary than the full variation in wages, so firms substitute between labor and capital to a lesser extent.

I also examine two other instruments for wages. One potential reason for differences in wages across local areas is differences in cost of living. If everyone could costlessly adjust their location, high cost of living areas would have high wages to keep people indifferent on their location choice. Many of the high wage areas such as New York and San Francisco also have a high cost of living. The proxy I use for cost of living differences is housing and rental prices at the local area level. I take housing and rental price data from the Population Census 5% samples and produce a MSA housing and rental price as the MSA average residual after controlling for housing quality. To avoid cyclical variation due to booms and busts, I instrument the quality adjusted MSA wages with both current housing prices and rents and their ten year lead and lag values. I obtain an estimate of the elasticity of substitution of .34 in 1987 and .38 in 1997, slighly lower than the OLS estimates.

The final instrument is simply the ten year lag and lead values for the MSA wage. Instrumenting with previous and future wages isolates the permanent component to the wage variation across local areas. In these regressions, my estimates of the elasticity of substitution are .57 in 1987 and .40 in 1997, similar to the baseline OLS estimates.

[6] The SIC industry definitions changed between 1977 and 1987 from 1972 SIC to 1987 SIC, so I could not do the same procedure in 1987.

Table 13: Instrument Based Estimates of the Elasticity of Substitution between Labor and Capital

Instrument:	Labor Demand Interaction Using Non Manufacturing Industries, 2 digit SIC	Labor Demand Interaction Using Non Manufacturing Industries, 4 digit SIC	Housing Prices and Rents, this year and 10 year lag/lead	10 year lag/lead wages
1987	.17 *(.34)*	NA	.34 *(.03)*	.57 *(.03)*
1997	.07 *(.42)*	.33 *(.09)*	.38 *(.05)*	.40 *(.03)*
Wage Variation:	County Level, Employer Data		MSA Level, Worker Data	
Source of Wage Data	Longitudinal Business Database		Census 5% individual samples	
Controls	Yes			
N	~125,000			

Note: All regressions include industry dummies, age fixed effects, and a multiunit status indicator. Labor demand instruments are based on the interaction between 10 year lag industry composition of employment in the MSA and nationwide changes in labor demand at the industry level. 10 year lag wage instrument is just the 10 year lagged wage. Housing Price and Rent instruments are based on Census data on prices and rents across local areas, controlling for housing composition differences across areas. For 1997, I can only include the lag housing price, rent, or wage as the 2010 Census has not yet been released. Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages . For the Census 5% sample data, I calculate the residual wage after controlling for composition differences, where the wage is wage and salary income over total number of hours worked. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry.

Another potential concern is that the plant's choice of its level of labor augmenting productivity B depends on the local wage. If plants adjust their level of labor augmenting technology because of the local wage, B will be related to the wage. When wages are high, labor augmenting technology that saves on labor is more valuable. Such wage based technology adoption would cause high wage areas to have high levels of labor augmenting technology and high capital shares. This form

of technology adoption would bias the estimate of the elasticity of substitution towards one.

4.2.3 Capital Robustness

I also check to see whether my estimates of the elasticity of substitution change if I switch my measure of capital to the perpetual inventory measure of capital I constructed in Subsection 2.3.1. Classical measurement errors in capital should not affect my results as the factor cost ratio is the dependent variable. Indeed, I estimate that the elasticity of substitution is .35 in the regressions without controls and .47 in the regressions with controls using quality adjusted MSA wages, and .63 in the regressions without controls and .70 in the regressions with controls using unadjusted county level wages. These estimates are fairly close to the estimates for 1987 using book value measures of capital.

Table 14: Capital Robustness: Elasticities of Substitution between Labor and Capital for All Manufacturing using a Perpetual Inventory Measure for Capital

	MSA Level		County Level	
1987	.35 *(.07)*	.47 *(.07)*	.63 *(.04)*	.70 *(.04)*
Source of Wage Data	Census 5% individual samples		Longitudinal Business Database	
Controls	No	Yes	No	Yes
N	∼30,000		∼50,000	

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages and the wage is wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry. Controls are age fixed effects and a multiunit indicator.

4.3 Productivity

Since I have estimated the elasticity of substitution, I can now examine the productivity of each plant. I do not want to assume that productivity is enter only as Hicks neutral or labor augmenting. Thus, I estimate both a Hicks neutral productivity parameter A and a labor augmenting productivity parameter B. Here, I assume that the quantity production function is a CES production function with constant returns to scale. I subsume α into labor augmenting productivity B. Since the elasticity of substitution is less than one, this is without loss of generality. I can calculate the labor augmenting productivity of the plant straight from its factor allocations and payments, without using any data on output. Manipulating equation 3.3, I have that:

$$\log B = \log(w/r) + \frac{1}{1-\sigma}\log(\frac{rK}{wL}) \tag{4.2}$$

This measure of labor augmenting productivity is the residual from the elasticity of substitution regression multiplied by $\frac{1}{1-\sigma}$, which about doubles the factor cost ratio given my measures of the elasticity of substitution. I use the county level wage for w. I set σ to .52, the elasticity of substitution I estimate in the MSA level wage regressions for all manufacturing plants in 1987, controlling for age fixed effects, industry fixed effects, and a multi unit indicator.

To estimate the Hicks neutral productivity A, I take the equation for the average product of capital and impose cost minimization conditions on the factors. Given measures of quantity, the Hicks neutral productivity A is then:

$$\log A = \log(Y/K) - \frac{\sigma}{1-\sigma}\log(\frac{rK}{rK+wL}) \tag{4.3}$$

For most of the plants in my sample, I only have revenue measures and not quantity measures. Given an isoelastic demand curve, I can only measure neutral productivity together with demand

shocks:

$$\log A + \log D = \frac{\varepsilon}{\varepsilon - 1}\log(PY/K) - \frac{\sigma}{1-\sigma}\log(\frac{rK}{rK+wL}) + \frac{1}{\varepsilon - 1}\log K \quad (4.4)$$

I then need an estimate of the elasticity of demand ε to obtain an estimate of $\log(A) + \log(D)$. An isoelastic demand function implies that the markup over marginal cost is a function of the elasticity of demand. I calculate the markup as revenue divided by total costs (capital, labor, and materials). Measuring the elasticity of demand from the average markup, I find that the elasticity of demand is 3.94 in 1987 and 4.04 in 1997. I use the 1987 value for the productivity estimates. Oberfield and Raval (2011) provide some more details on estimating the elasticity of demand in this setup.

I first look at correlations between these measures of productivity. In 1987, the correlation between $\log A + \log D$ and $\log B$ is -.87. One explanation for this negative correlation is measurement errors in capital, which push up $\log B$ and pull down $\log A + \log D$. Another reason is that a purely capital augmenting and labor augmenting parameter are distributed independently, which would induce a negative correlation between neutral productivity and labor augmenting productivity as defined here. If I redefine productivity as a capital augmenting productivity and labor augmenting productivity, the correlation between the two is close to zero at .14. This low correlation is consistent with the theoretical foundation ofJones (2005), which assumes that a capital augmenting productivity and labor augmenting productivity come from independent Pareto distributions.

I now reexamine some of the standard relationships between productivity and plant level variables found in the literature.

4.3.1 Persistence

I first examine the autocorrelation of productivity, which is an important parameter for theory. For example, Moll (2010) finds that the effect of misallocation frictions on plant welfare depends

Table 15: Autocorrelation of Productivity

	Ten Year	Implied One Year	Ten Year	Implied One Year
Log(A)+Log(D)	.27 *(.004)*	.87	.31 *(.004)*	.89
Log(B)	.34 *(.004)*	.90	.43 *(.004)*	.92
Weights	No	No	Value Added	Value Added

All regressions contain 4 digit SIC industry dummies. The implied one year coefficient is the ten year coefficient to the 1/10 power.

critically on how persistent productivity shocks are. Table 15 contains estimates of the 10 year autocorrelation of productivity measures between the 1997 and 1987 Manufacturing Censuses. All three measures are positively correlated autocorrelated over time. My measure of neutral productivity and demand shocks has a 10 year autocorrelation of .27 in the unweighted regressions and .31 in the weighted regressions. Labor augmenting productivity has higher correlations over time, with a 10 year autocorrelation of .34 in the unweighted regressions and .43 in the weighted regressions. These results show that productivity is very persistent over time.

4.3.2 Correlation with Size

Profit maximization implies that plants with higher neutral productivity, demand shocks, and labor augmenting productivity have higher sales. The effect on unemployment for labor augmenting productivity is not as strong because more productive plants are using more capital intensive techniques. Table 16 examines how productivity varies with the value added of the plant. In the table, each cell is a separate regression with a log productivity measure as the dependent variable and a log value added measure as the independent variable, along with industry dummies as controls. I use productivity measures as the dependent variable so measurement errors in capital do not bias the coefficients.

Table 16: Correlations between Productivity and Value Added

	1987		1997	
Log(A)+Log(D)	.31 *(.002)*	.26 *(.002)*	.33 *(.003)*	.26 *(.002)*
Log(B)	.15 *(.003)*	.20 *(.002)*	.12 *(.003)*	.26 *(.002)*
Weights	No	Value Added	No	Value Added

All regressions contain 4 digit SIC industry dummies. Log(VA) is the independent variable and a measure of productivity the dependent variable.

Both measures of productivity are highly correlated with value added as predicted by theory. A 100% increase in value added increases the neutral productivity and demand shock measure by 31% in the unweighted regressions and 26% in the weighted regressions in 1987. The same increase in value added will increase labor augmenting productivity by 15% in the 1987 unweighted regressions and 20% in the 1987 weighted regressions. Thus, both measures of productivity are correlated with plant value added. Table 17 examines correlations with total employment. A 100% increase in employment increases the neutral productivity and demand shock measure by 28% in the unweighted regressions and 22% in the weighted regressions in 1987. The same increase in employment will increase labor augmenting productivity by 5% in the 1987 unweighted regressions and 12% in the 1987 weighted regressions. Thus, correlations with employment are stronger for the neutral productivity and demand shock measure than for the labor augmenting productivity measure.

In the above results I can not separate the effect of demand shocks from neutral productivity, making it unclear whether neutral productivity itself is correlated with plant size. Foster, Haltiwanger, and Syverson (2008) point out that demand shocks are an important component of conventioanlly measured productivity. To examine demand shocks from neutral productivity, I look separately at the set of homogeneous product industries examined by Foster, Haltiwanger, and Syverson (2008). Here, I can measure sales and quantity separately, and so construct a quantity

Table 17: Correlations between Productivity and Employment

	1987		1997	
Log(A)+Log(D)	.28	.22	.29	.18
	(.003)	*(.002)*	*(.003)*	*(.003)*
Log(B)	.05	.12	.07	.21
	(.003)	*(.003)*	*(.003)*	*(.003)*
Weights	No	Value Added	No	Value Added

All regressions contain 4 digit SIC industry dummies. The dependent variable is a measure of productivity and independent variable the log of employment.

measure of productivity A. These regressions are pooled over 1987, 1992, and 1997 and include product year fixed effects. Given that only a few products are included, the unweighted regressions represent mostly concrete plants and the weighted regressions gasoline refining plants.

Table 18 contains the estimates from these regressions. Labor augmenting productivity B is correlated with quantity produced. A 100% increase in quantity produced leads to a 24% increase in labor augmenting productivity in the unweighted regressions and a 55% increase in labor augmenting productivity in the weighted regressions. Just as in the earlier regressions, labor augmenting productivity is also positively correlated with sales. Neutral productivity, on the other hand, is negatively correlated with quantity produced, as a 100% in quantity produced has a -3% fall in neutral productivity in the unweighted regressions and a 29% fall in neutral productivity in the weighted regressions. The results with sales are similar, with a 100% rise in sales leading to a 4% rise in neutral productivity in the unweighted regressions and a 15% fall in neutral productivity in the weighted regressions. The combination of productivity and demand shocks is positively correlated with both quantity and sales. For these homogenous products, large plants have high labor augmenting productivity and high demand shocks, but either a small increase or drop in neutral productivity. Neutral productivity thus may not be an important part of productivity differences among firms relative to demand shocks and labor augmenting productivity.

Table 18: Correlations between Productivity and Size Measures, using Quantity Data Subsamples

	Log(Y)		Log(PY)	
Log(A)+Log(D)	.80	.12	.67	.15
	(.02)	*(.02)*	*(.02)*	*(.02)*
Log(A)	-.03	-.29	.04	-.15
	(.02)	*(.02)*	*(.02)*	*(.02)*
Log(B)	.24	.55	.24	.44
	(.02)	*(.03)*	*(.02)*	*(.03)*
Weights	No	Value Added	No	Value Added

All regressions contain 4 digit SIC industry dummies. The dependent variable is a measure of productivity and independent variable the log of quantity produced or log of sales. Standard errors here are clustered at the manufacturing plant level.

4.3.3 Exports

The trade literature has found that exporting plants are more capital intense and also more productive (Bernard, Jensen, Redding, and Schott (2007)). The main models of trade, such as the Bernard, Eaton, Jensen, and Kortum (2003) and Melitz (2003) models, are based upon firms with high productivity selecting to export. I examine the impact of various measures of productivity on the intensive margin of exports in Table 19. The dependent variable is the log of exports, so I exclude all non exporters. Exporters that export a lot have high labor augmenting productivity. In 1987, a plant with a 100% increase in exports has on average a 15% increase in labor augmenting productivity B in the unweighted regressions and an 11% increase in the weighted regressions. These same correlations are much lower for the measure of neutral productivity and demand shocks, with a 4% rise in the 1987 unweighted regressions and a 5% rise in the 1987 weighed regressions. In the 1997 unweighted regression, this increase in neutral productivity is almost nonexistent. Thus, labor augmenting productivity appears to be more important than neutral productivity for exporting behavior as well.

Table 19: Correlations between Productivity and Exports

	1987		1997	
Log(A)+Log(D)	.04	.05	.008	.06
	(.004)	*(.003)*	*(.004)*	*(.003)*
Log(B)	.15	.11	.18	.13
	(.004)	*(.004)*	*(.004)*	*(.003)*
Weights	No	Sales	No	Sales

All regressions contain 4 digit SIC industry dummies. The log of exports is the independent variable and a measure of productivity the dependent variable.

5 Caveats

So far, I have placed a number of strong assumptions on both the production function and firm behavior. In this section, I relax some of these assumptions to examine how the main conclusions of this paper change. The first major restriction I have placed upon production technology is a constant elasticity of substitution between labor and capital. More general production functions, such as the translog production function of Christensen, Jorgenson, and Lau (1973), allow the elasticity of substitution to vary across plants. Below, I test the constancy of the elasticity of substiution through wage spline and quantile estimations.

I have also assumed a value added production function using only capital and labor as factors. I now examine the role of materials in a gross output production framework. Finally, I have assumed that manufacturing plants face competitive input markets. I also look at a non competitive market for labor through unions in my framework.

5.1 Non Constant Elasticity of Substitution

5.1.1 Wage Splines

Under a CES production function, the factor cost ratio is a linear function of the local area wage. A non linear relationship between the factor cost ratio and the wage could thus indicate a non

constant elasticity of substitution. I test for nonlinearities in this relationship through wage spline regressions. Each segment of the wage distribution is free to have its own slope and so its own elasticities of substitution. I take segments of the wage distribution by either splitting the wage distribution at its midpoint, or at the points a third of the way and two thirds of the way from the endpoint. This procedure will either select the first half and second half of the wage distribution or the first third, middle third, and last third of the wage distribution. I allow a common intercept but the slope of the relationship can vary by the segment of the wage distribution. Table 20 contains these wage spline estimates using the MSA worker quality adjusted wage variation for 1987 and 1997. For the 1997 wage splines with two segments, the elasticity of substitution from the first half is almost identical to the second half, at .45 and .46. Only in this case can I not reject that the segments have different slopes.

However, the other wage spline regressions have only modest departures from a constant elasticity. In all of the other spline regressions, the elasticity of substitution varies within a narrow band from .40 to .70. Thus, labor and capital are always complements. The elasticity of substitution for the first half segment is .66 and for the second half segment .42 in the 1987 two segment regression, for example. The wage spline regressions do not reveal economically significant violations of the constant elasticity assumption.

Table 20: Elasticities of Substitution between Labor and Capital for All Manufacturing at Different Segments of the Wage Distribution

	1987		1997	
Segment of Wage Distribution:				
First Half	.66 *(.07)*		.45 *(.08)*	
Second Half	.42 *(.06)*		.46 *(.06)*	
First Third		.50 *(.07)*		.63 *(.07)*
Middle Third		.69 *(.07)*		.41 *(.06)*
Last Third		.45 *(.05)*		.40 *(.06)*
P-value from Equality Test	.03	.02	.95	.04
Wage Variation	MSA Level, Worker Data			
Controls	Yes			
N	∼120,000			

Note: All regressions include industry dummies and have standard errors clustered at the two digit industry-area level (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as wage and salary income over total number of hours worked for the Census 5% sample data. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry. Controls are age fixed effects and a multiunit indicator. Here I take spline estimates- with knots at the midpoint of the wage distribution in one specification and at a third and two thirds of the distance between the endpoints of the wage distribution in another specification.

5.1.2 Quantiles

A constant elasticity production function also implies that the elasticity is constant for every quantile of the factor cost ratio distribution. Other more general production functions allow for the elasticity of substitution to vary by the plant's level of capital and labor. For example, the Variable Elasticity of Substitution (VES) production function developed by Revankar (1971) has an elasticity of substitution that is monotone in the plant capital labor ratio. Thus, a VES produc-

tion function should have quantile estimates of the elasticity of substitution that are monotonically increasing or decreasing in the quantile. The basic labor augmenting productivity facts will not change with these production functions, however, as neutral productivity changes will continue to have no effect on the factor cost ratio.

Conditional quantile regressions should uncover the elasticity of substitution at each quantile of the factor cost ratio distribution. Before running these regressions, I remove industry, age, and multiunit variation from both the local area wage and the factor cost ratio through OLS regressions in a manner analogous to partitioned regression. I then estimate quantile regressions of the residual factor cost ratio on the residual wage for every decile from the 10th to the 90th and the 25th and 75th percentiles. I perform this procedure to have the regression controls shift only the location of the distribution: otherwise fixed effects for 459 industries will change at each quantile. [7]

Figure 5.1 displays the quantile estimates, where the blue line is the 1987 estimates, the red line the 1997 estimates, and the dashed lines 95% confidence bands. First, all of these elasticities are significantly less than one. Labor and capital are always complementary, which implies that the non-neutral productivity differences across plants detailed earlier are labor augmenting and not capital augmenting in nature. Second, the elasticity of substitution is not constant across quantiles but is not monotone in the quantile either. Instead, the elasticity of substitution has an inverted U shape with the lowest estimates at the upper and lower quantiles. The median quantile estimate is .75 in 1987 and .69 in 1997, higher than the mean estimates from the earlier regressions. The quantile estimates are even smaller for the lower quantiles than the upper quantiles. The 1997 lower quantile estimates are particularly low with the estimate of the elasticity of substitution not statistically significantly different from zero at the 10th quantile.

[7] Koenker (2004) has a detailed discussion of the problems of fixed effects in quantile regressions.

Figure 5.1: Quantile Estimates of the Elasticity of Substitution in 1987 and 1997

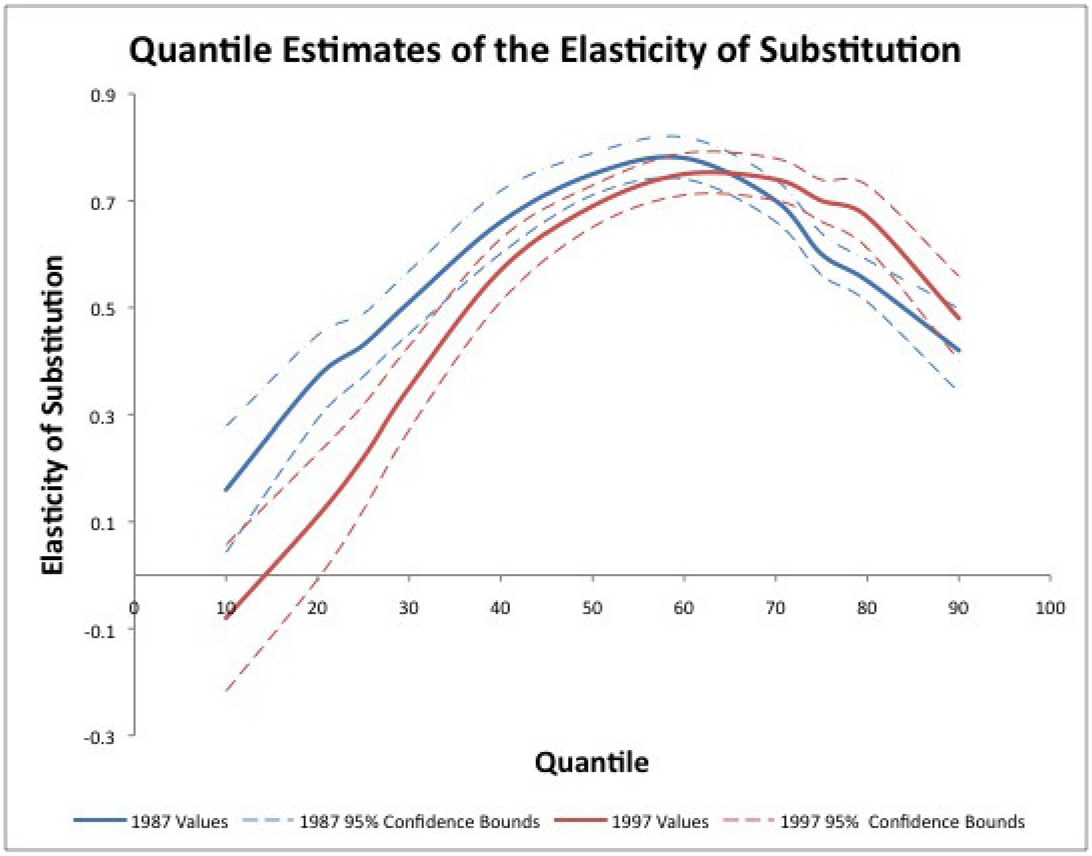

The X axis is the quantile of the conditional factor cost ratio distribution. The Y axis is the estimated value of the elasticity of substitution for that quantile, from a quantile regression of the factor cost ratio on the log MSA wage after taking industry and age fixed effects and a multiunit indicator out of both variables through regression. The dashes lines are 95% confidence bands formed by jointly estimating all the quantiles of a given year and bootstrapping the standard errors.

5.2 Additional Factors of Production

5.2.1 Materials

I have up until now 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 still hold with a gross output production function, as long as materials are separable from the capital-labor aggregate. For example, let $AF(K,BL)$ represent the production function for the capital-labor aggregate. A separable gross output production function will then be:

$$Y = G(AF(K,BL),h(M))$$

The quantity or price of materials will affect only the levels of capital and labor but not the factor cost ratio. Thus, the cost minimization conditions will imply the same estimation procedure for the elasticity of substitution, and the same formula for labor augmenting productivity B. Given expressions for the G and h functions, we can invert the above function for $AF(K,BL)$ and so solve for a measure of neutral productivity. For example, a Cobb Douglas production function between the capital-labor aggegate $AF(K,BL)$ and materials will mean that:

$$\begin{aligned} Y &= (AF(K,BL))^{1-\alpha_m} M^{\alpha_m} \\ AF(K,BL) &= \left(\frac{Y}{M^{\alpha_m}}\right)^{\frac{1}{1-\alpha_m}} \end{aligned}$$

I have examined the productivity correlations with size adjusting the measure of neutral productivity and demand shocks to 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 or amount of materials used can affect the ratio of capital costs to labor costs. To check for this, I control for the plant's share of materials in total costs and in total shipments in some of my main regressions. Table 21 contains the estimates of the regressions of the log factor cost ratio on log value added after controlling for materials intensity in this way. Plants with high materials shares also tend to have high capital shares, as plants with a 100% materials cost share have an 18% higher factor cost ratio, and those with a 100% materials share of shipments have a 16% higher factor cost ratio. However, the correlations with value added are unchanged. Similarly, Table 22 reports estimates of the elasticity of substitution after controlling for materials intensity.

Table 21: Correlations with Size for Factor Cost Ratio with Materials Intensity Controls

	1987		1997	
Cost Share:				
Log(Value Added)	.05 *(.001)*	.09 *(.005)*	.02 *(.001)*	.11 *(.006)*
Log(Materials Cost Share)	.18 *(.004)*	.11 *(.014)*	.08 *(.005)*	-.04 *(.020)*
Shipments Share:				
Log(Value Added)	.06 *(.001)*	.10 *(.005)*	.03 *(.001)*	.11 *(.005)*
Log(Materials Shipments Share)	.16 *(.004)*	.11 *(.013)*	.17 *(.005)*	.13 *(.017)*
Weights	No	Value Added	No	Value Added

All of these coefficients are from regressions with the log of the factor cost ratio as the dependent variable and log of value added as the independent variable. Controls include dummy variables for age and state, single establishment status and 4 digit SIC industry. I use robust standard errors. The materials cost share is the cost of materials over the cost of materials, capital, and labor. The materials share of shipments is the share of materials cost over total shipments.

The estimates of the elasticity of substitution are not affected by the materials controls. Thus, adding materials to the production function does not change the main conclusions of this paper.

Table 22: Estimates of the Elasticity of Substitution between Labor and Capital with Union Intensity and Materials Intensity Controls

Robustness Check:	Control for Materials Cost Share	Control for Materials Share of Sales	Controls for State Union Intensity	Controls for MSA Union Intensity
1987	.64 *(.02)*	.65 *(.02)*	.48 *(.04)*	.52 *(.04)*
1997	.66 *(.01)*	.66 *(.01)*	.40 *(.04)*	.41 *(.04)*
Wage Variation:	County Level, Employer Data		MSA Level, Worker Data	
Controls	Yes			
N	~180,000		~120,000	

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 (so for state-level regressions, two digit sic-state, etc.) Wages used are the average log wage for the geographic area, where the wage is computed as payroll/number of employees at the establishment level for the LBD wages . For the Census 5% sample data, I calculate the residual wage after controlling for composition differences, where the wage is wage and salary income over total number of hours worked. The average log wages using worker data are adjusted for differences in education, experience, race, occupation, and industry.

5.3 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 has already declined considerably from its peak. 25% of manufacturing workers are covered by a union in 1987, far below the 37% covered just ten years earlier in 1977. Only 17% of manufacturing workers are covered by a union in 1997. Given that only a minority of manufacturing workers are covered by unions, union biases may be relatively unimportant.

A union can affect 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 will appear in the cost minimization conditions as follows:

$$\log(rk/wl) = -(1-\sigma)\log(w/r) - (1-\sigma)\log(u_p) + (1-\sigma)\log B + \sigma \log \frac{\alpha}{1-\alpha} \quad (5.1)$$

A union premium will only affect the elasticity of substitution if the unionization of a plant varies with the wage of a plant's location. This bias will bias the estimate of the elasticity of substitution downwards. Union plants that pay higher wages than common in their local area will also have higher labor augmenting productivity than my estimates, as the plant's true labor input is lower than would be expected from its wage bill. 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 premium from worker data (in Hirsch's case, from the Current Population Survey (CPS)). Schmitz (2005) and Dunne, Klimek, and Schmitz (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. [8] The within state estimates detailed in Table 7 are only slightly different than the estimates allowing across state variation. The elasticity of substitution rises from .65 or .67 to .69 in the county level estimates, and actually falls for the MSA level estimates.

Ideally, I would also control for the plant level union status in the elasticity of substitution

[8] 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.

regressions. Since I do not have data on plant level union status, I control instead for the local area union intensity calculated by Hirsch and Macpherson (2003) from CPS data. Table 22 controls for the state level union coverage of manufacturing workers in one specification and for the MSA level union coverage of private sector workers in another specification in MSA level wage regresssions. The estimates of the elasticity of substitution are .48 in 1987 and .40 in 1997 controlling for state union intensity and .52 in 1987 and .41 in 1997 controlling for MSA level union intensity. The 1987 estimates are essentially unchanged from before and the 1997 estimates actually fall. Thus, it does not seem that correcting for unionization substantially affects the conclusions of this paper.

6 Application to Misallocation

A proposed explanation for the vast differences in TFP between rich and poor countries is that resources are not allocated well in poor countries. In this view, some highly productive firms in a poor country do not have enough capital, while other less productive firms have too much capital. This low allocative efficiency can then cause countries to have low aggregate TFP.

A number of recent papers have quantitatively examined whether misallocation can generate large losses in aggregate TFP. Hsieh and Klenow (2009) examine a static setup where profit maximizing manufacturing plants face output and capital wedges in their first order conditions. These wedges are meant to generalize many different mechanisms for misallocation. They then find that eliminating these wedges can increase aggregate TFP by 40% in the US and more than 100% in China and India. Midrigan and Xu (2009) explore a dynamic model with adjustment costs of capital and financial frictions and try to match the model to Korean plant data. They find that adjustment costs of capital and financial frictions can explain the observed variation in the time series marginal product of capital but not the large cross-section differences in the marginal product of capital.

All of these models assume that the production function is Cobb-Douglas and all differences in productivity are Hicks neutral. A CES production setup with labor augmenting productivity

will affect the misallocation problem in two major ways. First, a low elasticity of substitution increases the cost of misallocation frictions because plants can less readily substitute away from a factor facing frictions. By the same token, plant factor cost ratios will vary less across plants with different levels of frictions. Thus, a lower estimate of the elasticity of substitution will raise both estimates of frictions and the cost associated with them.

Second, differences in labor augmenting productivity will cause dispersion in the capital cost share independent of misallocation frictions. The Hsieh and Klenow model provides a clear illustration of the different implications of a model of misallocation. In their model, each plant faces exogenous capital wedges τ_k and output wedges τ_y and isoelastic demand. Thus, each plant faces the following maximization problem:

$$\pi = (1-\tau_y)PY - wL - (1+\tau_k)rK \tag{6.1}$$

$$Y = AK^{\alpha}L^{1-\alpha} \tag{6.2}$$

Firms that face high capital taxes optimally choose low capital shares of cost as the capital tax discourages them from purchasing capital. Firms with higher capital taxes have a higher marginal cost and so have lower revenue. Firms that face high output taxes also have less revenue as the output tax discourages them from producing more. These firms will also have a lower labor share of revenue, as their output restrictions mean higher prices and so higher revenue per unit produced. Thus, firms with high output taxes will have low revenue and a low labor share of revenue. Hsieh and Klenow identify the frictions in the micro data as follows:

$$1+\tau_k = (\frac{rK}{wL})^{-1}\frac{\alpha}{1-\alpha} \tag{6.3}$$

$$1-\tau_y = \frac{wL}{PY}\frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha} \tag{6.4}$$

Capital taxes are inversely proportional to the factor cost ratio, while output taxes are proportional to the inverse of the labor share of revenue.

A setup with labor augmenting productivity and an elasticity of substitution less than one explains both the factor cost ratio and labor share of revenue with one unobservable: labor augmenting productivity. A setup with labor augmenting productivity has two main different implications from the Hsieh Klenow misallocation model. First, labor augmenting productivity differences would imply that high revenue plants also have a low labor share of revenue. Output taxes in the misallocation model would imply a positive correlation between revenue and the labor share of revenue. Second, plants with a low factor cost ratio (or capital share of cost) should also have a high labor share of revenue. Under the misallocation setup, both of these are proportional to exogenous firctions hitting plants and so have no predicted relationship.

In the US data, I find that high revenue firms have a lower labor share of revenue, with correlations of -.33 in the 1987 unweighted regressions, -.26 in the 1987 weighted regressions, -.33 in the 1997 unweighted regressions, and -.49 in the 1997 weighted regressions after controlling for industry effects. I also find that firms with a high labor cost to capital cost ratio (the inverse of the factor cost ratio) have a high labor share of revenue, wo Firms facing capital taxes would have low revenue, as labor augmenting productivity would predict, but firms facing high output taxes would have lower revenue with correlations of the .29 in 1987 unweighted regressions, .33 in the 1987 weighted regressions, .20 in the 1997 unweighted regressions, and .23 in the 1997 weighted regressions. Thus, labor augmenting productivity B can explain patterns in the data that a misallocation theory can not, while using one less unobservable to explain the data.

So far, I have shown that labor augmenting productivity can generate micro data patterns that would otherwise appear to be misallocation. Labor augmenting productivity can explain the dispersion in wedges that Hsieh and Klenow find in the US plant level data. A key result in Hsieh and Klenow, however, is that China and India have much higher TFP gains from reallocating efficiently. Misallocation is certainly a reasonable explanation for the cross country differences that Hsieh and Klenow find. However, another valid explanation is that non-neutral productivity is more disperse in China and India than in the US.

7 Conclusion

A CES production function with labor augmenting productivity can better explain the plant level manufacturing data than a Cobb Douglas production function. If the elasticity of substitution between labor and capital is less than one, labor augmenting productivity is labor saving. Plants with higher labor augmenting productivity then have higher capital shares. Given downward sloping demand, labor augmenting productivity improvements will increase a firm's revenue and average revenue product of labor, but not its average product of capital. This process induces a positive correlation between revenue and both the capital share and the average revenue product of labor, which I find in the data.

I then identify the labor capital elasticity of substitution using local labor market wage variation. Areas with higher wages have lower factor cost ratios, just as an elasticity less than one would predict. For manufacturing as a whole, I estimate the elasticity of substitution to be between .45 to .65, depending on the level of wage variation and year. When I estimate the elasticity of substitution for SIC 2 digit industries separately, I can reject the Cobb-Douglas specification for 15 out 19 industries with MSA level wages in 1987, with most of the estimates ranging between .4 and .7.

I can directly identify a measure of labor augmenting productivity from the conditions of cost minimization. Labor augmenting productivity is strongly persistent over time and positively correlated with the size of the plant and the level of exports conditional on exporting. Without price data, I can not separate neutral productivity from demand shocks. While a measure of neutral productivity and demand shocks is also highly correlated with plant size, when I use homogenous product data where I can separate the two a measure of neutral productivity alone is not correlated with plant size.

References

ACEMOGLU, D. (forthcoming): "When Does Labor Scarcity Encourage Innovation?," *Journal of Political Economy*.

ANTRAS, P. (2004): "Is the US Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution.," *Contributions to Macroeconomics*, 4(1).

BANERJEE, A., AND E. DUFLO (2005): "Growth theory through the lens of development economics," *Handbook of Economic Growth*, 1.

BASU, S., AND J. FERNALD (1997): "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 105(2).

BERNARD, A., J. EATON, B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93(4).

BERNARD, A., J. B. JENSEN, S. J. REDDING, AND P. SCHOTT (2007): "Firms in International Trade," *Journal of Economic Perspectives*, 21(3).

BLOOM, N., AND J. V. REENEN (2007): "Measuring and Explaining Management Practices Across Firms and Countries," *Quarterly Journal of Economics*, 122(4).

CHIRINKO, R. S. (2008): "Sigma: The long and short of it," *Journal of Macroeconomics*, 30, 671–686.

CHIRINKO, R. S., S. M. FAZZARI, AND A. P. MEYER (2004): "That elusive elasticity: a long-panel approach to estimating the capital-labor substitution elasticity," .

CHRISTENSEN, L. R., D. W. JORGENSON, AND L. J. LAU (1973): "Transcendental Logarithmic Production Frontiers Transcendental Logarithmic Production Frontiers," *Review of Economics and Statistics*, 55(1).

DINARDO, J., AND D. LEE (2004): "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001," *Quarterly Journal of Economics*, 119(4).

DUNNE, T., S. KLIMEK, AND J. SCHMITZ (2010): "Does Competition Spur Productivity? Evidence from U.S. Cement Manufacturing Industry in the Post WWII Period," Mimeo.

FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 98(1), 394–425.

GANDHI, A., S. NAVARRO, AND D. RIVERS (2009): "Identifying Production Functions using Restrictions from Economic Theory," Mimeo.

HICKS, J. (1932): *The Theory of Wages*. MacMillan and Co.

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

HIRSCH, B. T., AND D. A. MACPHERSON (2003): “Union Membership and Coverage Database from the Current Population Survey: Note,” *Industrial and Labor Relations Review*, 56(2).

HOLMES, T. (1998): “The Effects of State Policies on the Location of Industry: Evidence from State Borders,” *Journal of Political Economy*, 106(4), 667–705.

HSIEH, C.-T., AND P. KLENOW (2009): “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 124(4).

JONES, C. (2005): “The Shape of Production Functions and the Direction of Technical Change,” *Quarterly Journal of Economics*, 120(2).

KLUMP, R., P. MCADAM, AND A. WILLMAN (2007): “Factor Substitution and Factor Augmenting Technical Progress in the US,” *Review of Economics and Statistics*, 89(1).

KOENKER, R. (2004): “Quantile regression for longitudinal data,” *Journal of Multivariate Analysis*, 91(1).

LEÓN-LEDESMA, M. A., P. MCADAM, AND A. WILLMAN (2010): “Identifying the Elasticity of Substitution with Biased Technical Change,” *American Economic Review*, 100(4).

LEVINSOHN, J., AND A. PETRIN (2003): “Estimating Production Functions Using Inputs to Control for Unobservables,” *The Review of Economic Studies*, 70(2), 317–341.

MELITZ, M. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71.

MIDRIGAN, V., AND D. XU (2009): “Accounting for Plant-Level Misallocation,” Discussion paper, New York University.

MOLL, B. (2010): “Productivity Losses from Financial Frictions: Can Self-financing Undo Capital Misallocation?,” Discussion paper, Princeton University.

OBERFIELD, E., AND D. RAVAL (2011): “Aggregation, Micro Data, and the Macro Elasticity of Substitution,” .

OLLEY, G. S., AND A. PAKES (1996): “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64(6), 1263–1297.

RESTUCCIA, D., AND R. ROGERSON (2008): “Policy distortions and aggregate productivity with heterogeneous establishments,” *Review of Economic Dynamics*, 11(4).

REVANKAR, N. S. (1971): “A Class of Variable Elasticity of Substitution Production Functions A Class of Variable Elasticity of Substitution Production Functions A Class of Variable Elasticity of Substitution Production Functions A Class of Variable Elasticity of Substitution Production Functions A Class of Variable Elasticity of Substitution Production Functions A Class of Variable Elasticity of Substitution Production Functions,” *Econometrica*, 39(1).

ROBINSON, J. (1933): *The Economics of Imperfect Competition*. MacMillan and Co.

SCHMITZ, J. (2005): "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*, 113(3).