Does Input Quality Drive Measured Differences in Firm Productivity?

Published: *International Economic Review* (2011), 52(4), 961–989.[*]

Jeremy T. Fox
University of Michigan
& NBER

Valérie Smeets
Aarhus School of Business
Aarhus University

November 2011

Abstract

Firms in the same industry can differ in measured productivity by multiples of 3. Griliches (1957) suggests one explanation: the quality of inputs differs across firms. We add labor market history variables such as experience and firm and industry tenure, as well as general human capital measures such as schooling and sex. We also use the wage bill and worker fixed effects. We show adding human capital variables and the wage bill decreases the ratio of the 90th to 10th productivity quantiles from 3.27 to 2.68 across eight Danish manufacturing and service industries. The productivity dispersion decrease is roughly of the same order of magnitude as some competitive effects found in the literature, but input quality measures do not explain most productivity dispersion, despite economically large production function coefficients. We find that the wage bill explains as much dispersion as human capital measures.

JEL codes: D24, L23, M11

Keywords: productivity dispersion, total factor productivity, input quality

[*]Fox thanks the National Science Foundation, the NET Institute, the Olin Foundation, and the Stigler Center for generous funding. Smeets gratefully acknowledges financial support from the Danish Council for Independent Research in Social Sciences and the Marie Curie Program of the European Commission (MEIF-2003-501280). Thanks to helpful discussions with Daniel Ackerberg, Allan Collard-Wexler, Ulrich Doraszelski, Tor Eriksson, Jan de Loecker, Amil Petrin, Chad Syverson, Johannes Van Biesebroeck, and Frédéric Warzynski. We appreciate remarks by seminar participants at Aarhus University, CAED, the Econometric Society, IIOC, Jornadas de Economia Industrial, NYU Stern and Texas A&M. Thanks to the professional staff at the Center for Corporate Performance for hospitality and for help integrating the KØB and IDA datasets. Our email addresses are jtfox@umich.edu and vas@asb.dk.

1 Introduction

Measured differences in productivity across plants in the same industry are usually large. Bartelsman and Doms (2000) survey the literature and find many instances where the highest productivity firm has more than twice the measured productivity of the lowest productivity firm. Dhrymes (1995) studies American manufacturing and finds that the ratio of total factor productivity (TFP) of plants in the ninth decile to the TFP of plants in the second decile is 2.75. We find that the ratio of the 90th quantile of productivity to the 10th quantile of productivity is a mean of 3.27 across eight Danish manufacturing and service industries. For the same inputs, a firm at the 90th quantile of TFP produces 3.27 times the output of a firm at the 10th quantile of TFP.

These huge differences in cross-sectional, measured productivities have spawned a literature investigating why productivity differences are so large. One explanation is simply measurement error in output. However, measured productivity dispersion is similar in developed and developing countries, whereas measurement error might be expected to be larger in developing country datasets (Bartelsman and Doms, 2000). Also, productivity at the firm or plant level is persistent across time, meaning all components of measured productivity cannot be transient (Baily, Hulten and Campbell, 1992). Further, productivity dispersion decreases with competition, as theory predicts (Olley and Pakes, 1996; Syverson, 2004; Bloom and Van Reenen, 2007). The literature shows measured productivity predicts firm growth and firm exit (Baldwin, 1995), export success (Bernard and Jensen, 1995), and even transfers of plants between conglomerate firms (Maksimovic and Phillips, 2001; Schoar, 2002). Foster, Haltiwanger and Syverson (2008) use physical output instead of sales as the measure of output, and show that technological productivity dispersion is often even higher than revenue productivity dispersion. Further, both types of productivity are correlated with firm outcomes such as growth and exit. The consensus in the literature is that productivity dispersion is a real phenomenon with important consequences for economic efficiency and our understanding of how markets with heterogeneous producers operate.

This paper investigates whether failing to account for input quality drives productivity residuals. Economists since at least Griliches (1957) have argued that productivity dispersion reflects the quality of inputs across firms. Economists working with US manufacturing plant data typically measure inputs as the dollar value of physical capital and the number of workers at a firm. Sometimes, employees are separated into production and nonproduction workers. Not surprisingly, labor and capital vary in much greater detail. Two types of machines may have different uses and may not be perfect substitutes and two types of workers may not have the same contributions to firm output.

Input quality seems to us like a fundamentally different explanation for productivity dispersion than some other explanations, such as management competence, economic incentives, business strategy or other difficult-to-measure firm characteristics. Any firm can buy a higher quality machine or hire an abler worker simply by paying more money for the higher quality inputs. If input quality is the reason for productivity, then productivity is really an artifact of a measurement problem. Input markets can be used to reallocate "productivity" across firms: higher quality workers will switch to the firms that pay them the most, for example. Thus, there is no sense that the firm as an organization is playing an important role in productivity dispersion. If some mostly fixed firm characteristic such as business strategy explains productivity, then input markets will be less effective at reallocating or increasing productivity. Instead, a Schumpeterian process of creative destruction, where high productivity firms grow more quickly, may be needed to raise the economy's aggregate productivity. Measuring the role of input quality for productivity dispersion is essential precisely because the optimal policy responses differ depending on whether productivity dispersion is due to input quality or some fixed firm characteristic.

As physical capital is measured in monetary units rather than the number of machines, the quality of capital is likely somewhat better measured than the quality of workers in a typical production function regression, although we discuss literatures on capacity utilization and vintage capital below that provide caveats suggesting that even physical capital has important measurement issues. Because physical capital is measured in a more informative unit than the number of machines, our contribution is to disaggregate the labor input. We use matched employer-employee panel data from Denmark to precisely measure many characteristics of workers at a firm. We merge individual-level data on all Danish residents with firm-level data on value added and physical capital. We then construct firm-level statistics about worker characteristics.

We present productivity regressions with increasingly detailed input quality measures. First, we investigate a simple adjustment where we follow the literature on income inequality and disaggregate the labor inputs into "skilled" (college) and "unskilled" (noncollege) workers. Next, we include two regressions with much more detailed input quality measures. Schooling, sex, total experience and industry tenure proxy for general- or occupation-specific human capital. Tenure at a worker's current firm proxies for firm-specific human capital. Our production function includes a quality-weighting function that transforms firm-level measures of individual worker characteristics into efficiency units of labor. This labor quality function is embedded in the estimation of an otherwise standard Cobb-Douglas production function. The residual from this production function estimate is a firm's total factor productivity (TFP). We examine whether adjusting for labor input quality reduces the measured within-industry dispersion in TFP. We use two different functional forms for labor quality. One specification follows Griliches (1957) and multiplies the contribution of different worker characteristics. The second specification follows Welch (1969) and adds the contribution of each labor quality measure.

The human capital measures that we use may not capture all aspects of worker quality. One approach is to estimate the abilities of workers using wage data. We use panel data on worker wage histories to estimate time invariant worker fixed effects as proxies for worker abilities. We then include these worker fixed effects in production functions as alternative measures of labor quality.

Most researchers do not have access to detailed worker panel data to construct labor market histories or to estimate fixed effects. Therefore, we also investigate using proxies for labor quality that can be obtained from accounting data, such as the wage bill of the firm. We present two specifications with wages: the total wage bill instead of the number of workers, and the fraction of the wage bill spent for various human capital bins. We show that the wage bill reduces productivity dispersion as much as our detailed human capital measures. Including the wage bill is also interesting because human capital characteristics tend to have low explanatory power in wage regressions. Here we show that the wage bill does a little better at predicting output than our human capital measures do; the wage bill may be picking up some unobserved (in our data) input quality.

We present our results separately for two different production functions: Cobb-Douglas and the translog. As firms in different industries use different technologies, we present separate results for eight industries. Also, we present a benchmark for any decline in measured productivity dispersion: the decline in dispersion from adding past employment growth as a control. The previous empirical literature has emphasized that growth is correlated with productivity. Another benchmark compares our productivity declines from adjusting for input quality to the productivity declines from local product-market competition in Syverson (2004).

Our main empirical finding is that our detailed input quality measures, among the best one can hope to obtain, do reduce measured productivity dispersion somewhat, but there is still a large amount of productivity dispersion remaining after controlling for input quality. Averaging across our eight industries, the ratio of the outputs of the firm in the 90th quantile of TFP to the firm in the 10th quantile of TFP is 3.27. This declines to 2.68 with human capital and wage bill controls. 2.68 is 82% of 3.27, meaning that the 90/10 ratio of TFP quantiles declines by 18% from adding the most detailed human capital and wage bill controls. Labor input quality contributes to typical measures of TFP dispersion and explains about 18% of the dispersion.

Our finding of a 18% decline in productivity dispersion is not because human capital measures are unimportant in production. Indeed, for each industry we estimate usually economically large and sometimes statistically precise coefficients on the human capital measures. Rather, our finding is that the reason some firms are dramatically more productive than others is not only a simple failure to account for input quality. While many factors probably contribute to productivity dispersion and input quality is among them, adjusting for input quality does not explain most productivity dispersion. As stated above, any firm could hire, say, more college-educated workers simply by paying the market wage for workers with college degrees. Combining our empirical result with some of the findings from the literature mentioned earlier, the remaining explanations for the unexplained portion of firm productivity dispersion appear to reflect attributes that are hard to buy and sell in input markets. Explanations include managerial competence, business strategy, or some legally protected competitive advantage. Whatever the true relative importances of the various explanations, which the literature is slowly measuring, many of the attributes that determine productivity seems to be hard to define and perhaps hard to buy in a market. While discussing optimal policy is well beyond the scope of our paper, this does suggest product-market competition, rather than relying only on input markets, may be an essential force in raising aggregate productivity.[1]

Most of the productivity literature studies manufacturing industries, perhaps because of data availability. We compare our results for three traditional manufacturing industries (machinery, food and furniture) to three skill-intensive service industries: accounting, advertising and computer services, one less skill-intensive service industry, hotels, and one industry, publishing and printing, which is a hybrid of manufacturing and services.

[1] It is puzzling how low-productivity firms can remain in business at all. One explanation is product differentiation: each firm sells a slightly different product and so heterogeneous consumer demand supports a variety of firms.

Overall, we find declines in productivity dispersion from adjusting for labor input quality that are twice as high in services (and publishing and printing) as in manufacturing, although there is heterogeneity across industries, particularly within the services sector.

1.1 Literature comparison

Several recent papers use both worker data and firm output data, either to compare production and wage regression coefficients (Hellerstein and Neumark, 1999; Hellerstein, Neumark and Troske, 1999; Van Biesebroeck, 2007) or to control for worker ability in wage regressions (Frazer, 2006). We study productivity dispersion and do not compare our production function estimates to wages. Likewise, Haltiwanger, Lane and Spletzer (2007) regress TFP residuals on worker-quality controls using US unemployment-insurance data. They focus on the coefficients on labor quality rather than whether productivity dispersion can be explained by input quality.

Hellerstein and Neumark (2006) do remark on R^2 as a measure of productivity dispersion in one sentence in their paper. In their Table 6, they show that R^2 in a Cobb-Douglas production function estimate without labor quality measures is 0.938 while the R^2 with labor quality controls is 0.940. The base R^2 in their paper is higher than in our results because their dependent variable is total sales and not value added, and they include materials as a separate input. However, the change in R^2 is close to zero in the results of Hellerstein and Neumark. We will find more substantial decreases in productivity dispersion, perhaps because of our richer and more appropriate input quality measures, as we now discuss.

Compared to our data, the controls for labor quality in Hellerstein and Neumark are coarser: they use only the fractions of workers in a high and a low schooling group as well as the fractions assigned to four occupational categories. We use more detailed measures of schooling and use worker panel data to construct measures of human capital based on labor market histories, such as firm and industry tenure and total experience. Indeed, our 21 year worker panel on all citizens in Denmark is a major data advantage. Unlike Hellerstein and Neumark, we do not look at the assignment of workers to occupations within the firm, as we feel that job assignment is an intermediate decision (how to use inputs) rather than a characteristic of the labor inputs. The production function itself models how well firms make intermediate decisions, including job assignment. Finally, we explore alternative measures of labor quality such as using the wage bill and using worker fixed effects from wage regressions, which Hellerstein and Neumark do not consider.

Hellerstein and Neumark study only manufacturing and assume that firms in all industries use the same production function. We estimate production functions separately for each industry and consider industries in both manufacturing and services. Finally, we look at productivity dispersion after correcting for simultaneity and selection bias using the estimator of Olley and Pakes (1996), which Hellerstein and Neumark do partially (they do not control for selection bias) when investigating some issues other than productivity dispersion, which again is not the focus of their paper.

Denison (1962) and Jorgenson, Gollop and Fraumeni (1987) account for demographic change (age, race, sex, schooling) and labor quality (they weight demographic groups by wage rates) when decomposing aggregate productivity growth.[2] This pioneering work contrasts with newer empirical work using firm- or plant-level data, which usually does not control for worker quality.

We investigate whether using the wage bill as a measure of quality-adjusted labor reduces productivity dispersion as much as using detailed human capital measures. This is important in part for data availability reasons, as the total wage bill may be found in plant- or firm-level datasets while detailed human capital measures often are not. The wage bill has been used as the labor input in a number of other studies. For a recent example, see Collard-Wexler (2010). These studies do not compare the use of the wage bill to the use of human capital measures.

We focus on labor input quality, while others have focused on related questions for materials and physical capital. The materials input is typically measured as the monetary value of that input. Ornaghi (2006) suggests that the cost of materials varies across firms and that failing to adjust for heterogeneous input costs will cause measurement error in the materials input (Berndt and Hesse, 1986). If not all measured inputs are being used, then production function estimates will be biased. Yet another literature focuses on the fact that pieces of capital equipment are of different vintages (Whelan, 2002). Machines lose their value due to technological progress. Researchers typically measure physical capital as the monetary value of installed capital. If the monetary value is not accurately adjusted for technological obsolescence, the measure of physical capital will be inadequate.

2 Production, input quality and productivity dispersion

2.1 Production functions

Differences in output across firms can be decomposed into differences in measured inputs, differences in residuals and differences in production technologies. Using data from a single industry and assuming a common technology for all firms, the literature typically estimates the Cobb-Douglas production function

$$\log y = \beta_0 + \beta_l \log l + \beta_k \log k + e, \qquad (1)$$

where y is value added, l is the number of workers, k is the monetary value of physical capital, and e is the residual. β_l and β_k are the input elasticities of labor and capital. Between two firms with the same inputs l and k, the firm with the higher output y is said to have a higher measured total factor productivity (TFP), which is $\exp(\beta_0 + e)$ above. Our measure of output is a firm's value added, which is just total sales minus materials

[2] A related literature studies the dollar value of accumulated human capital in, for example, US states (Mulligan and Sala-I-Martin, 2000).

and other outsourced inputs, such as consulting services.[3] We focus on e, the productivity residual. We call e productivity throughout the paper.[4]

We also report separate results for the translog production function

$$\log y = \beta_0 + \beta_l \log l + \beta_k \log k + \beta_{l,2} (\log l)^2 + \beta_{k,2} (\log k)^2 + \beta_{l,k} (\log l)(\log k) + e, \tag{2}$$

where the second-order terms and the interaction add approximation flexibility (Christensen, Jorgenson and Lau, 1973). While not reported, our conclusions about TFP dispersion are robust to estimating a constant-elasticity-of-substitution (CES) production function.

2.2 Labor quality

2.2.1 College and noncollege workers

There is only limited work on adding input quality measures to firm- and plant-level production functions. Therefore, there is no consensus in the literature in how to incorporate input quality measures. One way is to define new inputs. The empirical literature on income inequality often focuses on "skilled" (workers with a college degree) and "unskilled" (all others) workers. We break the number of workers, l, into $l = l_{\text{college}} + l_{\text{noncollege}}$, where, for example, l_{college} is the number of college-educated workers at a particular firm. We then estimate the Cobb-Douglas production function

$$\log y = \beta_0 + \beta_{\text{college}} \log l_{\text{college}} + \beta_{\text{noncollege}} \log l_{\text{noncollege}} + \beta_k \log k + e.$$

There is a similar generalization of the translog production function to the case of college and noncollege workers,

$$\begin{aligned}\log y = \beta_0 + \beta_{\text{c}} \log l_{\text{c}} + \beta_n \log l_n + \beta_k \log k + \beta_{c,2} (\log l_c)^2 + \beta_{n,2} (\log l_n)^2 + \beta_{k,2} (\log k)^2 + \\ \beta_{c,k} (\log l_c)(\log k) + \beta_{n,k} (\log l_n)(\log k) + \beta_{c,n} (\log l_c)(\log l_n) + e,\end{aligned}$$

[3] Consistent with much of the literature, our production functions model the relationship between output and primary inputs like labor and physical capital. We do not have data on intermediate decisions, such as the use of a performance pay scheme for the workforce. These schemes may indeed raise output, but in production function language they are intermediate decisions that are concentrated out of the production function. The production function gives output conditional on a firm making appropriate choices for intermediate inputs. As we will find a large remaining productivity dispersion when adjusting for labor quality, our results will be consistent with a hypothesis that firms who choose good management practices are more productive.

[4] Like most other papers on productivity, for reasons of data availability the dependent variable y is measured in monetary units. Therefore, it incorporates an unmodeled pricing decision. Foster, Haltiwanger and Syverson (2008) do have price data for a set of industries with homogeneous products and show that dispersion in technological productivity is actually higher than dispersion in the revenue-productivity measures we work with. Katayama, Lu and Tybout (2009) suggest that supply-and-demand analysis may be more appropriate than productivity analysis when a pricing decision affects the dependent variable.

where c and n stand for college and noncollege, respectively.[5]

2.2.2 Human capital measures

Almost every firm in our data has at least one worker with a college degree and one without a college degree. But if there was a firm with no college workers, then $l_{\text{college}} = 0$, $\log l_{\text{college}} = -\infty$ and the firm would produce no output. In a Cobb-Douglas specification, all inputs are essential for production. However, the data show that many firms lack, say, a worker with 3–5 years of tenure at the firm. This means defining $l_{\text{firmtenure},3-5}$ as the number of workers with 3–5 years of tenure and including $l_{\text{firmtenure},3-5}$ as a separate input in a Cobb-Douglas production function contradicts the evidence. Many other types of labor can substitute for those with 3–5 years of tenure; it is not an essential input.

We take several approaches to incorporating more detailed measures of labor quality into the production function. The first approach follows a classic suggestion of Griliches (1957), who in a survey paper put forth mismeasured input quality as a major explanation for productivity dispersion. This approach views the total labor input as the number of workers times labor quality. Each worker is a bundle of measured characteristics. We unbundle workers so that labor quality is a function of the fraction of workers in a firm with each characteristic.[6] In a firm with 100 workers, hiring 1 more woman with a college degree will increase the fraction of workers who are women by 1% and the fraction of workers with college degrees by 1%. Let $x_{\text{female}} = l_{\text{female}}/l$ be the fraction of workers who are women, and $x_{\text{college}} = l_{\text{college}}/l$ the fraction with a college degree. Total labor quality has the multiplicative functional form

$$q_{\theta}^{\text{mult}}(x) = (1 + \theta_{\text{female}} x_{\text{female}})\left(1 + \theta_{\text{college}} x_{\text{college}}\right). \tag{3}$$

Here, efficiency units of labor are the relative productivity compared to a male high-school graduate, say. θ_{female} is how much more productive a woman is than man, and θ_{college} is how much more productive a college-educated worker is than a worker who did not attend college. A firm of all men where 100% of its workers attended college will have a per-worker quality of $1 + \theta_{\text{college}}$.[7]

Labor quality is not additively separable across workers. For example, expanding the specification of $q_{\theta}(x)$ above produces the interaction term $\theta_{\text{female}} x_{\text{female}} \theta_{\text{college}} x_{\text{college}}$. If the θ's are positive, adding a male college graduate will produce a greater increase in labor quality at a firm with more women. By contrast, Welch (1969) emphasizes a production technology where human capital attributes are additive. Therefore, our next functional

[5] In a Cobb-Douglas production function, college and non-college workers are complementary inputs: production cannot take place without both inputs. However, there is not a formal model of hierarchical or team production, where perhaps college workers supervise noncollege workers. The flexible translog specification may provide a better approximation to a hierarchical production function.

[6] An exception is total labor market experience, which enters the labor quality function as a continuous variable: the mean level of experience in the firm. The data appendix discusses some topcoding reasons why some other variables enter as fractions of the workforce. There is nothing about our production functions that prevents us from choosing continuous or discrete variables, as appropriate.

[7] A multiplicative labor-quality measure is also used in Hellerstein and Neumark (2006) and Van Biesebroeck (2007).

form for labor quality takes an additive functional form

$$q_{\theta}^{\text{add}}(x) = 1 + \theta_{\text{female}} x_{\text{female}} + \theta_{\text{college}} x_{\text{college}}. \tag{4}$$

Our results about productivity dispersion will be relatively consistent across $q_{\theta}^{\text{mult}}(x)$ and $q_{\theta}^{\text{add}}(x)$.

Let the total number of workers at a firm be l. The total labor input is then $l \cdot q_{\theta}(x)$. Substituting this expression for labor in the Cobb-Douglas production function (1) gives the estimating equation

$$\log y = \beta_0 + \beta_l \log(l \cdot q_{\theta}(x)) + \beta_k \log k + e. \tag{5}$$

The parameters θ in the labor quality function enter this equation nonlinearly, so estimation is by nonlinear least squares. This requires using a numerical optimization procedure to minimize the least squares objective function, as there is no closed-form solution for the least squares estimator of θ. We also estimate a version of the translog production function, (2), with quality-adjusted labor $l \cdot q_{\theta}(x)$ replacing the total number of workers, l, as in

$$\log y = \beta_0 + \beta_l \log(l \cdot q_{\theta}(x)) + \beta_k \log k + \beta_{l,2} (\log(l \cdot q_{\theta}(x)))^2 + \beta_{k,2} (\log k)^2 + \beta_{l,k} (\log(l \cdot q_{\theta}(x))) (\log k) + e. \tag{6}$$

The same parameters θ appear in multiple places in the production function.

2.2.3 Wage bill as a proxy for labor quality

Another approach to adjusting for labor quality is to use the wage bill as a measure of the quality of the workforce. Wages will reflect marginal products in a competitive labor market. Even if the labor market is not perfectly competitive, wages are still likely highly correlated with worker ability. Just as physical capital is measured in terms of monetary units to at least partially reflect the quality of the machinery employed, labor can be measured in terms of its expense in order to reflect its quality. Using the wage bill instead of the number of workers thus makes the methods of measuring physical capital and human capital more symmetric.

The wage bill may also be more commonly found in the type of data used in firm- and plant-level productivity studies. The total wage bill may be part of some accounting-based firm-level datasets where data on the characteristics of the workers are not available. If results from using the wage bill as the labor input are similar to those using detailed labor characteristics, then it will ease the burden of data collection for those wanting to control for labor quality.

The wage bill specification is also attractive because the explanatory power of human capital variables in wage regressions can be low, suggesting unmeasured worker characteristics are also important determinants of labor quality. Further, the wage bill using monthly salaries better weights the contributions of part-time and full-time workers than do measures like the number of workers.

Our production function with the wage bill is

$$\log y = \beta_0 + \beta_l \log w + \beta_k \log k + e,$$

where the wage bill $w = \sum_{i=1}^{l} w_i$ is the total of the monthly salaries paid to all workers. We also estimate a translog production function, with the wage bill w replacing the number of workers l in (2).[8]

Even if the choice of human capital inputs is randomly assigned to firms with heterogeneous productivities, adding the wage bill could introduce an endogeneity problem if more productive firms pay higher salaries for any of several reasons, including profit sharing. We address simultaneity bias in section 2.5.

2.2.4 Combining the wage bill and the human capital measures

We also combine the wage bill and human capital variables to attempt to account for input quality in as detailed a manner as possible. We use the wage bill w instead of the number of workers l as our base labor input. Then we construct a labor quality adjustment that uses, in part, the human capital measures. Keeping the same human capital categories as before, we calculate the total of the monthly wages for workers in each bin and then normalize by the total wage bill of the firm. For example, $\tilde{w}_{\text{female}} = (w)^{-1} \sum_{i=1}^{l_{\text{female}}} w_{i,\text{female}}$ is the fraction of the firm's wage bill that is paid to women. This is a similar measure to x_{female} above, as it represents the fraction of firm labor inputs coming from women. The difference with x_{female} is that the base unit for counting labor inputs is the total of the monthly wages, rather than the number of workers. We then adapt the Griliches (1957) multiplicative-quality-adjustment term, (5), to give

$$q_\theta^{\text{mult,wage}}(\tilde{w}) = (1 + \theta_{\text{female}} \tilde{w}_{\text{female}}) \left(1 + \theta_{\text{college}} \tilde{w}_{\text{college}}\right), \tag{7}$$

where $\tilde{w}$ is the vector of wage bill fractions for the different human capital categories. We then estimate (5) using nonlinear least squares, with the labor quality term $q_\theta^{\text{mult,wage}}(\tilde{w})$ multiplying the total wage bill w. The regression equation is

$$\log y = \beta_0 + \beta_l \log \left(w \cdot q_\theta^{\text{mult,wage}}(\tilde{w})\right) + \beta_k \log k + e.$$

There is also a translog specification equivalent to (6).

2.2.5 Worker fixed effects from wage regressions

Our detailed human capital measures and the wage bill are measured in the data. Both measures have potential drawbacks. Human capital measures such as experience and schooling may do a poor job of accounting for skills such as innate ability, the ability to work as part of a team and leadership skills. Likewise, the wage

[8] Value added may be formed from sales by subtracting materials costs but not the wage bill. Thus, the wage bill does not, in an accounting sense, enter the calculation of value added.

bill reflects not only labor input quality but also the compensation policy of each firm. The wage bill can vary across firms for reasons not related to input quality, including efficiency wages, rent sharing, and workers' bargaining power.

We turn to worker panel data in order to address unmeasured (in the case of human capital measures) and improperly measured (in the case of the wage bill) input qualities. In our panel data on all Danish citizen, we observe the same worker over many years and potentially many employers. Under the key assumption that worker ability α_i is time invariant, we estimate a worker i fixed effect α_i using the panel data wage regression

$$\log w_{i,t} = \alpha_i + x_{i,t}'\gamma + \varepsilon_{i,t},$$

where t indexes years, $x_{i,t}$ includes the time varying human capital measures of industry tenure, firm tenure, total labor market experience and worker age, γ is a vector of estimated slope parameters on the human capital measures in $x_{i,t}$, and $\varepsilon_{i,t}$ is an time- and worker-specific wage disturbance. We use worker data from all workers employed in the private sector for the years 1992–2001.[9] We estimate the parameters γ using the well known within estimator, and then use the estimated coefficients to back out the implied estimates for the worker fixed effects α_i. The estimates of γ are consistent as the number of workers grows large, while the estimates of each α_i are consistent as the number of years recorded for each worker grows large. We use data on 2 million workers who are observed an average of 6.4 years. We do not use wage observations for years in which the worker is employed in the public sector. A mean of 6.4 observations per worker may be low, but we mainly use firm-level averages of the un-logged $\exp(\alpha_i)$, in which estimation error at the worker level averages out to some degree. The sum of $\exp(\alpha_i)$ at the firm level (for those workers at the firm in year t) is our main measure for the labor input that adjusts for unobserved characteristics of workers.

If the true wage model has a fixed effect $\phi_{j(i,t)}$ for the employer $j(i,t)$ of worker i in year t, our estimate is consistent if there is no correlation between α_i and $\phi_{j(i,t)}$. The typical result in the literature that estimates both worker and firm fixed effects is that the correlation between α_i and $\phi_{j(i,t)}$ is indeed close to zero (Abowd, Kramarz and Margolis, 1999). If, as expected based on prior empirical evidence from a variety of countries, the correlation between α_i and $\phi_{j(i,t)}$ is close to 0 in Denmark, then we lose little by not including firm fixed effects.[10]

We also estimate a production function where we combine human capital and wage fixed effect data, as in specification (7) for the wage bill. Return to the example of gender. Instead of calculating $\tilde{w}_{\text{female}} = (w)^{-1}\sum_{i=1}^{l_{\text{female}}} w_{i,\text{female}}$ as the fraction of the firm's wage bill that is paid to women, we now work with $\tilde{\alpha}_{\text{female}} = \left(\sum_{i=1}^{l}\exp(\alpha_i)\right)^{-1}\sum_{i=1}^{l_{\text{female}}}\exp(\alpha_{i,\text{female}})$ as the percentage of total firm worker fixed effects that are paid to

[9] We do not use the full worker panel because of the limitations on the hardware and software available in the protected data servers at Statistics Denmark.

[10] The economic assumptions motivating wage regressions with both worker and firm fixed effects have been challenged using both frictionless and search theoretic models of matching and wage formation (Dupuy, 2010; Eeckhout and Kircher, 2009; Lise, Meghir and Robin, 2010; Lopes de Melo, 2009). Our use of worker fixed effects is just one exercise out of many and we do not want to defend it against these theoretical critiques.

women as a measure of the labor intensity and quality of women.

2.2.6 Within-firm human capital dispersion as in Iranzo et al. (2008)

Iranzo, Schivardi and Tosetti (2008) estimate a production function where an additional input is the standard deviation of some measure of human capital. The idea is that the standard deviation is part of an approximation to a production function where workers of higher ability supervise workers of lower ability. Iranzo et al. use Italian data and find that the standard deviation of worker wage fixed effects α_i, at the firm level, contributes positively to firm output. We have estimated production functions using the within-firm dispersion in standard deviation of various labor quality measures as inputs. Our focus is not on the input elasticities of these within-firm dispersion measures but on how much adding them reduces measured across-firm productivity dispersion. We used both the standard deviation of wage levels and the standard deviation of wage fixed effects, for all of our industries. We find that the increase in R^2 (reduction in productivity dispersion) from adding within-firm labor quality dispersion as an additional input is always well less than 0.01. For this reason, we do not further report the estimates for the production function specifications inspired by Iranzo et al..

2.3 Productivity dispersion

Although we do discuss estimates of production function parameters such as β_0, β_l, β_k and θ, our primary focus is on total factor productivity, or the residual e in (1). The parameters such as θ can be economically large and statistically significant despite the dispersion in e, the key puzzle to understand about productivity, remaining large. We focus on several related measures of the dispersion of e.

Productivity dispersion is intimately related to $R^2 = 1 - \frac{\text{Var}(e)}{\text{Var}(\log y)}$. Maximizing R^2 is the least-squares criterion. One attempt to explain productivity dispersion is to add observables to the model to see how much residual productivity dispersion declines. If a new variable reduces productivity dispersion, it will also increase the statistical fit of the regression. The change in statistical fit, R^2, from adding a single new regressor z to the production function (1) estimated by ordinary least squares is

$$\Delta R^2 = \left(1 - R^2_{\text{base}}\right) \left(\text{partialcorr}\left(\log y, z \mid \log l, \log k\right)\right)^2, \tag{8}$$

where $\text{partialcorr}\left(\log y, z \mid \log l, \log k\right)$ is the partial correlation between output $\log y$ and the new input z once the non-quality adjusted inputs, $\log l$ and $\log k$, are controlled for. To compute a partial correlation, one separately regresses $\log y$ and z on $\log l$ and $\log k$ and then forms the simple correlation of the residuals from the $\log y$ and z regressions. Equation (8) indicates that a variable will add a lot of explanatory power to a regression if it is correlated with the dependent variable but is not so correlated with the other independent variables.[11]

[11] The R^2 from nonlinear least squares (NLS) is not guaranteed to be between 0 and 1 (the derivation for OLS uses the first order

To examine the units of productivity rather than a statistical fit criterion, we also report the standard deviation of e, which enters the numerator of R^2. Both the standard deviation of e and R^2 use the logged instead of the unlogged levels of productivity. Our preferred measure of productivity dispersion in unlogged levels is q_{90}/q_{10}, where q_{90} is the 90th quantile of TFP in levels $\exp(e)$ and, likewise, q_{10} is the 10th quantile of $\exp(e)$. q_{90}/q_{10} is the ratio of outputs for the 90th quantile and 10th quantile firms, if those firms had used the same inputs. We also report the ratio q_{75}/q_{25}, which is less sensitive to outliers than q_{90}/q_{10}. Keep in mind that a productivity quantile such as q_{90} does not necessarily move monotonically with a moment such as the standard deviation of e.

2.4 Productivity dispersion decline benchmarks

There is no absolute metric for whether any given decline in productivity dispersion is large or small. First, we benchmark the productivity declines from adding human capital measures against the decline in productivity dispersion from adding firm growth. Baldwin (1995) and others show that firms that are more productive will on average have higher rates of employment growth.[12] Because of the prior literature relating productivity to employment growth, there are a priori reasons to suspect that including employment growth will decrease productivity dispersion substantially.

We add firm growth as an observed components of productivity, as in

$$\log y = \beta_0 + \beta_l \log l + \beta_k \log k + \beta_{\text{DHgrowth}} r_{\text{DHgrowth}} + e.$$

We use the Davis and Haltiwanger (1992) measure of firm-employment growth, which ranges from -2 to 2, instead of -1 to 1, in order to account for firm entry and exit. We see how much residual productivity dispersion declines after accounting for growth. This decline provides a benchmark for the decline in dispersion from controlling for input quality.

Another approach for benchmarking compares our decline in productivity dispersion to another decline that has shown to be important in the literature. Syverson (2004) regresses productivity dispersion (the interquartile range) in a local geographic market on a measure of the demand density (a proxy for product market competition) in that market. He finds that a "one-standard-deviation increase in logged demand density implies a decrease in expected productivity dispersion by approximately 0.042 log points—roughly one-seventh of the mean dispersion and over one-fourth of its standard deviation." Syverson's measure of productivity dispersion is the interquartile range of log TFP. We will compare our productivity dispersion declines from adjusting for input quality to those from Syverson from varying local product-market competition.

---

conditions of OLS to set the sample covariance of the residual and the predicted dependent variable to 0.) We define R^2 for NLS to be $1 - \frac{\text{Var}(e)}{\text{Var}(\log y)}$.

[12] In various specifications we control for firm age. Cabral and Mata (2003) and others show that older firms tend to be more productive. We use the log of firm age as firm age can have some extreme outliers (hundreds of years old) in Denmark.

2.5 Simultaneity and selection bias

Marschak and Andrews (1944) introduce the endogeneity concern that more productive firms may use more inputs, leading to overestimating the input elasticities. Griliches and Mairesse (1998) argue that traditional methods of correcting regressions for endogeneity, panel data and instrumental variables, work poorly for production function estimation because of measurement error (panel data) and data availability (instruments). Following the recent literature, we use investment to correct for input endogeneity using the Olley and Pakes (1996) estimator. The Olley and Pakes estimator also corrects for selection bias, the problem that low productivity firms may exit the sample using an exit rule where included right-hand side regressors (state variables in a dynamic program) such as physical capital and firm age are taken into account. This endogenous exit induces correlation between physical capital and productivity as well as firm age and productivity among the remaining firms.

We present production function estimates using the Olley and Pakes estimator because this is the most commonly used procedure in the recent literature, in part because the estimates of input elasticities from this procedure often appear more a priori plausible than those from alternatives, such as the dynamic panel methods in Arellano and Bond (1991) and Blundell and Bond (2000). For example, Ornaghi (2006) uses a dynamic panel estimator for manufacturing industries and finds negative coefficients on physical capital. One concern is that dynamic panel methods difference or quasi-difference the data across time, which magnifies the statistical bias from the transient components of measurement error. As our paper emphasizes that typical measurements of inputs may be flawed, we do not wish to use a dynamic panel method that magnifies the bias from transient measurement error.

The Olley and Pakes model theoretically decomposes e into true productivity ω and measurement error η. When we refer to the Olley and Pakes productivity term, we refer to the sum e and not either ω or η separately, which follows the practice in the original Olley and Pakes paper on the telecommunications equipment industry. The dispersion in e from the Olley and Pakes estimator will often be higher than the dispersion in e from least squares. The Olley and Pakes estimator allows ω to be correlated with the inputs l and k, to address simultaneity bias and selection bias.[13] A common empirical result is that some of the explanatory power of the measured inputs l and k found using a least squares estimator is transferred to the ω term in the Olley and Pakes estimator. More dispersion in ω increases the dispersion of e relative to the least squares case if the dispersion of ω is not completely offset with a decline in the dispersion of η. When we report R^2 for the Olley and Pakes estimator, we report $1-\frac{\mathrm{Var}(e)}{\mathrm{Var}(\log y)}$, where again $e=\eta+\omega$.

A key assumption we make in order to use the Olley and Pakes estimator is that, following the original paper, the choice of labor inputs is static, i.e. last period's labor choices are not a state variable in the dynamic investment decision. In general, treating labor as a static input may be implausible if there are any hiring or

[13] The Olley and Pakes estimator addresses only correlation of inputs and firm age with ω, not η. To the extent our empirical results using overidentification tests in Fox and Smeets (2010) suggest that the method does not accurately decompose e into true productivity ω and measurement error η, the Olley and Pakes estimator will still be biased, although perhaps less biased than least squares.

firing costs. Modeling labor as static is particularly worrisome in our investigation, as we use labor market history variables such as firm tenure that explicitly have dynamic components to them. Ackerberg, Caves and Frazer (2007b) introduce a new estimator that may be consistent when the labor variable is a dynamic input. We have experimented with the first stage of the Ackerberg et al. estimator for the simple specification of using the log of the number of college and the log of the number of non-college workers as separate, dynamic inputs. The first stage identifies the transient productivity component η and we find that the dispersion in η is nearly identical to the estimates from the first stage of the Olley and Pakes estimator. This leads us to believe that the dispersion in η will be relatively constant even if labor is modeled as a dynamic input. We do not use the Ackerberg et al. estimator in our reported results because, like Olley and Pakes, the method requires estimating a nonparametric function of as many arguments as the number of dynamic state variables. For our detailed human capital specifications, we would estimate a nonparametric function of 17 inputs. High levels of bias and variance will then arise from the well-known curse of dimensionality in the estimates of infinite-dimensional objects such as functions. High bias and variance would occur in any dataset normally used in productivity analysis and especially our relatively small Danish industries.

Even our input quality measures are imperfect. A standard matching model suggests that inputs and firms should assortatively match, if firm productivity and input quality are complements. High ability workers should be at firms with high productivities. If so, a standard omitted variable bias story suggests that the parameter estimates on the human capital variables should be biased upwards: there is a positive correlation between human capital and the true error term, productivity. Recall that equation (8) suggests that the decline in productivity dispersion from adding a variable to a regression involves the partial correlation of the new regressor with the dependent variable. If assortative matching between firms and workers increases this partial correlation, the decline in productivity dispersion from adding human capital variables will be overstated. Therefore, this bias in the parameter estimates works against finding that the decline in productivity dispersion is large. We will not emphasize that our empirical findings of declines in productivity dispersion are especially large.

3 Data overview

We start with detailed panel data on all Danish citizens for 1980–2001. These data provide us general human capital (experience, schooling), firm-specific and industry-specific human capital (firm tenure, industry tenure) as well as the monthly salary for each worker. We are careful with measuring firm tenure because of changes in firm identification codes. The underlying data for these variables come from government records and not subjective self-reports, like in US publicly available microdata. Thus, we feel that our data on worker characteristics are of higher quality than any found in the United States. We aggregate our human capital measures to the firm level to construct our labor quality measures, as in (3) and (7). We also compute the total number of workers as well as several wage bill measures. We estimate worker wage fixed effects and aggregate those to the firm level as well.

We then merge the firm-level human capital measures with data on value added, physical capital and investment.[14] These data come from a credit rating agency, for the years 1992–2001. To use the Olley and Pakes (1996) estimator, we need data on firms with investment data for two consecutive years, which causes a lot of particularly small firms to be dropped. More details on the data are found in the appendix.

Denmark is a small open economy, so there are not many distinct firms in narrowly defined industries. We strive to balance the competing needs to have more observations for precise statistical inference and to allow heterogeneity in the production functions for firms in different industries. We consider a medium level of aggregation because we include many detailed measures of human capital variables and therefore need a lot of observations per regression. We perform separate regressions for eight industries: machinery, food and beverages, furniture, publishing and printing, hotels and restaurants, accounting, computer services and advertising. To alleviate some forms of within-sector heterogeneity, we include fixed effects at the five digit industry level in each regression.

Table 1 lists summary statistics for four of our eight industries: those with production function estimates in Tables 2–5. Value added and inputs vary a lot across firms. Importantly, the human capital measures vary a lot across firms. As there is variation across firms in the sample, equation (8) suggests it is a priori possible that adding human capital quality measures to a production function will increase the R^2 and hence reduce the dispersion in measured productivity.

4 Production function estimates

The paper's focus is on productivity dispersion, which arises from the dispersion of the residuals from production function regressions. Before discussing productivity dispersion in detail, we will describe the production function estimates in order to provide context. As a warning, we do not feel that the production function parameter estimates are robust empirical findings. Our results on productivity dispersion are robust across functional form choices. In this section, we report only results from the Olley and Pakes (1996) estimator that controls for simultaneity bias and selection bias.[15]

Tables 2–5 report estimates of production functions for four industries. Table 2 covers the food and beverages industry. Column 1 is a base specification, with just the number of workers for the labor input. The coefficient on labor is 0.69 and the coefficient on physical capital is 0.18, resulting in an estimate of a decreasing return to

[14] We do not observe measures of inputs other than labor and physical capital and we do not observe sales for many firms.

[15] In the first stage of Olley and Pakes (1996), we allow for a third-order polynomial, including interactions, in investment, physical capital, and firm age. All terms are specific to the individual years of firm-level data, 1992–2001. Our survival equation is a probit with a third-order polynomial, including interactions, in investment, physical capital, and firm age. Our second stage involves a third-order polynomial, including interactions, in the survival probability and the model's predicted productivity level for the previous time period. See the original Olley and Pakes article or Ackerberg, Benkard, Berry and Pakes (2007a) for explanations of the various steps. All of our standard errors allow for heteroskedasticity and use clustering to allow for autocorrelation at the firm level.

scale.[16] The R^2 from the base regression is 0.801.[17]

In Column 2, we begin to account for labor quality. Column 2 uses the numbers of college and non-college workers as separate inputs. The coefficient on the number of skilled workers is 0.20, and the coefficient on the unskilled workers is 0.54. The coefficient on physical capital declines to 0.14. R^2 increases by only 0.006, from 0.801 to 0.807.

Column 3 shows the estimates from (5) with the multiplicative / Griliches (1957) labor quality term, (3). The coefficient on female is -0.343, which can be interpreted as saying that a firm with 10% more of its workforce being women will have $1-0.343 \cdot 0.10 = 0.966$ or 97% of the total labor inputs $l \cdot q_\theta(x)$ as another firm with the same number of workers, l. Schooling is one of our main measures of general human capital. The coefficient of 1.48 on the fraction of college-educated workers says that a firm with 10% higher fraction of college-educated workers (as opposed to the excluded category, workers who completed high school or below) will have 15% more labor inputs. The coefficient is marginally statistically significant at conventional sizes. The coefficients on the fraction of workers with community college and vocational degrees are not statistically distinct from 0, although both point estimates are economically large. The estimates show that a firm with 10% more workers with community college degrees will have 7% more labor inputs.

One of our data advantages is that we can construct detailed labor history measures using our worker panel data. Total experience in the labor market is exactly computed at the worker level from government records (since 1964). With little concern about topcoding for experience, we enter experience as the mean level of experience of workers at the firm, mostly to save space in the tables. A firm whose workforce has an extra 10 years of labor-market experience will have 45% more labor inputs.

We next look at firm tenure and industry tenure in column 3. The tenure measures approximate firm- and occupational-specific human capital. The measures are the percentage of workers in each tenure bin, and all coefficients should be evaluated relative to the residual category, newcomers with 0 years of tenure. We find a firm with 10% more workers with 1–2 years of tenure instead of newcomers has $0.10 \cdot 0.938 = 9.4\%$ more labor inputs, a potentially large effect. All four firm tenure categories have positive coefficients. Because of the large standard errors, the coefficients for the firm tenure categories are mostly consistent with a large, one-time training cost for newcomers.[18] Two of the four industry tenure coefficients are negative, and one of those is statistically distinct from 0 at the 90% level. The R^2 from the multiplicative-labor-quality specification

[16] As the dependent variable is sales and not physical output, Klette and Griliches (1996) suggest that the returns to scale will be biased downwards. This bias could be offset by other biases such as the usual bias that more productive firms use more inputs, which tends to bias the returns to scale upwards.

[17] The Olley and Pakes productivity components η and ω have standard deviations of 0.48 and 0.40, respectively. Under the interpretation from Olley and Pakes's model, the dispersion in measurement error η is larger than the dispersion in ω. We will not discuss the decomposition of e into η and ω further, although Tables 2–5 list such figures for those interested in the decomposition.

[18] A potential "training cost" pattern of coefficients may also reflect a measurement issue: workers hired during the year at a growing firm who are mistakenly counted as working the entire year. This is an issue for growing firms and not firms with simply higher levels of turnover. We reran the labor quality specification by adding the past 5-year firm employment growth using the Davis and Haltiwanger (1992) measure and the extra regressor increases the magnitude of the firm tenure coefficients, which goes against the growing firms explanation.

is 0.818. Overall, we have a couple of statistically significant coefficients and many coefficients with economically large magnitudes. Our finding of a relatively small decrease in productivity dispersion from labor quality controls will not be due to economically small or statistically insignificant estimates of human capital production function parameters.

Column 4 of Table 2 uses the Welch (1969) additive labor quality function. While the coefficients are not directly comparable in magnitude to those using the multiplicative specification in column 4, several of the coefficients do change sign. For example, one of the industry tenure coefficients changes from negative to positive. Because the estimated signs of the labor quality coefficients are sensitive to the functional form for the labor quality function, which occurs in more industries than just food, we do not view the signs of the point estimates of the labor quality terms as robust findings. Interpretations of the parameters require a convincing argument that the labor inputs are uncorrelated with the true error term, productivity. Studies that do not correct for endogeneity argue that more productive firms employ higher-quality workers (Haltiwanger et al., 2007). Despite adopting the advice of Griliches and Mairesse (1998) to use the Olley and Pakes (1996) estimator, we do not take a strong stand that our estimates of the human capital parameters are causal production function estimates. Our main focus is on the dispersion of the productivity residual e, which from experimentation we believe is relatively robust to reasonable biases on the production function parameters.

The R^2 from the additive labor specification in column 4 is 0.800. This is smaller than the R^2 of 0.815 from the multiplicative specification. The difference in R^2 is an artifact of the Olley and Pakes procedure, where the explanatory power of the ω term often varies across specifications. The standard deviation of the η term is nearly identical across columns 3 and 4; only the ω term's standard deviation changes. Using nonlinear least squares and not Olley and Pakes, the R^2 from the multiplicative and additive labor quality specifications are often nearly identical, as Tables 6 and 7 will show.

Column 5 attempts to adjust for the quality of the workforce by using the wage bill. The coefficient on the total wage bill at 0.83 is higher than the corresponding input elasticity when the number of workers is used instead, in column 1. The coefficient on physical capital decreases. The R^2 increases to 0.848, which is higher than the R^2 for the specifications with the detailed human capital measures.

Column 6 uses the specification that combines human capital and wage bill data, (7). Labor quality uses the multiplicative form. The R^2 increases to 0.870, which is to be expected for the specification using the most data. Most of the coefficients are not statistically distinct from 0 at the 95% level.

Column 7 uses the sum of the (unlogged) worker-specific wage fixed effects α_i as the labor input. This measure accounts for both the unmeasured aspects of human capital and the removes some of the influence of firm-specific compensation schemes that may be present in the wage bill. The coefficient on labor drops and the coefficient on physical capital increases. Explanations could include properly accounting for labor quality and adding measurement error to labor quality because of statistical sampling error in the estimated fixed effects. We use a mean of 6.4 observations per worker, so sampling error in the individual fixed effects could be an issue, although we hope aggregating to the firm level ameliorates the estimation error to some degree. R^2 drops noticeably, to 0.772. Worker fixed effects by themselves explain less output variation across firms than does the number of workers.

Column 8 combines the worker-specific wage fixed effects and the detailed human capital measures, where, for example, $\tilde{\alpha}_{\text{female}} = \left(\sum_{i=1}^{l} \exp(\alpha_i)\right)^{-1} \sum_{i=1}^{l_{\text{female}}} \exp(\alpha_{i,\text{female}})$ replaces $\tilde{w}_{\text{female}}$ in (7). The R^2 of 0.772 from column 7 increases to 0.795. With our data, worker fixed effects do not seem to predict variation in output across firms as well as other methods do.

Column 9 is a benchmark regression. Some of the productivity literature finds that firms that are older and that firms that have recently grown quickly are more productive. We have a direct measure of firm age, although it is included in all specifications following Olley and Pakes, and can construct the past five years of firm employment growth, using the Davis and Haltiwanger (1992) measure. We include the growth measure as an extra regressor in a standard Cobb-Douglas regression, (1). In column 9, we find that past firm growth is not very predictive of firm output in the food and beverages industry. The R^2 does not detectably increase over the base case in column 1 and the coefficient on firm growth is not statistically distinct from 0.

Table 3 reports the same set of nine production function estimates for the furniture industry. Column 1 shows that the furniture sector is labor intensive with a labor input elasticity of 0.82 and a capital coefficient of 0.10. The R^2 with no labor-quality measures is 0.727, which increases to 0.736 by adding college and non-college workers as separate inputs. The R^2 with the multiplicative labor quality function is 0.748 and the additive labor quality function has an R^2 of 0.747. In this instance, the Olley and Pakes R^2 is robust across the choice of labor quality functional forms. The R^2 of the pure wage bill specification is 0.785, higher than the R^2 from the human capital specifications. The specification in column 6 that combines wage bill and human capital data has an R^2 of 0.799. The worker-specific wage fixed effects give an R^2 of 0.701 in column 7 and 0.762 in column 8, again showing the weakness of estimated fixed effects in predicting firm output. The R^2 of the firm growth benchmark is column 9 is quite low, at 0.727.

Table 4 looks at the publishing and printing industry. For publishing, column 1 finds an estimate of a small increasing returns to scale. The R^2 from the baseline specification in column 1 of 0.810 increases to 0.865 with the specification with both human capital and wage bill measures in column 6. The sum of the input elasticities also increases. The coefficients on schooling in the human capital specification in column 3 show that college and community college degrees in particular raise firm-level output in publishing.

Table 5 looks at a skill-intensive, service-sector industry, advertising.[19] Column 1 shows that the input elasticity on the number of workers is 0.94 while the input elasticity on physical capital is -0.07, although it is not statistically distinct from 0. It is unsurprising that labor is the dominant input in advertising, although the

[19]Following Olley and Pakes (1996), in most industries we include the logarithm of firm age as a shifter of firm productivity and a state variable in the investment and exit decision rules. We did so initially in advertising as well. However, the coefficient on firm age was estimated to be so large that R^2 was estimated to be negative. Given the a priori implausible parameter estimates, we dropped the state variable firm age for advertising under the Olley and Pakes estimator. We also drop firm age for accounting and hotels when using Olley and Pakes.

negative point estimate for the input elasticity for physical capital appears anomalous. We say advertising is skill-intensive in part because the coefficient on college schooling in column 3 is a large 2.1. The coefficient on college schooling is also a large 33.6 in the additive labor quality specification in column 4. The R^2 from column 1 is a relatively small 0.567 and it increases to 0.825 with the column 6 specification that combines human capital and wage bill data. This increase in statistical fit exceeds that in the other industries we have looked at, which is a point we will return to.

For conciseness, we do not report the parameter estimates for the translog production functions or for the other four industries. Nor do we report point estimates for the specifications using least squares and not Olley and Pakes.

5 Productivity dispersion and input quality

Tables 6–8 are the main results of the paper. The tables report the R^2; the standard deviation of log TFP, e; our measure of productivity dispersion in levels of output instead of logs, q_{90}/q_{10}; and the similar ratio q_{75}/q_{25}, which is less sensitive to outliers. For each of our eight industries, the first row is a baseline specification with the usual measure of labor, the number of workers. The second row uses simple labor quality measures previously found in the literature: the (log) number of workers with college degrees and the (log) number of workers without college degrees, as separate inputs. The third row presents the estimates of (5) using the detailed measures of general and specific human capital. The fourth row tests the robustness of the findings on productivity to the choice of functional forms for labor quality by using an additive instead of a multiplicative specification.

The fifth row replaces the number of workers with the wage bill. Wages may proxy for worker labor quality in a competitive labor market. The sixth row uses both the wage bill and human capital measures. This specification typically is the one that maximizes statistical fit and hence minimizes productivity dispersion.

The seventh row uses the worker-specific wage fixed effects as the measure of labor quality. The eighth specification combines the worker fixed effects with human capital measures. Finally, the ninth row is a benchmark, where we use data on recent growth in firm employment. All specifications include fixed effects for five digit sub-industries and for years.

Tables 6 shows linear and nonlinear least squares estimates of productivity dispersion measures for the Cobb-Douglas production function without the Olley and Pakes (1996) correction. Table 7 reports linear and nonlinear least squares estimates of the translog production function without the Olley and Pakes correction. Table 8 reports estimates for the Cobb-Douglas production function using the Olley and Pakes correction. We discuss the Olley and Pakes results last because some of the results are slightly unintuitive as explanatory power is shifted from the measured inputs l and k to the ω component of total productivity e. The ω component of e

is allowed to be correlated with the inputs l and k, something that is ruled out by assumption in least squares estimators that try only to maximize statistical fit.

Table 6A contains the least squares estimates of productivity dispersion measures for the Cobb-Douglas production function for the four manufacturing industries. Consider machinery. Including only the number of workers gives a R^2 from (1) of 0.834. Using our preferred measure, the ratio of the 90th quantile of $\exp(e)$ to the 10th quantile of unlogged TFP is 3.02. A firm at the 90th quantile produces 3.02 times the output as a firm at the 10th quantile, for the same inputs. We also report the ratio q_{75}/q_{25} in the table, which is less sensitive to outliers, although for conciseness we do not discuss the q_{75}/q_{25} measure in the text. We also will not discuss R^2 or the standard deviation of e, again for conciseness.

Continuing with the machinery industry, we now explore the reductions in productivity dispersion from including input quality measures. Disaggregating workers into separate college and noncollege inputs, as is sometimes done in the literature, actually increases q_{90}/q_{10} slightly. Our most important specifications are the ones that use the detailed human capital measures. q_{90}/q_{10} is 2.90 for the multiplicative labor quality functional form and a similar 2.91 for the additive functional form. An alternative to using human capital measures is to use the wage bill. With the wage bill, q_{90}/q_{10} is 2.83. The wage bill gives lower productivity dispersion than the human capital measures. q_{90}/q_{10} is 2.77 with both the wage bill and human capital measures, the specification that typically minimizes productivity dispersion. 2.77 represents an 8% decline from the baseline of 3.02. The worker fixed effects specification has a q_{90}/q_{10} of 3.25, repeating the earlier finding that fixed effects do not explain as much variation in firm output as other labor input measures. Combining worker fixed effects with detailed human capital measures gives a q_{90}/q_{10} of 2.88. Finally, the benchmark of firm growth and firm age increases q_{90}/q_{10} from the baseline specification by 0.02 to give a q_{90}/q_{10} of 3.04. Any of the specifications using the detailed human capital measures or the wage bill do a better job at reducing dispersion than does firm growth.

We also estimated all nine specifications using a translog production function. Table 7A lists these results for machinery. Remarkably, the estimates of productivity dispersion as measured by $\mathrm{sd}(e)$ are quite similar whether the production function is a Cobb-Douglas or a translog. Adding additional nonlinear terms to a Cobb-Douglas can only weakly increase R^2 and hence will often decrease productivity dispersion measures such as $\mathrm{sd}(e)$ and q_{90}/q_{10}. The measure q_{90}/q_{10} is more sensitive to the functional form for the production function; it tends to be slightly higher for the Cobb-Douglas. Overall, though, the translog results in Table 7 closely track those for the Cobb-Douglas in Table 6.

Table 8A lists the estimates for the machinery industry of productivity dispersion using the Cobb-Douglas production function and the Olley and Pakes estimator. The q_{90}/q_{10} ratios in Table 8A are higher than in Tables 6A and 7A because the Olley and Pakes estimator transfers some of the explanatory power of the measured inputs l and k found using least squares to the ω component of productivity e. The inputs are allowed to be correlated with ω to address simultaneity and selection biases. As an example, the q_{90}/q_{10} from the baseline specification for the machinery industry is 3.42 for Olley and Pakes in Table 8 and 3.02 for least squares in

Table 6. Despite the higher levels of productivity dispersion, the Olley and Pakes estimates follow the same qualitative pattern that productivity dispersion decreases with richer input quality measures.

There are seven other industries listed in Tables 6–8. First consider the three other manufacturing industries of food and beverages, furniture, and publishing and printing in Tables 6A, 7A and 8A. For conciseness, we discuss the least squares estimates using Cobb-Douglas in Table 6A. For food, the q_{90}/q_{10} drops from 3.24 in the baseline row one to 2.85 in the richest specification in row six. That is a 12% drop. For furniture, the decrease in q_{90}/q_{10} productivity dispersion from 3.40 to 2.99 is also a percentage decrease of 12%. For publishing and printing, a hybrid between a service and a manufacturing industry, the decrease in q_{90}/q_{10} is 19%, the largest in the so-called manufacturing sector.

Table 6B reports productivity dispersion estimates for hotels and restaurants, a non-skill-intensive service industry as well as three presumably skill intensive industries: accounting, advertising and computer services. The percentage q_{90}/q_{10} productivity dispersion decline ranges from row one (baseline) to row six (wage bill and human capital measures) is 20% for hotels and restaurants, a small 2.2% for accounting, a large 46% for advertising, and 21% for computer activities.

The manufacturing industries excluding publishing and printing (machinery, food and furniture) have an average drop from the baseline specification to the specification with both the wage bill and the human capital measures, row six, of $(0.08+0.12+0.12)/3 = 11\%$. Publishing and printing, which shares aspects between manufacturing and services, has a drop of 19%. The four service industries have a productivity dispersion decline from including input quality measures of $(0.20+0.022+0.46+0.21)/4 = 22\%$, which is twice the magnitude of the effect for the three pure manufacturing industries and about the same magnitude of the effect for the hybrid industry of publishing and printing. Overall then, the role of labor input quality in productivity dispersion is twice as important by this measure in services as in manufacturing. This is reasonable: the input elasticities of labor in some of these service industries are close to 1, as Table 5 shows for advertising. Another possibility for the larger decline in services is that the production technology is more heterogeneous in the service sector, and the human capital and wage bill measures pick up heterogeneity across firms in the production function. This is somewhat consistent with findings from other data that service industries have higher overall dispersion (Oulton, 1998). On the other hand, every industry is different and there was no reduction in productivity dispersion from adding measures of labor input quality in accounting.

The qualitative pattern of productivity dispersion levels is the same in seven of the eight industries we looked at, with accounting being different. For most industries, adding college and noncollege workers as separate inputs decreases dispersion only a little. Adding detailed human capital controls decreases productivity dispersion by more. In all industries, the productivity dispersion is roughly invariant to whether a multiplicative or additive labor quality functional form is used. The wage bill is potentially a more accurate measure of input quality than the detailed human capital measures. Indeed, the wage bill specification usually gives less dispersion than the human capital specifications. Unsurprisingly, the specification with both wage bill and human capital data decreases dispersion the most.

In unreported results, we have shown that the main results about productivity hold when using a CES production function. Altogether, our results about productivity dispersion are mostly invariant to the functional form of the production function as well as the functional form of the labor quality function. This finding about productivity dispersion contrasts with the signs of the parameter estimates of the production functions, which we argued above are sometimes but not always robustly estimated across functional forms.

Based on prior research, we used each firm's employment growth over the last five years as a benchmark for productivity dispersion decline. Firm growth did decrease productivity dispersion in most but not all of the eight industries, but never by a particularly large amount compared to the drops for labor quality measures. Syverson (2004) studied local demand density (a proxy for competition) and productivity dispersion and found that a one-standard deviation increase in demand density lowered the interquartile range of e by -0.042 log points. For machinery, the interquartile range of productivity e is 0.545 for the base case without labor quality adjustment and 0.483 for the specification with both wage bill and human capital data. Syverson (2004) studied narrow geographic markets for a homogeneous product, concrete. It is not surprising that the mean level of productivity dispersion of 0.275 in his paper is roughly half of our base value of 0.545. Starting from a higher base dispersion, the decrease in productivity dispersion from adding human capital and wage bill data is $0.545 - 0.483 = 0.06$ log points, or the equivalent of a one standard-deviation increase in demand density across local markets in the concrete industry. For advertising, the industry with the highest decline in productivity dispersion from adding input quality, the decline in the interquartile range is $0.658 - 0.348 = 0.31$ log points, or the equivalent of a seven standard deviation increase in concrete demand densities. Our interpretation is that adding human capital variables produces productivity dispersion declines roughly on the same order of magnitude as within-sample changes in concrete demand densities, particularly in manufacturing.

6 Conclusions

Since at least Griliches (1957), economists have speculated that productivity dispersion may arise because firms use inputs of varying qualities. We study labor inputs in part because physical capital is already quality adjusted (up to the caveats about vintage capital and capacity utilization mentioned earlier), as physical capital is usually measured in monetary units. By contrast, researchers often use the number of workers for the labor input. We use detailed data on all Danish citizens to construct human capital measures at the firm level. Human capital inputs do vary across companies in Denmark and our production function parameter estimates show human capital inputs raise firm output considerably. For some industries, the human capital coefficients are statistically precisely estimated.

Adding these quality-adjusted inputs decreases within-industry productivity dispersion. Averaging across the three manufacturing industries of machinery, food and furniture, productivity dispersion as measured by q_{90}/q_{10} declines by 11%. For the hybrid industry of publishing and printing, the ratio q_{90}/q_{10} declines by

19%. For the four service industries, the q_{90}/q_{10} ratio declines by an average of 22%, although this effect averages out the large decrease in advertising and the small decrease in accounting. Overall, it seems like input quality plays a greater role in explaining productivity dispersion in services than in manufacturing, although it certainly plays a measurable role in most industries.

The decline in productivity dispersion from adding controls for firm growth, a measure emphasized in the literature, was noticeably less than the decline in productivity dispersion from adding controls for input quality. The decline in productivity dispersion from input quality is roughly the same order of magnitude as the competitive effects studied in local geographic markets by Syverson (2004).

Input quality is one of perhaps many factors that contribute to productivity dispersion. However, labor quality does not explain most productivity dispersion. Returning to an issue we raised in the introduction, our results suggest that productivity mostly represents some attribute of a firm that cannot easily be bought and sold on the market for inputs. Possibilities include management quality, business strategy, the appropriate use of new technologies and heterogeneous production technologies. If a large portion of productivity cannot be traded, then the performance of product markets may be as important for economic efficiency and aggregate productivity growth as the performance of input markets.

A Danish labor and accounting data

We use accounting data for capital, value added and investment. The accounting data come from Købmanstandens Oplysningsbureau (KØB), a Danish credit-rating agency. The accounting data are an unbalanced panel that roughly covers the period 1992–2001 and uses each firm's proprietary accounting period. We rescale the accounting variables to a twelve-month, calendar-year basis.

We use value added as a measure of output and our measure of physical capital is tangible assets net of depreciation. Value added is reported for many more firms than total sales, perhaps because of the role of value added in value added taxes. We disregard firms that lack rescaled accounting information on valued added and fixed assets for a twelve-month period. For the labor input, we count the total number of workers in IDA, which is described below. Firm age is directly reported in the accounting data. We include the log of firm age in some specifications.[20]

To construct labor quality variables, we use the Danish Integrated Database for Labor Market Research (IDA), one of the central registers of Statistics Denmark. IDA combines several types of data. One dataset provides information at the individual level on demographics (age, sex, marital status, family status) and schooling for

[20]We construct investment from the accounting data in order to control for the endogeneity of the labor input using the Olley and Pakes (1996) approach. Investment is computed using the formula $i_t = k_{t+1} - k_t + \delta_{t+1}$, where δ_t is the total amount of depreciation and t is the year. Investment cannot be missing. The accounting data report the total depreciation δ_{t+1}. We use the same sample for Olley and Pakes estimates as for the pure least squares estimates. Therefore, we drop firms with missing i_t or i_{t-1}. The timing assumption is that investment in year t, i_t, is made in year t, so it informs us about year t's productivity.

all Danish citizens for 1980–2001. Each individual is given a unique identification number that can be further used for matching with the other datasets of IDA. Another IDA dataset's unit of observation is an individual's job. It contains information on individual labor earnings, some other variables and the number of years of labor market experience. Labor market experience is computed since 1964 by Statistics Denmark.

Both full- and part-time jobs are included, but in the rare case of a worker with three or more jobs, only the primary and secondary jobs are reported. The data also contain a unique identification number for each job's establishment. IDA's establishment dataset provides a firm identification number that can be use for matching with other firm-level data.

We use IDA for 1980–2001 to compute labor-market-history variables such as firm tenure and industry tenure. We compute firm tenure as the number of years a worker has been attached to a given firm. As we are concerned with spurious changes in firm identification codes, a worker's tenure is reset to zero only if both his firm and establishment identifiers change at the same time. We construct industry tenure using the following eight broad sectors: (1) agriculture and mining, (2) manufacturing, (3) construction and transport, (4) retail, hotels and restaurants, (5) finance, real estate and R&D activities, (6) public sector, (7) private households and extraterritorial activities and (8) others. These sectors encompass all Danish firms and are not equivalent to the industries for our estimation sample.

Industry is recorded at the establishment level. For our regressions, a multi-establishment firm's industry is the weighted (by number of workers) modal establishment industry.

All inputs are constructed at the firm level. We construct firm-level fractions of workers who have a given characteristic, say a college degree or 6–9 years of firm tenure. The intervals are simple to interpret as each measure is a fraction between 0 and 1. The intervals allow us to examine nonlinearities, and they handle topcoding from not observing firm and industry tenure for spells starting before 1980.

We estimated production functions for two samples: all firms with nonmissing variables and a sample with outliers removed. We are worried about possibly non-classical measurement error in the accounting data, so we removed the firms in the top and bottom 1% of the ratios of output to labor and also physical capital to labor. Removing these outliers increases the base R^2's substantially, but does not change the ΔR^2's from adding labor quality much. We report specifications with the outliers removed, but our main conclusions about ΔR^2's are similar if we include the outliers.

References

Abowd, J.M., F. Kramarz, and D.N. Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes, "Econometric Tools for Analyzing Market Outcomes.," in James J. Heckman and Edward E. Leamer, eds., *Handbook of Econometrics*, Vol. 6A, Elsevier, 2007.

_, Kevin Caves, and Garth Frazer, "Structural Estimation of Production Functions," 2007. working paper.

Arellano, M. and S. Bond, "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations," *The Review of Economic Studies*, 1991, *58* (2), 277–297.

Baily, Martin Neil, Charles Hulten, and David Campbell, "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity: Microeconomics*, 1992.

Baldwin, John R., *The Dynamics of Industrial Competition: A North American Perspective*, Cambridge University Press, 1995.

Bartelsman, Eric J. and Mark Doms, "Understanding Productivity: Lessons from Longitudinal Microdata," *Journal of Economic Literature*, September 2000, *38* (3), 569–594.

Bernard, A.B. and J.B. Jensen, "Exporters, Jobs, and Wages in US, Manufacturing: 1976-1987," *Brookings Papers on Economic Activity: Macroeconomics*, 1995.

Berndt, E.R. and D.M. Hesse, "Measuring and assessing capacity utilization in the manufacturing sectors of nine OECD countries," *European Economic Review*, 1986, *30* (5), 961–989.

Bloom, Nick and John Van Reenen, "Measuring and Explaining Management Practices Across Firms and Countries," *Quarterly Journal of Economics*, November 2007.

Blundell, R. and S. Bond, "GMM Estimation with persistent panel data: an application to production functions," *Econometric Reviews*, 2000, *19* (3), 321–340.

Cabral, Luis M. B. and José Mata, "On the Evolution of the Firm Size Distribution: Facts and Theory," *American Economic Review*, November 2003, *93* (4), 1075–1090.

Christensen, Laurits R., Dale W. Jorgenson, and Lawrence J. Lau, "Transcendental Logarithmic Production Functions," *The Review of Economics and Statistics*, February 1973, *55* (1), 28–45.

Collard-Wexler, Allan, "Productivity Dispersion and Plant Selection in the Ready-Mix Concrete Industry," 2010. working paper.

Davis, Steven J. and John C. Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, August 1992, *107* (3), 819–863.

Denison, Edward Fulton, *The Sources of Economic Growth in the United States and the Alternatives Before Us*, Committee for Economic Development, 1962.

Dhrymes, Phoebus, *The Structure of Production Technology: Productivity and Aggregation Effects*, Edward Elgar, 1995.

Dupuy, Arnaud, "Sorting on Skills and Preferences: Tinbergen Meets Sattinger," February 2010. Maastricht University working paper.

Eeckhout, Jan and Philipp Kircher, "Identifying Sorting – In Theory," 2009. University Pompeu Fabra working paper.

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

Fox, Jeremy T. and Valérie Smeets, "Three Empirical Checks for the Validity of Olley and Pakes Productivity Estimators," 2010. University of Michigan working paper.

Frazer, Garth, "Heterogeneous Labour in Firms and Returns to Education in Ghana," 2006. University of Toronto working paper.

Griliches, Zvi, "Specification Bias in Estimates of Production Functions," *Journal of Farm Economics*, February 1957, *39* (1), 8–20.

_ and Jacques Mairesse, "Production Functions: The Search for Identification," in S. Strøm, ed., *Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium*, Cambridge University Press, 1998.

Haltiwanger, John, Julia Lane, and James Spletzer, "Wages, Productivity, and the Dynamic Interaction of Businesses and Workers," *Labour Economics*, 2007, *14*, 575–602.

Hellerstein, J.K. and D. Neumark, "Sex, wages, and productivity: An empirical analysis of Israeli firm-level data," *International Economic Review*, 1999, *40* (1), 95–123.

_ , _ , and K.R. Troske, "Wages, productivity, and worker characteristics: Evidence from plant-level production functions and wage equations," *Journal of Labor Economics*, 1999, *17* (3), 409–446.

Hellerstein, Judith and David Neumark, "Production Function and Wage Equation Estimation with Heterogeneous Labor: Evidence from a New Matched Employer-Employee Data Set," in "Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches," University of Chicago, 2006.

Iranzo, S., F. Schivardi, and E. Tosetti, "Skill dispersion and firm productivity: an analysis with employer-employee matched data," *Journal of Labor Economics*, 2008, *26* (2), 247–285.

Jorgenson, D.W., F.M. Gollop, and B. Fraumeni, *Productivity and US Economic Growth*, Harvard University Press, 1987.

Katayama, Hajime, Shihua Lu, and James R. Tybout, "Firm-Level Productivity Studies: Illusions and a Solution," *International Journal of Industrial Organization*, 2009, *27*, 403–413.

Klette, Tor Jakob and Zvi Griliches, "The Inconsistency of Common Scale Estimators when Output Prices are Unobserved and Endogenous," *Journal of Applied Econometrics*, 1996, *11*, 343–361.

Lise, Jeremy, Costas Meghir, and Jean-Marc Robin, "Matching, Sorting and Wages," 2010. University College London working paper.

Lopes de Melo, Rafael, "Sorting in the Labor Market: Theory and Measurement," December 2009. University of Chicago working paper.

Maksimovic, V. and G. Phillips, "The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?," *The Journal of Finance*, 2001, *56* (6), 2019–2065.

Marschak, Jacob and William H. Andrews, "Random Simultaneous Equations and the Theory of Production," *Econometrica*, July-October 1944, *12* (3/4), 143–205.

Mulligan, C.B. and X. Sala-I-Martin, "Measuring Aggregate Human Capital," *Journal of Economic Growth*, 2000, *5* (3), 215–252.

Olley, G. Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, November 1996, *64* (6), 1263–1297.

Ornaghi, C., "Assessing the effects of measurement errors on the estimation of production functions," *Journal of Applied Econometrics*, 2006, *21* (6), 879–891.

Oulton, N., "Competition and the dispersion of labour productivity amongst UK companies," *Oxford Economic Papers*, 1998, *50* (1), 23–38.

Schoar, A., "Effects of Corporate Diversification on Productivity," *The Journal of Finance*, 2002, *57* (6), 2379–2403.

Syverson, Chad, "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 2004, *112*, 1181–1222.

Van Biesebroeck, Johannes, "Wages Equal Productivity: Fact or Fiction?," June 2007. NBER working paper No. 10174.

Welch, Finis, "Linear Synthesis of Skill Distribution," *The Journal of Human Resources*, Summer 1969, *4* (3), 311–327.

Whelan, K., "Computers, obsolescence, and productivity," *Review of Economics and Statistics*, 2002, *84* (3), 445–461.