Labor productivity, age and education in Swedish mining and manufacturing 1985-1996*

Björn Andersson and Thomas Lindh†

Version 3, but still preliminary and not to be quoted without consent from the authors

May 1, 2008

Abstract

We study the productivity effects of changes in education and age of the labor force in Swedish manufacturing over the eleven years 1985-1996. Our data allow us to examine the composition of the workforce with respect to demographics and education at each plant. Education and age correlate due to cohort effects since the younger workers have higher education. This makes identification of education effects precarious. We investigate this problem by conditioning on the age composition. The tentative result is that tertiary education has less positive effects than secondary education in these industries.

JEL Codes: C23, J24, L60

Keywords: plant productivity, demographics, education

---

*We are grateful for financial support from the Swedish Council for Work Life Research (RALF) who funded the purchase and cleaning of data for this study. We are also very grateful to Fredrik Andersson, Hans Schwarz, Christian Svendsen and Max Halvarsson for their patience in cleaning and organising the data.

†Björn Andersson, Sveriges Riksbank, 103 37 Stockholm, Sweden. Tel +46-8-787 03 88. Fax +46-8-21 05 31. Email bjorn.andersson@riksbank.se. Thomas Lindh, Institute for Futures Studies, Box 591, SE-101 31, Stockholm, Sweden. Tel +46-8-402 12 16. Email thomas.lindh@framtidsstudier.se. Lindh is also affiliated to Växjö University.

1 Introduction

During the end of the last century more detailed micro-level data sets made possible studies of productivity at the plant level. With the use of the new data it has been possible to investigate and test assumptions in many different research fields.[1] One of the main results of the new research has been the finding of a persistent productivity dispersion over plants (Baily et al. 1992, Bartelsman and Dhrymes 1998), something which has called the conventional aggregate production-function approach into question. However, there is still a large uncertainty about how much of this productivity dispersion that is a systematic feature of the data and how much that is simply a consequence of unobserved heterogeneity.

It has been difficult to investigate this since there has been a dearth of data where individual characteristics of labor can be matched to plant-level productivity. A few papers e.g. Entorf and Kramarz (1998), Entorf et al. (1999) and Doms et al. (1997) have had access to longitudinal data over plants where worker characteristics are known, but the longitudinal dimension has in general been of rather limited length.[2]. By now more data sets of matched employer-employee matched data over longer time periods have emerged, see for example . In recent years the availability of employer-employee matched data sets have multiplied and there are now a number of studies attempting to study how age and education affect productivity at the plant level. A great deal of studies have focused on the old seniority pay discussion, e.g. Beffy et al., Ilmakunnas and Maliranta, Aubert and Crépon etc. Results are, however, rather mixed, due to a host of estimation difficulties.

We have at our disposal a longitudinal data set where all Swedish manufacturing plants with more than 5 or 10 employees can be followed for the period 1970-1996. For the years 1985-1996 there are in addition educational and demographic data on the individuals working in manufacturing that can be matched to the plant data. Unfortunately, our data lack technological information at the individual level, which some of the papers mentioned above have access to. However, our focus is on the plant rather than the individual; whereas the aforementioned studies have been directed towards

[1] See Bartelsman and Doms (2000) for an overview of previous research.

[2] While Doms et al. (1997) have only one cross-section of worker characteristics and a rather small survey sample of larger plants, Entorf et. al. (1999) can follow both workers and firms 1991-1993 so they also have a longitudinal dimension in worker characteristics.

issues of technology use and wage formation, our focus is on productivity and human-capital effects on productivity. Of course, human capital is not an unambiguous concept. It could be argued that it is in fact synonymous with labor, since there is no such thing as "raw labor" that does not require any initial investments. Every child requires huge investments in time, care, food, housing, clothing and basic training before it reaches an age and maturity where any kind of productive labor can be had. Such a comprehensive definition may not be very useful, however, since it is difficult to measure differences in that kind of investment as distinguished from innate capabilities. The conventional view that human capital consists of investments in formal education and work experience is easier to handle.

As Bartelsman and Doms (2000) remark there is some uncertainty surrounding the question of how much of the very large productivity differences between plants that are due to random measurement errors and how much that is due to unobserved heterogeneity. There is, however, substantial evidence that productivity differences are quite persistent. As Mellander (1999) shows in a factor demand study of Swedish industries, productivity differences between industries can to some extent be explained by the composition of the labor force with respect to education and age. Börsch-Supan et al. (2007) show that even at the work station level age composition effects can be found along an assembly line with very standardised tasks. This is, of course, not unexpected, but due to lack of data these dimensions have often been ignored in traditional productivity studies. Our paper contributes to the literature by examining how productivity varies with the age and education composition of the workforce in Swedish manufacturing over the years 1985-1996. In particular we focus on the less studied effects of educational composition.

Previous literature has not ignored this question. Labor economists have always been interested in the connection between productivity and the skill of the work force and the human capital. Growth-accounting studies have traditionally used wage differences as proxies for average productivity differences in order to decompose the labor input of different groups. This is, however, problematic since research has shown ((Forslund and Lindh 1997)) that wage differences are only weakly related to productivity differences at the plant level and thus are very unlikely to catch the average productivity effect of human capital. After all, wages are supposed to measure the *marginal* worker contribution. While it is clear from countless studies that education confer a wage premium as expected and macro productivity studies have started to come to grips with measurement errors (Lindahl and Krueger 2001, Lutz et al.) showing that there is indeed a positive effect from education on productivity growth, the evidence from the plant level is less clear. In fact, there is evidence that at least in Swedish manufacturing it is secondary schooling rather than tertiary that is important for productivity growth (Gunnarsson et al.). Also in our data a straight regression of education on productivity shares at the plant level indicates that upgrading from secondary to tertiary education is associated with a decrease in productivity. This is rather puzzling in view of the wage premium evidence. While diminishing returns from education might be expected, negative returns is clearly at odds with both theory, intuition and other emprirical evidence. Interpreted as a composition effect a more plausible explanation may therefore be that the addition of tertiary education to the optimal composition is comparatively of less importance in our manufacturing and mining data. In a more formal way we would thus interpret the evidence as a lower elasticity of output with respect to tertiary education.

Thus, our primary aim here is to investigate the reasons for this unexpected result. There are a number of difficulties in identifying plant level effects on productivity. First, the composition of the workforce is not independent of other decisions and shocks affecting productivity, hence endogeneity bias through correlation between residuals and right hand variables is an issue. Second, considerable heterogeneity problems could be suspected since technology may differ considerably according to size, specific industry, age of plants and region. A host of other identification problems involving collinearity and non-linearities are also present. In order to start probing these issues we in particular focus on the interaction of age and education, where we have an obvious problem in identifying cohort effects, since young workers are in general more highly educated than old workers.

The next section of the paper contains a description of our data set as well as some descriptive graphics. In Section 3 the econometric model specification is discussed and the results are presented in Section 4. In Section 5 concluding remarks and suggestions for further research are made.

2 Matched individual and plant data

Our data is a specially matched set of establishment data from the Industrial Statistics 1970-1996 and individual data from the Annual Regional Employment Statistics (ÅRSYS) 1985-1996, later renamed Regional Labour Market Statistics (RAMS), both produced by Statistics Sweden. The population for the plant data are all workplaces within mining and manufacturing in Sweden. However, only firms with at least five employees (ten after 1990) are included in the total survey; smaller plants have been included with some data imputed based on a sample survey. These plant data have been matched to individuals at the plant level in order to get measures of the educational and demographic composition of the labor force at each plant. Labor productivity is measured as value-added per employed person deflated using the producer-price index. There are some matching problems, forcing us to exclude some observations where we cannot unambiguously identify the plant to which individuals belong. These are mostly small firms that we would have excluded anyway.

The measure of productivity per employed is, as all variables taken from the Industrial Statistics, an average over the year. This leads to some compatibility problems with the ÅRSYS data since these data refer to the employment in a particular week in November. The employment variables in the two data sets might therefore differ quite a lot for plants that have experienced dramatic changes during the year. We have solved this problem by interpolating the employment values in ÅRSYS - including individual age/education cells - for these plants based on the difference in employment between the two data sets.[3]

The sample of plants we use in the empirical investigation is selected as follows: since data on age and education of the employees only are available from 1985, all observations for the years 1970-1984 are dropped. In order to get a somewhat more homogeneous sample we in some cases exclude plants with less than 50 employees on average over the years the plant is present in the data. There are two problems with smaller plants. One is that data *per se* often are less reliable since survey answers often are lacking, another is that shares of age/education groups are less likely to have a systematic relation to productivity in smaller firms. The fewer people the harder it is to achieve an optimal mix due to indivisibilities. The exact limit we have chosen here is, admittedly, a bit arbitrary and chosen in order to be on the safe side. This exclusion reduces the number of plants with about 75 percent and the number of observations with about 70 percent. It might seem radical to throw away that many observations but our sample still comprises well over half of the employment in manufacturing, leaving us with an unbalanced

[3] To be precise, we linearly interpolate values on total employment in years when the relative difference in employment between the two sets is deemed too large - defined to be the 10% largest differences. When extrapolations at the beginning or the end of the period are needed, e.g. for periods T and T-1, values are calculated sequentially as:

$x_{T-1} = 0.5(x_{T-2} + x_{T-3})$

$x_T = 0.5(x_{T-1} + x_{T-2})$

When a value on total employment has been interpolated for a plant a particular year this aggregate number is divided up on an individual age/education cell according to a ratio determined by the average of the weight of the cell before and after the interpolated year. Plants for which the employment in ÅRSYS is zero all of the years have been dropped since imputation of age/education values then becomes impossible.

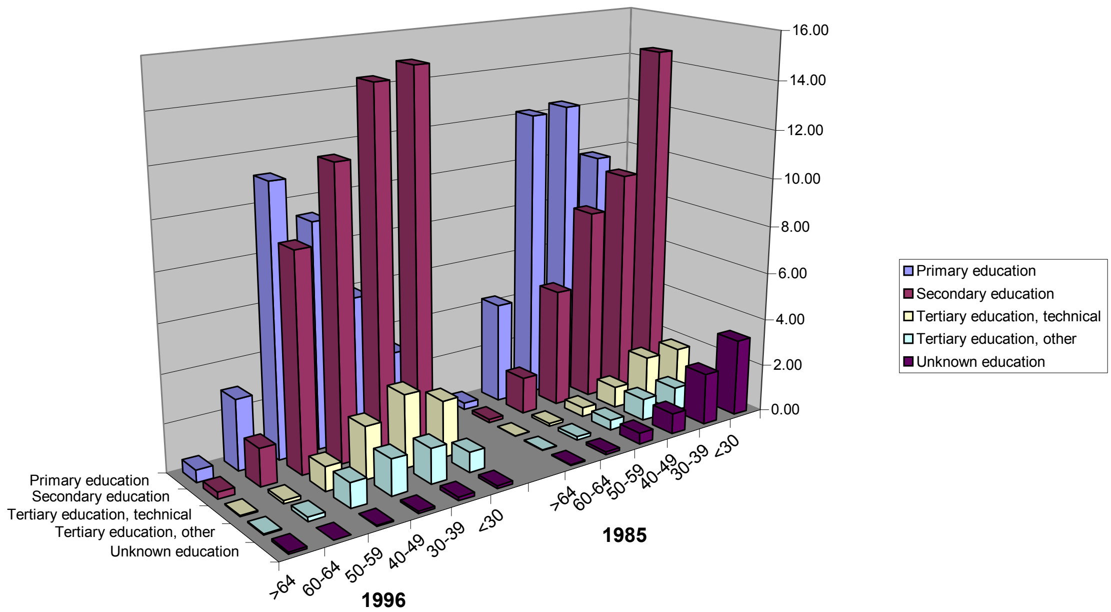

Figure 1: Comparison of the general level of education in different age groups 1985 and 1996.

panel with roughly 28 500 observations on 4 000 plants over a period of 12 years.

In our analysis we focus on employment shares of different age/education groups. We subdivide the workforce into the age groups 16-29, 30-39, 40-49, 50-59, and those 60 and above. The educational categories we use are those with at most compulsory primary education, those with at most secondary education, and for those with tertiary education we sometimes subdivide into two groups: education with technological content and other content. There is also a group where the educational code regarding length and/or subject has not been established. This group is considerably larger before 1990 than after due to better data collection routines being introduced at that time. In the appendix, Table A1 presents some descriptive statistics for the different variables. Figure 1 below shows the distribution of categories in the whole of manufacturing in 1985 and 1996. As can be seen the general level of education has increased in each age group over the period making the risk of confounding education and age effects obvious.

As was mentioned in the introduction it is quite common in empirical investigations to use wage differences as proxies for marginal productivity differences between different groups. Figure 3 shows the danger of this on plant data. In the figure observations have been ranked by productivity and the corresponding wage per employee has been plotted for a representative industry, the food industry—the picture is basically the same regardless of industry—in four different years. What we see are Salter curves (after Salter, 1960). The naked eye cannot discern any relation at all between productivity and wage level. In fact, a regression reveals that the wage line may have a faint upward, but it is quite indisputable that this cannot be very important. In Forslund and Lindh (1997) it is shown that a similar picture holds not only for our sample period but even back into the 1960s. Not only does it hold for each industry in each year, but the actual productivity distribution has a very distinctive and stable form; a truncated Cauchy distribution with time-varying parameters provides an excellent fit for all cases. Although the general shape of the productivity distribution has been known for a very long time, economists, educated in the tradition of representative agents and firms, have a strong tendency to ignore the fact that the average plant in this distribution is in fact not very representative of the aggregate at all. Rather what matters are the properties of inflow and outflow in the distribution. One pertinent observation in the literature is that local labor markets seem to improve when there is a large inflow of young workers (Nordström Skans, 2002, studies Swedish local labor markets and Shimer, 1999, studies U.S. states). One hypothesis why this is so could be that a larger inflow improves on matching efficiency in the labor market. This would be expected to also have productivity effects and indeed Nordström Skans (2007) finds that this is the case. Manufacturing productivity is higher in the local labour market areas where the share of workers above 50 is higher.

Figure 4 shows the time series of the average number of employees in different age groups with different educational levels. It is clear from the graph for the age group 16-29 that the recession in the beginning of the 1990s is associated with a decrease in the number of young employees with primary and secondary education, of which a large part was employed during the previously overheated business cycle in the end of the 1980s. We can see a similar drop in the groups 30-39 and 40-49, but it is interesting to note that this decrease only occurs for the group with at most primary education. There is actually an expansion of the number of people with secondary education in all ages 30-59 during the recession, and for the age group 50-59 this is even true for those with primary education. In all age groups we also see a steady increase in the number of employees with higher education. This leads to the conclusion that the shedding of employment during the recession was upgrading the education level, not only by firing those with low education and low experience but also by actually hiring more people with higher

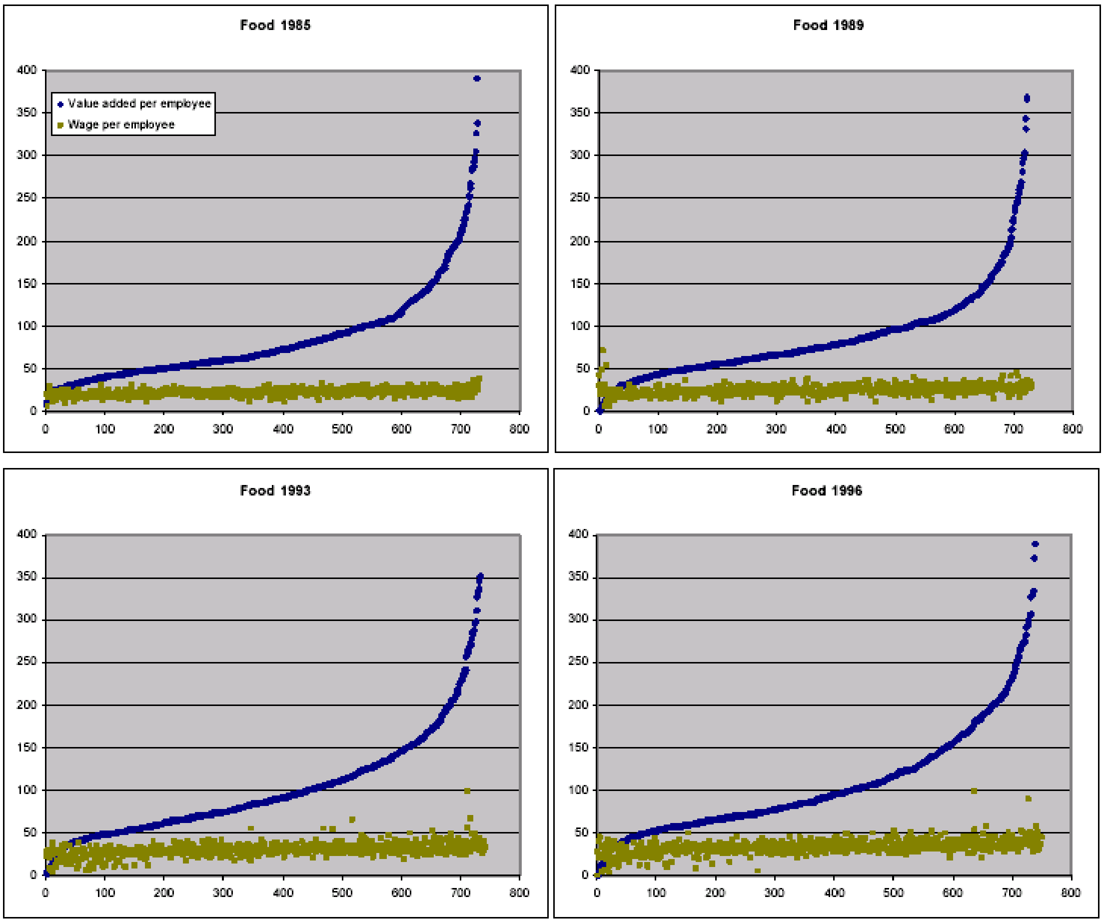

Figure 2: Rank plot of productivity with corresponding wage costs per employee in the lower line for the food industry in 1985, 1989, 1993, and 1996. The horizontal scale is the rank in the distribution of plants, while the vertical scale are amounts in thousand 1968 SEK (approximately 0.7 times this figure equals current EUR).

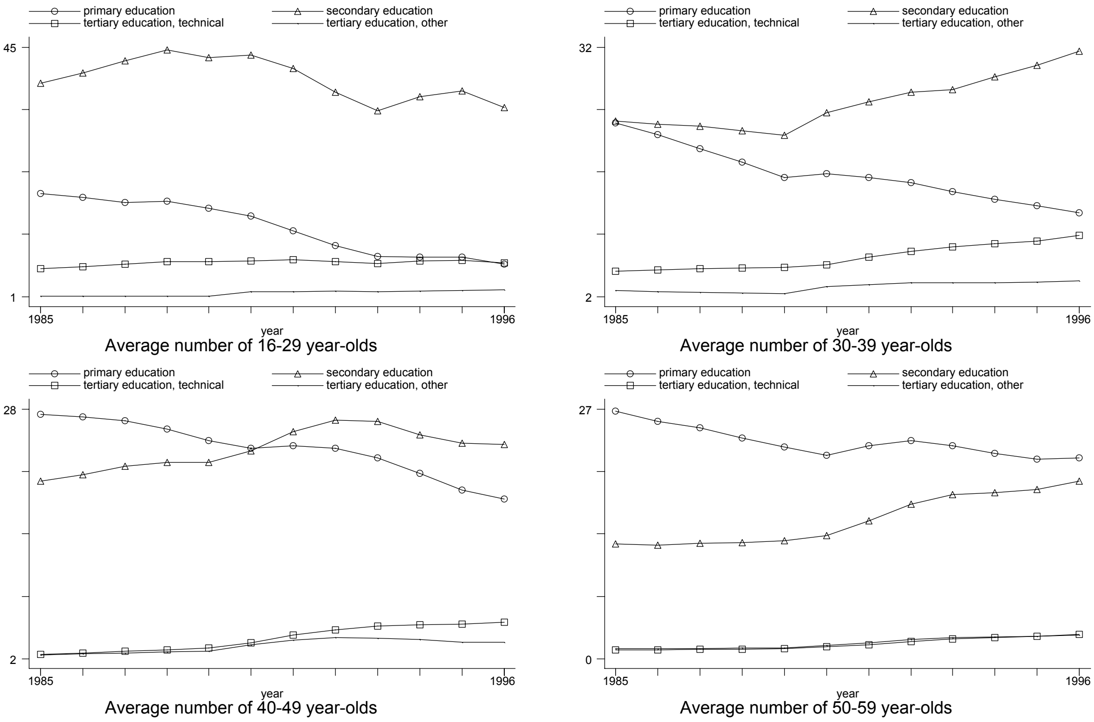

Figure 3: Plots of the average number of employees in age groups at different educational levels 1985-1996.

education or more experience.

In Figure 5 the mean, median, and interquartile range of value-added per employee are shown for four different industries. Even though there are similarities between the plots it is also clear that there are noticeable differences in business-cycle sensitivity between industries, both regarding the center and spread of the productivity distribution. Saw- and planing mills and the pulp industry both display more volatility over time than the two machinery industries with two distinguishing peaks at the end of the 1980s and the middle of the 1990s. The machinery industries, particularly other machinery and equipment, are less sensitive to these business-cycle fluctuations. As Figure 5 also demonstrates, industries that share a time-pattern for the center of the distribution do not necessarily have similar variations in the spread. Unlike the saw- and planing mill industry that has procyclical fluctuations of the tail of the productivity distribution during the

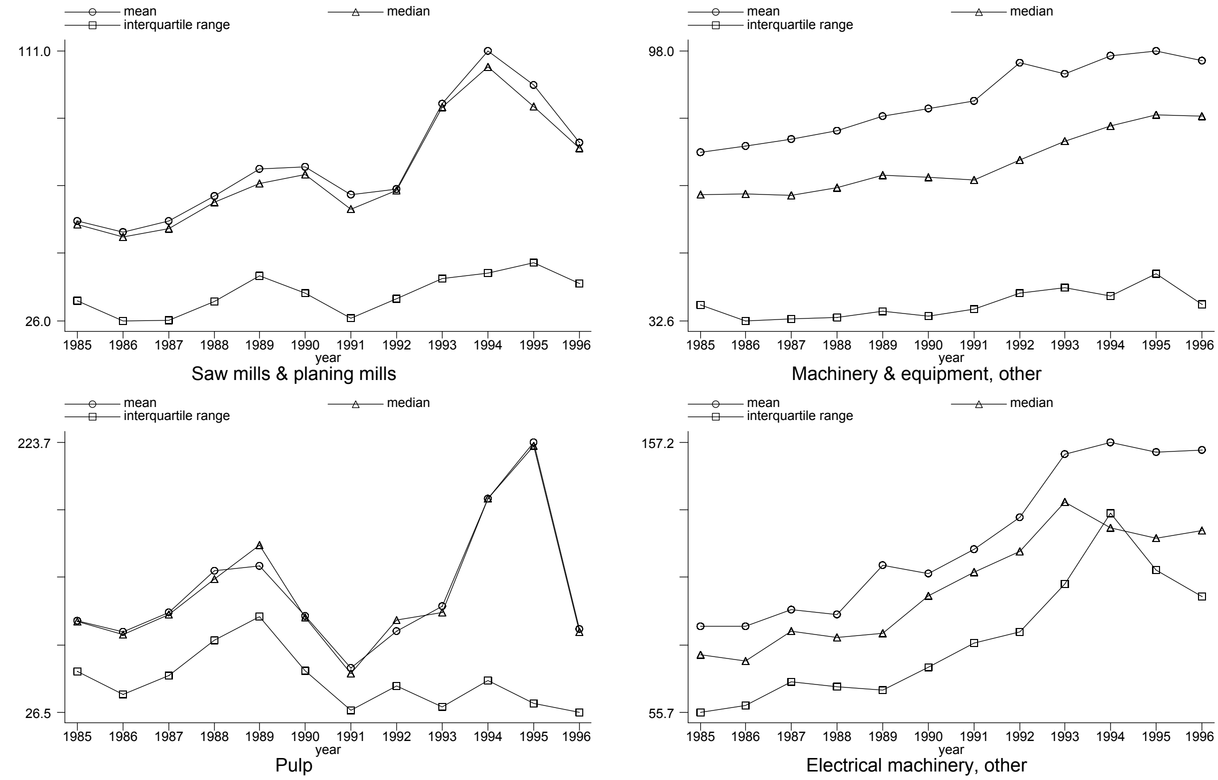

Figure 4: Plots for the mean, median, and interquartile range of value-added per employee for four different industries.

whole period, the effect of the slump in the beginning of the 1990s resulted in a more stable shape of the distribution in the pulp industry. The spread of the distribution in the two machinery industries also have quite dissimilar time patterns, particularly in the 1990s.

Accounting for this type of heterogeneity of the different industries is difficult to do in a traditional panel-estimation framework with fixed effects. Type of industry is a plant-specific effect (by construction of our sample), which means that we could rely on the standard within estimator to purge the data of the industry effects and then let time effects pick up period effects that are common to all plants - if we believed that the shape of the time pattern is the same in all industries. However, as shown in Figure 5, the industry effect is clearly different in e.g. the recession. In order to rid the data of these effects we will, in the estimations, run standard fixed-effect regressions on data where the industry/time-specific mean have been removed. Thus we are in effect using variables that are deviations from the average in each industry and for each year. There is a risk that we then have removed too much variation, in particular some of the variation between contracting and expanding industries. This variation is likely to be correlated to the composition of the workforce so variation due to inflow and outflow of labor between industries has been removed.

3 Model specification

Productivity estimation generally starts off from an assumption that all firms have access to the same technology and maximize profits given the constraints of factor supply and demand for the product. Under an assumption of constant returns to scale perfect competition can be assumed and it follows that in market equilibrium we can treat the outcome as an equalization of marginal factor unit costs to product price. This has always been a controversial assumption but the improved access to plant level data and studies of the gross in- and outflow of plants (Davis et al. 1996) has made this "representative plant" assumption even more questionable. The very large and persistent productivity dispersion that we observe within industries is very hard to reconcile with a representative firm view.

Of course, there is no shortage of explanations for this stylized fact. Some degree of market power seems wide-spread in the real world. Economies of scale and non-malleable capital assets are all factors making it more likely that there is a high degree of inertia and at least temporarily inefficient allocation of production resources. In Sweden that has actually been an implicit assumption behind important parts of the labor market policies implemented in the post-war period. The so called Rehn-Meidner model explicitly saw solidarity wage policies as a means to force resources out of firms working with obsolete capital equipment into more dynamically efficient firms. The formal modelling of such structures through vintage or putty-clay models under imperfect competition is, however, no easy task. All too often the radical simplifications necessary to make models tractable stifle the objective of achieving a more realistic model of production.

To these difficulties we have to add the untenable assumption of more or less homogeneous labor. It is all too obvious that labor is in reality a very heterogeneous input even with respect to easily measured dimensions as formal education and experience, not to speak of unobservables as task-specific ability and tacit knowledge.

Since it is impossible for us to construct reliable capital measures at the plant level for the whole sample we focus on labor productivity in the analysis here. To control for differences in capital structure to some extent we assume that there are industry-specific and plant-specific characteristics which catch the effects of a partly nonmalleable capital stock and differences in technology.[4] Labor productivity π is then viewed as a simple linear function

$$\pi_{it} = \alpha + \sum_{k=1}^{m} \beta_k n_{kit} + \varepsilon_i + \nu_t + \eta_{bt} + \mu_{it} \tag{1}$$

where n_k is the share of different types of labor used in the production plant i in period t within industry b. We assume that there are specific intercepts for each production unit, industry and time period. The plant-specific fixed effects will catch any systematic differences in capital level that are persistent over time. The time-specific effects on the other hand will catch common business-cycle variations in investment activity and thus most of the short-run variation in the capital stock. Finally the industry-specific effects should account for industry-specific demand variations as well as some of the technological differences.

3.1 A production function interpretation?

From the viewpoint of production theory equation (1) corresponds closest to a production-function approach to estimation. Inserting the logarithms of labor categories and labor productivity corresponds to the assumption that

$$\frac{y_{it}}{N_{it}} = A_{it} \left(\prod_{k=1}^{m} n_{kit}^{\beta_k} \right) e^{\varepsilon_i \nu_t \eta_{tb} u_{it}} \tag{2}$$

Thus, if the β_k sum to one the production function will be a constant returns to scale Cobb-Douglas function with labor inputs L_k.[5] Any general technological shocks will be caught in the time effects. But so will demand shocks and it is hard to identify these separately. Time-varying plant effects are theoretically quite plausible but are clearly impractical to estimate.

[4] Some experimentation with capital measures for a sub-sample of data where we had investment data has been undertaken in a companion paper focusing on age structure, Malmberg et al. (2008), and has been shown to have surprisingly small effects on the age composition effects. It is possible that capital and education has a more complementary relation, however.

[5] This is not strictly correct when we are using overlapping labor inputs, but for the sake of argument we make this simplification. Note that this is not at all unusual to see in the literature where human capital is often added as a separate input in spite of the fact that it cannot be physically separable from labor.

Industry-specific time effects might however be possible to estimate. We will, however, restrict ourselves to common elasticities for the labor categories since the weak relation of wage levels to productivity indicate that marginal productivities are on average more or less equalized over plants and industries. First-order profit maximization conditions under perfect competition with this technology are

$$\frac{\partial y}{\partial L_k} = \frac{\beta_k y}{L_k} = \lambda w_k$$

so the wage share $\frac{w_k L_k}{y}$ is proportional to the elasticity β_k. It would seem that, at best, this could hold at the industry level, or at more aggregated levels. The partially decentralized wage bargaining in Sweden during the period in question is mainly performed at the industry level.

To assume that each firm performs this optimization at the plant level requires us to believe that technological elasticities by some coincidence happen to vary so that wages become more or less equalized, or that the shadow values of labor in different plants vary in such a way that common wages can be accommodated within the same technological structure. To see this consider that

$$\beta_{ki} \left(\frac{y}{L_k} \right)_i = \lambda_i w_k \qquad \text{for each } i$$

must hold in each firm i. Since labour productivity varies dramatically in an industry the ratio λ_i / β_{ki} must match labour productivity variation more or less exactly. In practice we know that average wages are more or less constant so this should mean that labour elasticities vary approximately proportional to labour productivity. Since the form of the productivity distribution is highly persistent it is highly unlikely that this parameter follows such a strict rule in every industry. The Cobb-Douglas form also assumes that the elasticity of substitution between different categories of labor is equal to one. This might be a too harsh restriction on data since there are obvious differences in complementarity between young and old of the same educational status (Kremer and Thomson 1998).

The conclusion is that we can only interpret the results of estimating the equation above as an aggregate production function and not take it to imply that this holds strictly at the plant level. Plausible assumptions to validate a plant interpretation would require allowing for different parameters for each plant and also allow for time-varying parameters to add to the problem. Some kind of vintage structure or putty-clay function might be able to handle this, but since we know that plants also are migrating both up and down in the distribution, we would need some hybrid model. At this stage we find it more important to preserve transparency and view the estimation more as a data description tool.

It is common to use cost-function approaches at the industry level for productivity research (for example Mellander 1999 and Morrison 1988) since price data at this level are often more widely available and more reliable than output and input data. However, at the plant level the situation is rather the reverse. It is practically impossible to get any plant-specific price indexes while, at least in manufacturing, output and input data are available. Separate wage data for the different categories are very hard to construct. As previously remarked the lack of capital data is not satisfactory but with regard to the uncertainty of such data we might as well use a fixed effect approach to correct for it.

There are, of course, countless dimensions along which labor is differentiated. In some sense every human being contributes a unique mix of abilities, experiences and knowledge to the production process. In order to avoid the curse of dimensionality we have to limit ourselves to a tractable set of characteristics. Gender, age and education are three dimensions in which labor is clearly distinguished. Subdividing labor into categories along these dimensions is, therefore, a natural choice. Still, to obtain a manageable set of characteristics, it is necessary to keep the categories rather rough. With five educational categories, six age groups and two genders we have 60 different kinds of labor to deal with. If we are to take account of interactions between these groups the number of parameters quickly becomes unmanageable. Since female labour is clearly underrepresented in manufacturing and mining we skipped this category as a first step.

But 30 variables may be too much also. However, we started looking at estimations using this full set of disjunct variables. Very few significant parameter estimates made it difficult to interpret the results. An attempt to impose polynomial restrictions on the coefficients in order to deal with excessive collinearity resulted in even fewer significant parameter estimates and very strange age and education patterns implying that the share of young people with only primary education had the most positive effect on productivity. We therefore had to resort to use overlapping categories. That is age shares and education shares taken as separable factors.

Another issue that we need to deal with is the fact that educational achievement to some extent is cohort-specific (Ohlsson 1986, Easterlin 1961, Macunovich 1998) and thus creating a potential identification problem since the cohort-specific education will transfer from age group to age group and thus may be confounded with time effects. The fact that we aggregate age groups in fairly wide segments should, however, alleviate that problem. (On this issue see further in Fienberg and Mason 1985 and Andersson 2001).

How much of a problem this is also depends on to what extent the workforce composition at a given plant is stable. The less turnover there is the more of a problem this issue can become. *A priori* we can only note this. We know, however, that both job flows and worker flows are quite substantial. The former are on average annually in manufacturing around 20 percent and the latter may reach 30-40 percent with a rather wide spread over business cycles. There are quite dramatic differences in these flows for different age and educational categories.

The basic estimation equation 1 assumes that we can treat the shares of different worker categories as independent variables with a loglinear or linear relation to labor productivity. Besides an implicit assumption of the absence of scale economies at the plant level and at least weak separability of factor inputs, current composition is likely to be correlated with productivity shocks, thus we have the possibility of endogeneity bias. In addition to identification problems, omitted variables and unaccounted heterogeneity in technology this creates some non-trivial estimation problems that we cannot at this point solve to our satisfaction. To some extent fixed effects take care of omitted variables and constant differences in technology (and capital stocks). Instrumental variable techniques are precarious since it is difficult to find good instruments that works at the plant level (see Malmberg et al. 2008). Here we have so far left endogeneity issues aside and thus cannot claim any causal interpretations.

4 Results

The results reported here are still preliminary and have not been thoroughly tested. Thus the results should be viewed with some caution but they may give some hints for further work along these lines.

Age-specific education levels

Education is unevenly distributed over age groups with the younger cohorts being generally more educated (and also more recently, which may be of importance in handling modern ICT technology for example). Thus we have a correlation between the variables used as independent variables that may cause a problem in identifying the effects. This problem is related to the familiar problem of identifying age, cohort and period effects separately, which is impossible in any strict sense without imposing restrictions of some kind on the data since we have an exact linear relation between these three kinds of effects. That is a consequence of age simply being time minus the date of birth.

Therefore, since younger cohorts are more educated and older cohorts less educated (and also since the value of the education of older cohorts may have depreciated) controlling for the level of education may actually confound the estimates since an increase in education generally will be correlated with a decrease in age. In Table 1 we report a specification where we have used overlapping variable cells and skipped some categories of workers. We report two variants one using only plant-specific effects and a regime dummy to account for the drop in overall employment due to the crisis in the beginning of the 1990s and one where we add time-specific effects. This makes very little difference (except for the R^2). Hours worked per employee has been added as a control for the fact that we cannot infer part-time or only part-of-year work for the individuals. For age shares we get a reasonable hump-shaped pattern with a peak in the age group 30-39. Although only the coefficient for primary education is significant and negative the expected pattern does appear with higher education giving relatively more positive effects.

Adding interactions in Table 2 yields partly confusing results. The age pattern now is all negative. Since the variables are logged we cannot strictly interpret 60+ as reference category and intercepts are not well defined. The pattern is that 30-39 now becomes the minimum, however, while 16-29 and 40-49 have more positive effects. The estimate for hours worked remains stable as do the constant and regime dummy parameters.. To make the table legible only significant education coefficients and interaction coefficients have been included.but it is hard to make sense of these estimates. Whether we use a full set of time dummies or only the regime dummy does not matter, however, so obviously the interaction terms tend to reverse age patterns The strongly negative effect from secondary education is also unexpected. Looking at the significant interaction coefficients it becomes clear that the most negative age groups 30-39 and 50-59 also have the most negative coefficients for interaction with secondary education. To make sense of this you must note that the logged shares are all negative variables and thus the interaction variables become positive. Thus the effect of for example 30-39 with secondary education is computed as

$$(-0.268 - 0.170 * \log \text{secondary education share}) * \log \text{age share 30-39}$$

thus the interaction offsets the negative effect of the 30-39 age group. An increase in the group will per se nevertheless have a negative effect per se since the log function is monotonically increasing but dependent on the concurrent change in educational composition it may be offset. Without a closer analysis of how this actually works out it is difficult to say whether the patterns are

Regressions made on industry/year-demeaned data				
Method	plant-specific effects		plant- and time-specific effects	
Dependent variable				
log value-added per employee	Coef.	*Std. Error*	Coef.	*Std. Error*
log hours worked per worker	0.327**	*0.013*	0.327**	*0.012*
log ageshare 16-29	0.025*	*0.010*	0.025**	*0.009*
log ageshare 30-39	0.045**	*0.011*	0.046**	*0.011*
log ageshare 40-49	0.020	*0.014*	0.019	*0.013*
log ageshare 50-59	0.005	*0.010*	0.004	*0.010*
log ageshare 60+				
log share primary education	-0.054**	*0.018*	-0.057**	*0.017*
log share secondary education	-0.022	*0.025*	-0.021	*0.023*
log share tertiary education, tech.	0.002	*0.002*	0.002	*0.002*
log share tertiary education, other	0.002	*0.002*	0.002	*0.002*
log share unknown education				
constant	0.028***	*0.003*		
regime dummy	-0.061***	*0.004*		
# of obs	27770		27770	
Adj R^2	0.004		0.029	

Table 1: Regressions with age and education shares as regressors. ** marks estimates that are significantly different from zero at the 1 percent level, * marks significance at the 5 percent level

reasonable or not. It is however not straight-forward how to do such an analysis since we have a complicated dependence between shares as such that one also has to take into consideration.[6] The cohort dependence of educational level also prevents us from asserting that we only pick up clearly identified age-eduction interactions. Our interpretation of the results in Table 2 is that interaction effects clearly are important but it is difficult to come up with any coherent explanation of the rather peculiar pattern we see.

The problems of interpreting the estimates when using overlapping categories motivates an attempt to cut down both the number of age groups and the number of education categories in the hope of getting clearer patterns. In Table 3 results are reported where age categories are restricted to young (-29), prime age (30-49) and old workers (50+) and education is only divided into primary, secondary and tertiary level. Since hours worked also correlate with age (younger workers tending to be more absent from work) we also skip that variable. This results in the sample with 50 or more employees increasing somewhat since we get fewer missing observations that way. In order to increase the precision of estimates we have also used the full sample in the first three columns of Table 3.

In the first column the pooled sample indicate clear hump shapes with older and less educated workers being most negative for productivity. As we control for plant-specific effects the coefficient estimates decrease and young and old workers shift place, the younger becomes more negative. The education ranking remains the same, however. Controlling further for time-specific effects decreases parameter estimates even more but there are no further changes in the pattern although differences are no longer all significant. In columns 4 and 5 we revert to a more limited sample of larger firms. With plant-specific effects this again strengthens the pattern but as we add time effects coefficients decrease and the pattern becomes more diffuse again. Thus our initial hunch that too small firms should not be used, because they either have less reliable data or are too small to be able to optimize workforce composition, seems vindicated. The hump shape of age coefficients is more or less expected from other studies. It is somewhat more puzzling to find in Table 5 that also the education productivity profile is hump-shaped with secondary education yielding clearly more positive effects than tertiary education. One potential explanation for this is that shares of tertiary educated people are very small in most plants. Taking the log of this then generates large negative values that may act as outliers if the functional form of the

[6] Since shares sum to one an increase in one share implies decreases in other shares that prevents a straight-forward calculation of marginal effects.

Regressions made on industry/year-demeaned data				
Method	plant-specific effects		plant- and time-specific effects	
Dependent variable				
log value-added per employee	Coef.	Std. Error	Coef.	Std. Error
log hours worked per worker	0.327**	*0.013*	0.327**	0.012
log ageshare 16-29	-0.113*	*0.057*	-0.114*	0.053
log ageshare 30-39	-0.268**	*0.076*	-0.271**	*0.070*
log ageshare 40-49	-0.002	*0.079*	-0.005	*0.073*
log ageshare 50-59	-0.203**	*0.055*	-0.203**	*0.051*
log ageshare 60+				
log share primary education	-0.393**	*0.133*	-0.398**	*0.123*
log share secondary education	-0.838**	*0.194*	-0.848**	*0.180*
log share unknown education				
16-29 * secondary edu	-0.123**	*0.035*	-0.124**	*0.033*
30-39 * primary edu	-0.159**	*0.037*	-0.161**	*0.034*
30-39 * secondary edu	-0.170**	*0.050*	-0.173**	*0.046*
30-39 * tertiary edu, other	-0.012*	*0.006*	-0.012*	*0.005*
50-59 * primary edu	-0.056	*0.023*	-0.055*	*0.021*
50-59 * secondary edu	-0.183**	*0.039*	-0.183**	*0.036*
constant	0.028**	*0.004*		
regime dummy	-0.060**	*0.003*		
# of obs	27770		27770	
Adj R^2	0.004		0.031	

Table 2: Regressions with age and education shares as regressors. ** marks estimates that are significantly different from zero at the 1 percent level, * marks significance at the 5 percent level

Dep var: log VA/empl	1	2	3	4	5
log age share -29	0.015	-0.025	-0.004	-0.029	0.014
	*(4.04)***	*(8.87)***	*-1.52*	*(3.65)***	*-1.71*
log age share 30-49	0.11	0.038	0.030	0.139	0.096
	*(17.24)***	*(6.89)***	*(5.56)***	*(6.77)***	*(4.68)***
log age share 50+	-0.018	0.013	0.001	0.078	0.015
	*(5.69)***	*(4.91)***	*-0.26*	*(7.35)***	*-1.34*
log share primary edu	-0.127	-0.029	-0.004	-0.088	0.004
	*(28.24)***	*(6.35)***	*-0.83*	*(6.71)***	*-0.28*
log share secondary edu	0.034	0.13	0.066	0.324	0.168
	*(5.61)***	*(22.94)***	*(11.20)***	*(18.90)***	*(8.41)***
log share tertiary edu	-0.021	0.004	0.003	0.021	0.012
	*(14.32)***	*(2.88)***	*(2.25)**	*(6.72)***	*(3.80)***
constant	4.139	0	0	0	0
	*(298.85)***	*0*	*0*	*0*	*0*
Plant size restriction				>49 empl	>49 empl
Fixed plant effect		Yes	Yes	Yes	Yes
Fixed time effect			Yes		Yes
Observations	93641	93641	93641	28583	28583

Table 3: Regressions with age and education shares as regressors. t-values within parentheses. ** marks estimates that are significantly different from zero at the 1 percent level, * marks significance at the 5 percent level.

relation is misspecified.[7] Another potential problem is the previously mentioned correlation between education and age. As younger age groups tend to have more tertiary education they may confound an age effect with an education effect. Of course, we also have to worry about the simultaneity in the determination of workforce composition and productivity in this case as well. A more positive interpretation would be to actually interpret the coeffients as elasticities of education. In a Cobb-Douglas production function. This would then correspond to factor cost shares and the low elasticity for tertiary education would not seem so strange since the share of tertiary education in manufacturing and mining indeed is small. At this stage we are still rather dubious that this interpretation really can be made consistent with other facts (as for instance in Table 1). Our approach in this data set thus seems to be limited to a rather coarse division of age groups and education in order to yield reasonably stable results. We have, however, experimented also using categories combining age and education level. In this specification we subdivide into age groups 16-29, 30-39, 40-49, 50-59, 60+ and then further subdivide according to education those with only primary education, secondary education or having tertiary education. With six age groups and three education categories this generates 18 different disjunct cells that an employee can belong to (in fact we have some with unknown education, due to classification problems especially in the 1980s, thus adding another six categories to that number). It is then definitely not very meaningful to include small plants in the sample, since they will in general have a lot of empty categories. Only plants with at least 50 employees have therefore been considered with this specification. As remarked above we get a lot of insignificant parameter estimates making any interpretation of the point estimates rather loosely grounded. Anyway, the most stable configuration turned out to be a regression of logged productivity on shares that were not logged.[8] Tables start to get very hard to read if we include all coefficients so we only report the coefficients graphically in Figure 5. The two estimations control for plant-specific effects and time-specific effects in one case. Also the sample has been trimmed by removing the lowest and highest percentile of the productivity distribution in order to avoid extreme values. Figure 5

[7] Some as yet unconfirmed results using straight age shares instead of logging indicate that this education pattern tends to disappear. This supports this hypothesis but has to be further checked.

[8] Such an interpretation would require a production function of the form

$$F(x_1, ..., x_n) = LA \exp(\beta_1 x_1 + ... + \beta_n x_n)$$

being strictly convex in all factor shares, implying that doubling all inputs would correspond exactly to doubling the output, no matter what the β_i sums to.

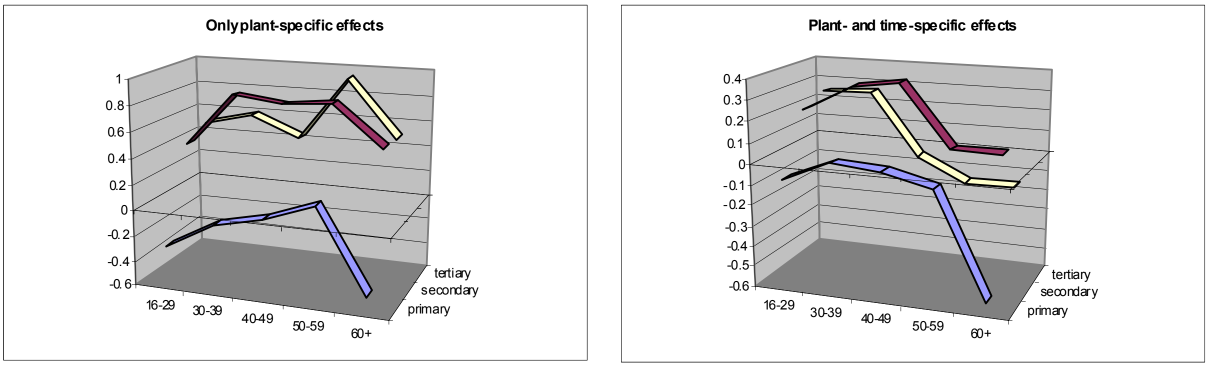

Figure 5: Estimated coefficient values with disjunct age-education shares. Log value added per employee regressed on shares in the workforce of each age-education category. 40-49 years old with only primary education has been chosen as reference category.

indicates much more reasonable age patterns in the specification with only plant-specific effects. However, secondary education still seems to dominate or at least be fairly equal to tertiary education. Adding time effects we get a rather implausible decrease with age in the effects of secondary and tertiary education. The problems encountered in achieving stable and significant estimates of education-specific age patterns makes it necessary to delay final conclusions until further research has been undertaken. Two tentative conclusions still seems to be warranted. One is that inclusion of time effects seems to be confounding the estimates such that the results hardly can be interpreted in a reasonable way. This could be due to our inability to efficiently control for cohort effects. Second, the peaks of education-specific patterns indicate productivity effects being more dependent on secondary education. This is not a unique result on Swedish data (see references above) but rather different from what we in general find in the literature. At this point we are reluctant to make any definite interpretation of this. The instability here may also be a result of an oversimplified view of how the optimal combination of age and education groups is determined. We have been working with specifications that are in effect crude linearized production relations. In practice this may have only local validity for the particular supply and demand conditions during the 12-year period we observe. Including interaction terms to catch cross-effects seems more confusing than enlightening though, so if that is the case we are not likely to find out using this approach.

5 Conclusion (Tentative)

One conclusion is rather clear from this study. Age and education composition explains only a minor part of the productivity variation in mining and manufacturing. The direct approach that we use here—to regress productivity on age group and education shares—is of limited use in order to actually resolve the productivity effects. Concerning age the seemingly most reliable estimates indicate that an ageing workforce would be less of a problem for productivity than a rejuvenating one, but there is no clear indication that education higher than secondary would provide much productivity impact at the plant level in manufacturing and mining. Although perhaps not wholly unexpected these results must still be regarded as highly uncertain awaiting further studies to resolve difficulties in the estimation. It would be premature to say much more than so at this stage of our investigation.

References

Andersson, B.: 2001, Portfolio allocation over the life cycle: Evidence from-Swedish household data, *Working Paper 2001:4*, Dept of Economics, Uppsala University, Uppsala. Work in progress.

Aubert, P. and Crépon, B.: 2003, La productivité des salariés âgés: Une tentative d'estimation, *Economie et Statistique* (368), 95–119.

Baily, M. N., Hulten, C. and Campbell, D.: 1992, The distribution of productivity in manufacturing plants, *Brookings Papers on Economic Activity: Microeconomics* pp. 187–249.

Bartelsman, E. J. and Dhrymes, P. J.: 1998, Productivity dynamics: U.S. manufacturing plants 1972-1986, *Journal of Productivity Analysis* **9**(1), 5–34.

Bartelsman, E. J. and Doms, M.: 2000, Understanding productivity: Lessons from longitudinal microdata, *Journal of Economic Literature* **38**(3), 569–594.

Beffy, M., Buchinsky, M., Fougère, D., Kamionka, T. and Kramarz, F.: 2006, The returns to seniority in france (and why are they lower than in the united states?), *Discussion Papers 5486*, CEPR.

Börsch-Supan, A., Düzgün, I. and Weiss, M.: 2007, Productivity and the age composition of work teams: Evidence from the assembly line. manuscript.

Davis, S. J., Haltiwanger, J. C. and Schuh, S.: 1996, *Job Creation and Destruction*, MIT Press, Cambridge, Massachusetts.

Doms, M. E., Dunne, T. and Troske, K.: 1997, Workers, wages and technology, *Quarterly Journal of Economics* **112**(1), 253–290.

Easterlin, R. A.: 1961, The American baby boom in historical perspective, *American Economic Review* **51**(5), 869–911.

Entorf, H., Gollac, M. and Kramarz, F.: 1999, New technologies, wages and worker selection, *Journal of Labor Economics* **17**(3), 464–491.

Entorf, H. and Kramarz, F.: 1998, The impact of new technologies on wages: Lessons from matching panels on employees and on their firms,, *Economics of Innovation and New Technology* **5**(2-4), 165–197.

Fienberg, S. and Mason, W.: 1985, Specification and implementation of age, period and cohort models, *in* S. Fienberg and W. Mason (eds), *Cohort Analysis in Social Research. Beyond the Identification Problem.*, Springer-Verlag, New York.

Forslund, A. and Lindh, T.: 1997, Bargaining, capital formation and unemployment: A putty-clay approach, *Working Paper 1997:18*, Dept of Economics, Uppsala University, Uppsala.

Gunnarsson, G., Mellander, E. and Savvidou, E.: 2004, Human capital is the key to the IT productivity paradox, *Working paper 2004:13*, Institutet för Arbetsmarknadspolitisk Utvärdering (IFAU), Uppsala.

Ilmakunnas, P. and Maliranta, M.: 2005, Technology, labour characteristics and wage-productivity gaps, *Oxford Bulletin of Economics and Statistics* **67**(5), 623–645.

Ilmakunnas, P., Maliranta, M. and Vainiomäki, J.: 1999, The roles of employer and employee characteristics for plant productivity, *Working Paper W-223*, Helsinki School of Economics and Business Administration, Helsinki.

Kremer, M. and Thomson, J.: 1998, Why isn't convergence instantaneous? Young workers, old workers, and gradual adjustment, *Journal of Economic Growth* **3**, 5–28.

Krueger, A. B. and Lindahl, M.: 2001, Education for growth: Why and for whom?, *Journal of Economic Literature* **39**(4), 1101–1136.

Lutz, W., Goujon, A., Samir, K. and Sanderson, W.: 2007, Reconstruction of populations by age, sex and level of educational attainment for 120 countries for 1970-2000, *Vienna Yearbook of Population Research* **2007**, 193–235.

Macunovich, D. J.: 1998, Relative cohort size and inequality in the United States, *American Economic Review* **88**(2), 259–264.

Malmberg, B.: 1994, Age structure effects on economic growth: Swedish evidence, *Scandinavian Economic History Review* **42**, 279–295.

Mellander, E.: 1999, The multi-dimensional nature of labor demand and skill-biased technical change, *Working Paper 518*, Industrial Institute of Economic and Social Research, Stockholm.

Morrison, C. J.: 1988, Quasi-fixed inputs in U.S. and Japanese manufacturing: A Generalized Leontief restricted cost function approach, *Review of Economics and Statistics* **74**, 275–287.

Nordström Skans, O.: 2002, Age effects in swedish local labour markets, *Working Paper 2002:11*, Institute for Labour Market Policy Evaluation, Uppsala.

Nordström Skans, O.: 2007, How does the age structure affect regional productivity?, *Applied Economics Letters* p. forthcoming.

Ohlsson, R.: 1986, *Högre utbildning och demografisk förändring*, Ekonomisk-historiska föreningen, Lund.

Shimer, R.: 1999, The impact of young workers on the aggregate labor market, *Working Paper 7306*, National Bureau of Economic Research, Cambridge, Massachusetts.

A Descriptive statistics (including the mean measures)

The descriptive statistics in Table 4 is for the initial sample with 50 employees or more, where plants with missing observations for some variable have been deleted. These data are computed on disjunct categories, i.e. without overlapping.

Variable	Mean	Std. Dev.	Min	Max
Value added per employee (1000's of SEK 1968)	85.4	63.6	0.1	1466.4
Hours worked per employee	1.1	0.3	0	4.1
Age 16-29 primary education	13.8	30.6	0	914
Secondary education	40.2	92.8	0	2978
Tertiary education, tech.	6.9	26.5	0	620
Tertiary education, other	1.5	5	0	146
Unknown education	3.2	14	0	769
Age 30-39 primary education	17.2	30.8	0	778
Secondary education	25.3	51.3	0	2086
Tertiary education, tech.	6.4	24.9	0	956
Tertiary education, other	2.9	9.8	0	301
Unknown education	3.1	14	0	731
Age 40-49 primary education	24.4	42.6	0	993
Secondary education	23.7	48.9	0	1135
Tertiary education, tech.	3.9	14	0	521
Tertiary education, other	3.1	9.9	0	299
Unknown education	1.2	6.4	0	353
Age 50-59 primary education	23.8	40	0	781
Secondary education	14.8	34.6	0	945
Tertiary education, tech.	1.5	6.3	0	255
Tertiary education, other	1.5	5.3	0	206
Unknown education	0.5	2.2	0	89
Age 60+ primary education	8.3	14.2	0	307
Secondary education	3.9	9.6	0	300
Tertiary education, tech.	0.4	1.8	0	74
Tertiary education, other	0.4	1.3	0	40
Unknown education	0.2	0.7	0	23

Table 4: Descriptive statistics for the reduced sample's 2770 observations on 3823 plants. Statistics have been calculated from the time averaged plant data.