Productivity and the age composition of work teams: Evidence from the assembly line*

Axel Börsch-Supan† and Matthias Weiss§

September, 7th, 2009

Abstract

This paper studies the relation between workers' age and their productivity in work teams. We explore a unique data set that combines data on errors occurring in the production process of a large car manufacturer with detailed information on the personal characteristics of workers responsible for the errors. We do not find evidence that productivity declines with age.

JEL codes: J24, J14, D24

Keywords: Age-Productivity Profiles

* We are grateful to Silke Januszewski, Arthur Kennickell, Melanie Lührmann, Jürgen Maurer, Steffen Reinhold, and Joachim Winter for valuable comments and suggestions. We thank Ismail Düzgün for his invaluable assistance in collecting the data. Verena Arendt, Melanie Gräsle, Klaus Härtl, Simon Hilpert, and Finn Körner provided excellent research assistance. Funding by the State Baden-Württemberg, the Gesamtverband der Deutschen Versicherungswirtschaft, and the Hans-Böckler-Stiftung is gratefully acknowledged. We are especially grateful to the Daimler truck assembly plant in Wörth, Germany. Without the great support from our contact persons in the different departments of the plant and from the side of the management and the works council, this project would not have been possible.

† MEA, Universität Mannheim, Germany and NBER.

§ Corresponding author. MEA, Universität Mannheim, Germany, E-Mail: *weiss@mea.uni-mannheim.de.*

1 Introduction

The age composition in most of the developed world has been shifting towards older age groups for more than 100 years now. The steady population aging has intensified dramatically in consequence of the direct succession of the post war “baby boom” and the “baby bust” of the late 1960s. This accelerated aging process will have far-reaching economic consequences. Most prominent in the public discussion are the consequences for the pay-as-you-go financed social security systems. But before the babyboomers are going to retire (with their pensions having to be financed by the babybusters), extensive changes are to be expected on labor markets and in production: In Germany, e.g., the share of workers aged 55 years and older will more than double from 12% in 2005 to almost 25% in 2035. In view of this looming evolution, it is important to better understand the relation between workers’ age and their labor productivity.

Estimating age-productivity profiles has been on the agenda of labor economists for a long time. The main problem with estimating age-productivity profiles is that it requires a valid measure for productivity. There are many studies in occupational medicine, cognitive psychology, and gerontology that look at how different abilities and skills of humans evolve over their life-cycle. They look at muscle strength, sight, retentiveness, the functioning of lungs, kidney, and the heart, and many other measurable indicators. More or less concordantly, they find that from the age of 25 onwards, physical and mental fitness are deteriorating.[1] But there is certainly more to labor productivity than muscle strength, sight, and cognitive ability. Experience plays a role and is increasing with age. This is illustrated in Figure 1. Physical and cognitive ability are hump-shaped over the life-cyle and peak early. Experience is increasing monotoneously in age. Productivity as a function of both, physical and cognitive ability on the one hand and experience on the other therefore peaks later in life. Hence, there is a need for more direct measures of productivity.

Regarding the measurement of productivity, the existing literature can be broadly divided into four branches: (i) studies relating plant level productivity to the age of the plants’ employees,[2] (ii) studies using individual’s wages as a productivity

[1] This literature is surveyed in Skirbekk (2004) and Börsch-Supan et al. (2005).

[2] E.g., Hellerstein, and Neumark (2004), Hellerstein et al. (1999), Haltiwanger, Lane, and Spletzer (1999) and (2007) for the U.S., Aubert (2003), Crépon, Deniau, and Pérez-Duarte (2003), Aubert and Crépon (2007) for France, Hellerstein, and Neumark (1995) for Israel, Grund and Westergård-Nielsen (2005) for Denmark, Ilmakunnas and Maliranta (2005) and (2007), Daveri and Maliranta (2007) for Finland, Malmberg et al. (2005) for Sweden, Dostie (2006) for Canada, Prskawetz et al.

measure,[3] (iii) studies using interviews of managers on their employees' performance,[4] and (iv) studies using direct measures of individual productivity like, e.g., the number and quality of publications in academic research,[5] the value of artists' paintings (in terms of auction proceeds),[6] or performance in sports and chess.[7]

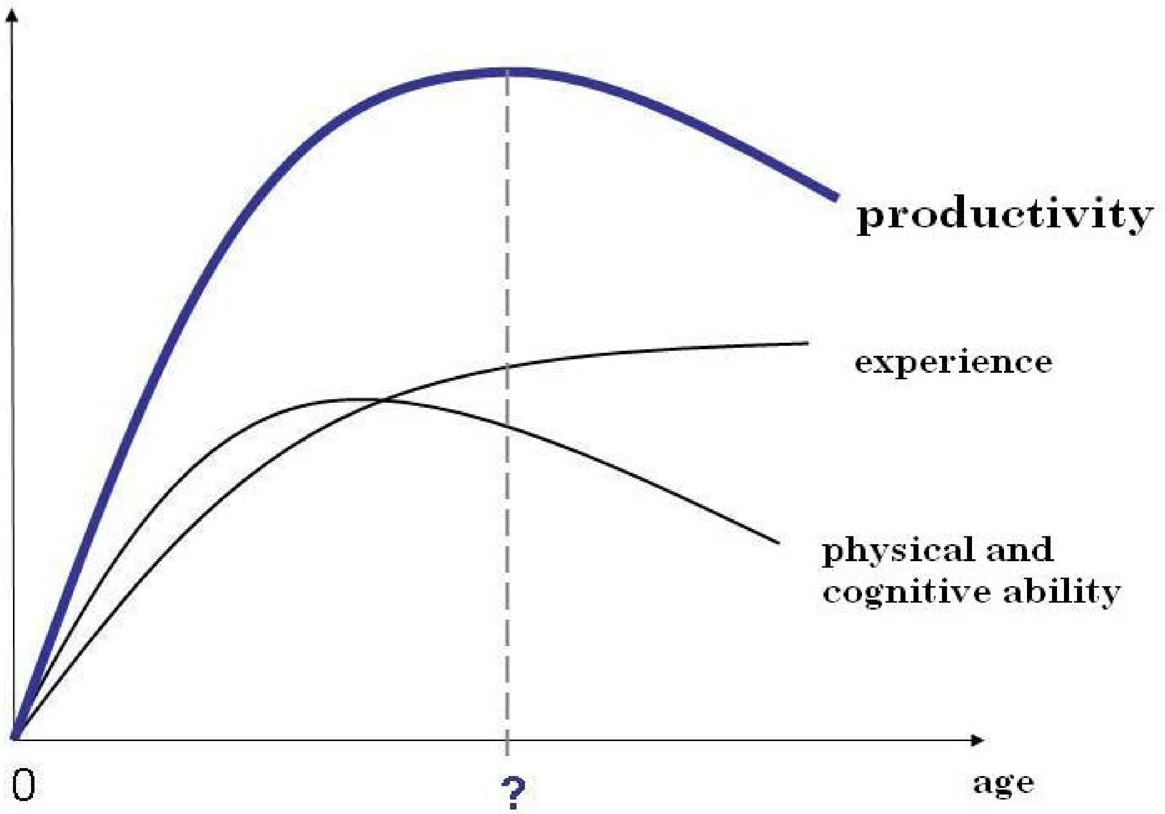

Figure 1: Schematic representation of productivity as a function of physical and cognitive ability and experience.

These different approaches all have their vices and virtues. Plant level productivity can be measured easily and reliably but the level of aggregation is quite high when

(2006) for Austria and Sweden, Lallemand and Ryckx (2009) for Belgium, van Ours (2009) for the Netherlands, Schneider (2007) and Goebel and Zwick (2009) for Germany.

3 E.g., Kotlikoff and Wise (1989), Kotlikoff and Gokhale (1992), and Laitner and Stolyarov (2005).

4 E.g., Medoff and Abraham (1980), Hunter and Hunter (1984), McEvoy and Cascio (1989), Salthouse and Maurer (1996), and Schneider and Stein (2006).

5 Jones (2005), Weinberg and Galenson (2005), van Ours (2009).

6 Galenson and Weinberg (2000) and (2001), Galenson (2005) and Bayer et al. (2009).

7 Fair (1994), (2005), and (2007), van Ours (2009), and Castellucci et al. (2009).

the goal is to study the relation between productivity and age. Furthermore, the age structure of firms is probably not exogenous.

Wages are the obvious productivity measure in many applications (returns to schooling, inter-personal comparisons, etc.) but when it comes to age profiles, the problem is that in many occupations, wages increase with age and/or seniority independently of productivity. Wage *de*creases are extremely rare.[8] Therefore, Kotlikoff and Wise (1989) look at life-time earnings profiles of insurance salesmen whose wages are proportional to the number of insurance contracts they sell. In addition, they look at starting rates of pay in cross-sections of office workers, managers and sales men who were hired at different ages. Thus, they obtain earnings age profiles that are independent of seniority wage schemes.

Supervisors' assessments are problematic as they might reflect prejudices about age productivity profiles.

The studies subsumed as approach (iv) are able to measure productivity relatively exactly. Therefore, they can estimate age-productivity profiles quite precisely. But the occupations where this approach is feasible are rare and particular so that the results can hardly be generalized. Another downside of this approach is that most of the studies refer to top performances of, eg, athletes or scientists which peak relatively early in life. These results cannot be generalized to everyday work life. Most workflows are organized such that they do not depend on top performance. Assembly lines, eg, run at a speed that is low enough to guarantee a largely error-free operation. In other words, the workflow is customised to everyday performance rather than top performance.

In addition, approaches (ii) through (iv) cannot take into account the fact that workers often work in teams and thereby affect one another's productivity. More specifically, if, e.g., older workers devote some of their working time to helping younger workers, the individual approach will underestimate older workers' productivity. Related aspects are workers' contributions to their team's work climate or how they deal with hectic situations (which again affects the productivity of the entire team).

In this paper, we look at productivity at the level of work teams. This takes into account the individual worker's contribution to his co-workers' productivity. We exploit a unique data set that we have compiled from an assembly plant of the

---

[8] Lazear (1979) and (1981) explains the increasing age-earning profiles with incentive effects. Loewenstein and Sicherman (1991) and Frank and Hutchens (1993) show in experiments that workers have a preference for increasing wage profiles and explain this with loss aversion and problems of self-control.

German car manufacturer Mercedes-Benz. At this plant, trucks are assembled by work teams on an assembly line. We have data on the number and severity of production errors that occur in the assembly process. We use the daily variation in the team composition of 100 work teams over 973 work days to estimate age productivity profiles. We deal with endogeneity of the age composition of work teams and with non-random sample selection.

2 The data

2.1 Our productivity measure and main explanatory variables

. As economists, we are used to defining labor productivity as the ratio of some measure of output to some measure of labor input. In the assembly plant we consider, the quantity of output is determined by the speed of the assembly line. The assembly line is divided into 50 workplaces that are located one after another. At each workplace there is one work team in the early shift and one work team in the late shift.[9] If work teams differ in productivity, this is not going to show up in differences in the quantity of output because the assembly line has the same speed for all teams. But production quality differs across work teams as they can make errors. Variation in productivity thus becomes manifest *only* in variation in production errors.

At the end of the assembly line, a quality inspector checks a random sample of about 10% of the trucks. The quality inspector is able to assign every error to the workplace where it occurred. At any time, there is exactly one work team at any workplace. Every error is given a weight (between 5 and 95) that specifies the severity of the error which depends on the costs respectively the time it takes to make up for the error. From this record of errors, we know which team has made how many errors of which severity on any day in 2003 through 2006. We observe 3824 workers in 100 work teams at 50 workplaces on 973 days. The number of teams is double the number of workplaces because on every day, there is an early and a late shift. Our productivity measure is the sum of errors per team per day where the errors are rated with their respective weights. E.g., if a team makes two

[9] At every workplace there is an „Team A“ and a „Team B“. A-Teams work early in even work weeks and late in uneven weeks. B-Teams respectively work early in uneven weeks and late in even weeks.

errors on a day with weights 5 and 30, our inverse productivity measure for this team for this day takes the value 35.

The information on errors is matched with personnel data that inform us about the daily composition of the work teams, personal characteristics of the workers such as age, sex, education, nationality, job tenure, and whether or not a worker is in his regular team. Note that errors are assigned to workplaces (and thus to work teams). The relation between errors and *individual* workers cannot be identified. We do not think that this is a disadvantage. After all, workers work in teams at this plant and they make errors in teams. If worker A "makes" an error, it might be his fault. But it might just as well be the fault of worker B who failed to do the preliminary work properly or the fault worker C who assisted inadequately.

In addition, we have data on the daily production plan which gives us information on the work load.

2.2 Matching error data and personnel data

The error data contain information on the work team where the error occurred. This information allows matching the error data with personnel data. However, not every error in our data set can be related to one single work team. For many errors, the quality inspector specified as locus delicti an area of the assembly line that encompasses the workplaces of several work teams. In other cases, the quality inspector was able to unambiguously specify the workplace but not whether the error occurred during the early shift or the late shift. In these ambiguous cases, we created an observation for each possible outcome and attributed weights to these observations according to their probability. For example, if an error is uniquely attributed to a workplace but cannot be related to early or late shift, we create one observation where we attribute the error to the team that worked at this workplace in the early shift and an additional observation where we attribute the error to the team that worked at this workplace in the late shift. Each of these two observations enters our regressions with weight 0.5. Figure 2 shows the distribution of these observation weights. Roughly one half of the observations have a weight equal to 1. The observation weight can have so many different values because the observation unit is a team-day. Suppose, e.g., there are two errors that were potentially made in team j on a certain day, one with probability 1/2 and one with probability 1/3. In this case, we create three observations, one, where team j makes no error (observation weight = probability that none of the two errors were made by team j =

$(1-1/2)\cdot(1-1/3)=1/3$), one with 1 error (observation weight = probability that one of the two errors occurred in team j = $1/2\cdot(1-1/3)+(1-1/2)\cdot 1/3=1/2$), and one with 2 errors (observation weight = probability that both errors occurred in team j = $1/2\cdot 1/3=1/6$). As work teams in our sample make up to eight errors per day, the number of possible values for the observation weight is large. Obviously, the observation weights must sum to 1 for each team-day.

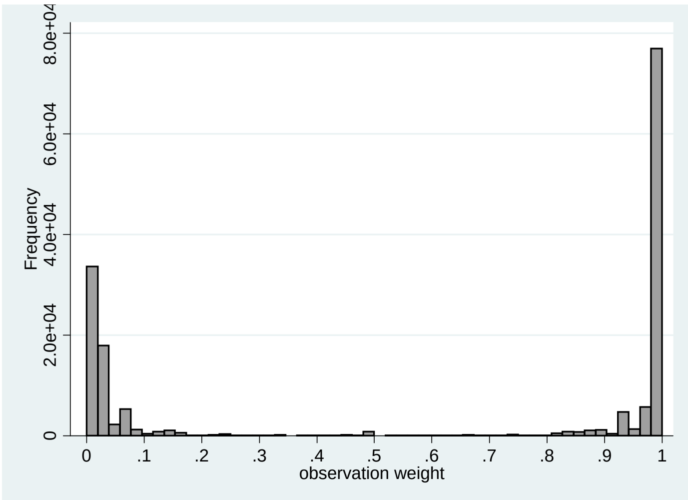

Figure 2: Distribution of observation weights

2.3 Some descriptives

As the data set we use is quite unique, this section gives a brief description of the main variables we use. Table 2 in Appendix A reports descriptive statistics of all variables used in the paper.

Errors

We observe 8564 errors in 100 teams on 973 days. Hence, errors are rather rare. The distribution of error weights (only for those days and teams for which we observe errors) is given in Figure 3. It shows that most errors are rather light.

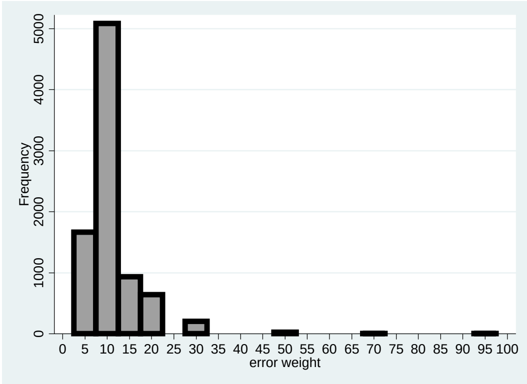

Figure 3: Distribution of error weights conditional on the observation errors

Age

The age composition in the plant is fairly representative for the German workforce in that workers older than 55 are rare. Figure 4 shows the age distribution in the plant (black) in comparison to the age distribution of the German population (grey).

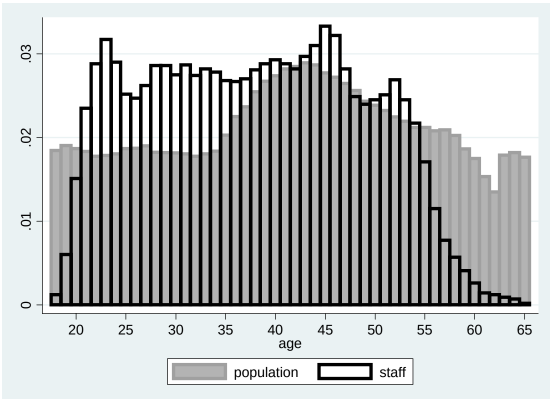

Figure 4: Age distribution in the plant (black) and in Germany (grey)

People younger than twenty are underrepresented because they are still in education or training. The share of workers aged 55 and over is low at the assembly line because many are already retired or have moved to better jobs. Only about 5% of the workers are older than 55 years. But still, they represent some 89,000 worker-day-observations (out of a total of 1,767,030 observations). Figure 5 shows the distribution of average age of work teams which constitute the observation unit in our regression analysis.

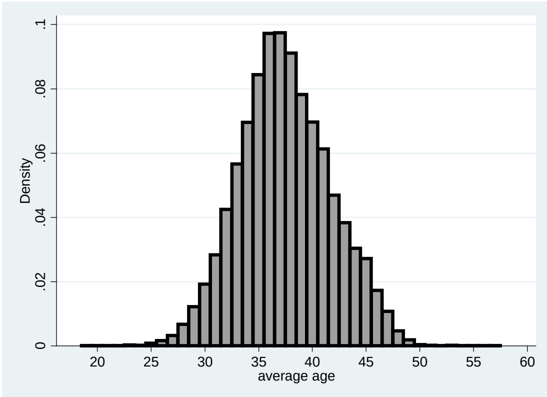

Figure 5: Distribution of average age of work teams

Job tenure

In addition to age, we have information on many characteristics of the workers. Workers' job tenure is a particularly important one indicating experience. Job tenure increases with age but the two variables are not perfectly correlated as workers are hired at different ages. The distribution of job tenure in the plant is shown in Figure 6. The spikes show hiring waves roughly every 5 to 10 years, the most recent having been just within the observation period (at job tenure=0).

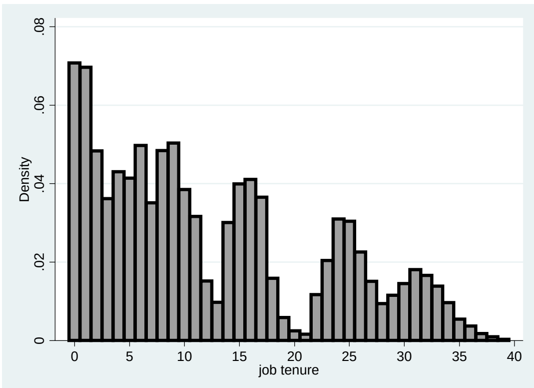

Figure 6: Distribution of job tenure in the plant

It is remarkable, that—as the distribution of average job tenure in work teams in Figure 7 shows—at hiring waves, the newly hired workers have been spread evenly over existing work teams: The histogram of average job tenure does not exhibit any comparable spikes.

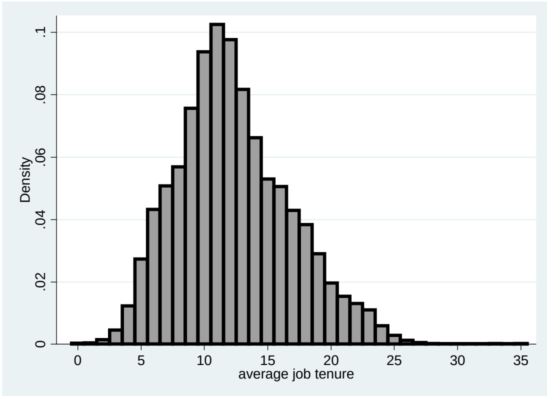

Figure 7: Distribution of average job tenure of work teams

Figure 8 shows the relation between age and job tenure in the plant. For any individual worker, age and job tenure are perfectly correlated over time, but as workers are hired at different ages, the overall correlation (over time *and* across workers) is “only” 0.78.

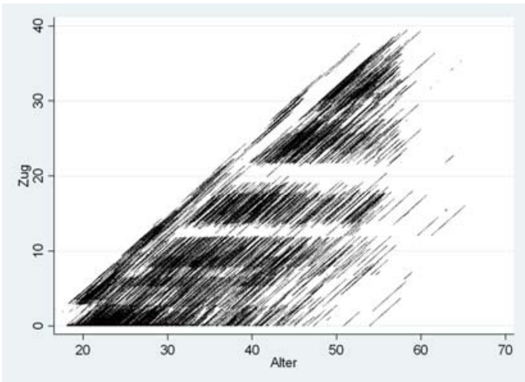

Figure 8: Scatter plot of job tenure (vertical axis) vs. age (horizontal axis)

The relation is tighter (correlation = 0.94) at the team level (see Figure 9 where *average* job tenure is plotted against *average* age of the work teams). This means that within teams, the correlation between age and job tenure is lower (0.75 on average).[10]

[10] One can think of the correlation between *average* age and *average* job tenure *between* work teams and the correlation between *individual* age and *individual* job tenure *within* work teams as a decomposition of the correlation between individual age and individual job tenure in the entire plant. If work teams are composed such that the correlation between age and job tenure within the teams is low than the correlation between average age and average job tenure between work teams must be high and vice versa.

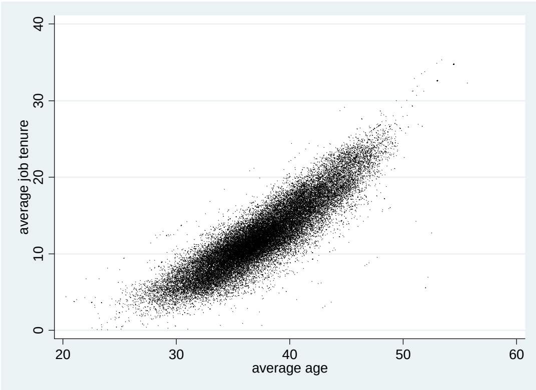

Figure 9: Scatter plot of average job tenure (vertical axis) vs. average age (horizontal axis) of work teams

Team size

The size of work teams varies between 4 and 35 workers. 90% of work teams have between 8 and 21 members (see Figure 10). The average team size is 14.

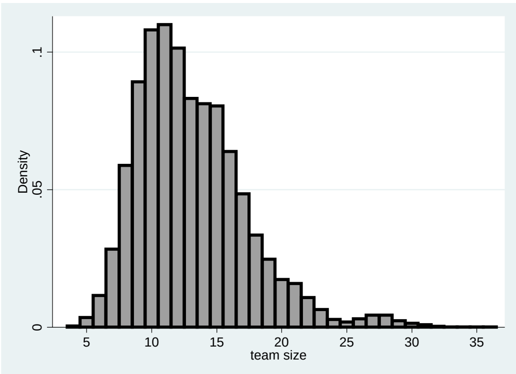

Figure 10: Distribution of team size in the plant

Sex

The share of women in the plant is very small and only 4.2%. In 63% of all work teams, there are only men (see Figure 11). In the other 37% of teams, women's share is 11.4% on average. The distribution of the number of women per work team is given in Figure 10. Within the sample period, the female share has increased by remarkable 87% from 3.0% in 2003 to 5.6% in 2006.

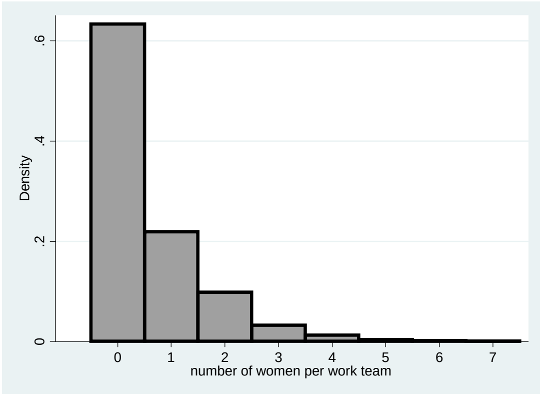

Figure 11: Distribution of the number of women per work team

Nationality

The composition of the personnel with respect to nationality is given in the following table:

nationality	German	French	Turkish	other
share	65.3%	26.0%	4.1%	4.6%

Workload

The production program and thereby the daily volume of work for every team varies over time (see Appendix D for a sample of the production program in 2003): Truck type A may be especially laborious for work team X (which assembles the axle suspensions) while truck type B may require complex and tedious work in team Y (which mounts the driver's cabs). So, on days where many trucks of type A (and few of type B) are produced, workload for work team X is high while on days where trucks of type B are superior in number, work team Y has a high workload. The required number of workers implied by the production program does not always exactly match the actual manning. We have daily information on the actual volume of work (measured in the number of required workers) and on actual manning for every day and every team. We use the percentage deviation of actual volume of work from actual manning as a measure of excess workload per worker. Figure 12 shows that the variation in excess workload is substantial.

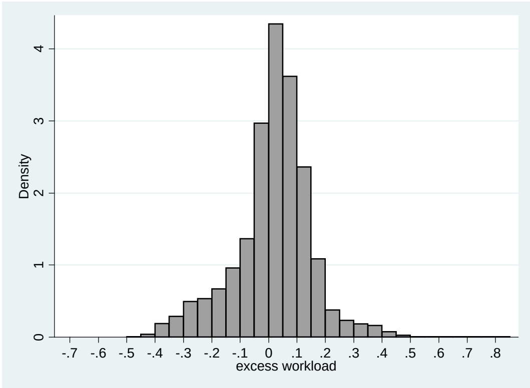

Figure 12: Distribution of excess work load (as a share of actual manning)

Cycle Time

One of the key variables in the assembly process is the cycle time, i.e. the time that workers have to perform their tasks on one car before the next car arrives at the workplace. The variation of cycle time within the plant is quite substantial.

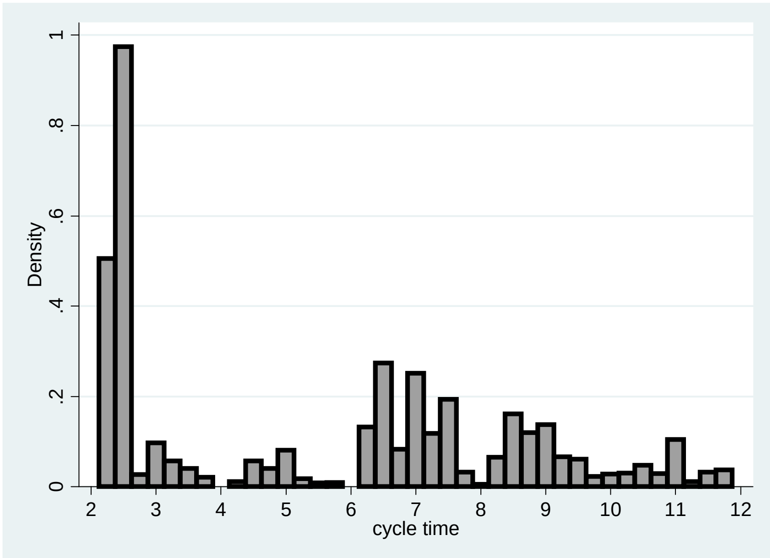

Figure 13: Distribution of cycle time across workplaces and over time

Figure 13 shows the distribution of cycle time over time and across workplaces in the plant. The part of the assembly line where the finishing of the driver's cabs is done exhibits very short cycle times (2 – 3 minutes) while the part where trucks are actually assembled is characterized by longer cycle times (6 – 10 minutes depending on the production program).

External workers

Each worker is assigned to one team as his "regular" team. But—due to fluctuations in team composition and workload—workers work outside their regular team 6% of the time on average. As can be seen in Figure 14, roughly one third of the workers always worked in the same team. The other two thirds of the workers change work teams over time.

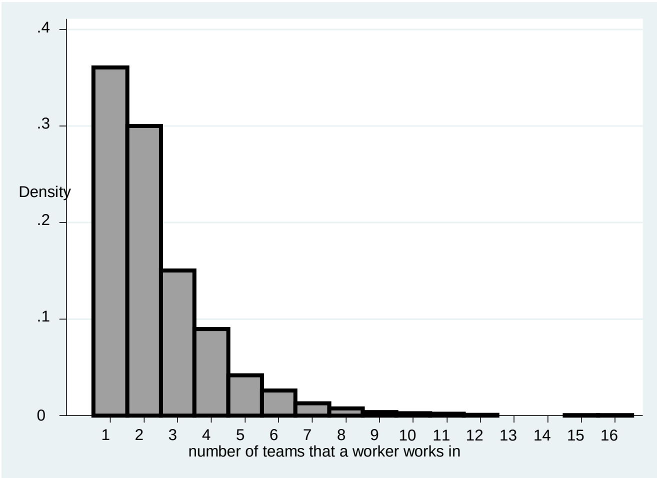

Figure 14: Distribution across workers of the number of teams that a worker works in over time

Those workers who do change work teams over time do so quite often. Figure 15 shows that work team changes occur up to 179 per worker within this 4-year period. The average number of team changes is 12 (18 for those who work in more than one team), the median number is 2 (8 for those who work in more than one team).

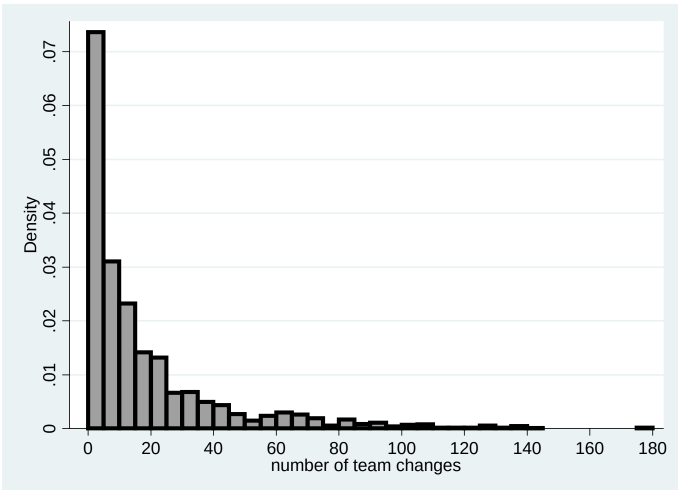

Figure 15: Distribution across workers of the number of team changes that a worker undergoes over time given that he works in moren than one work team

Figure 16 displays the distribution over time and across work teams of the share of workers external to the work team.

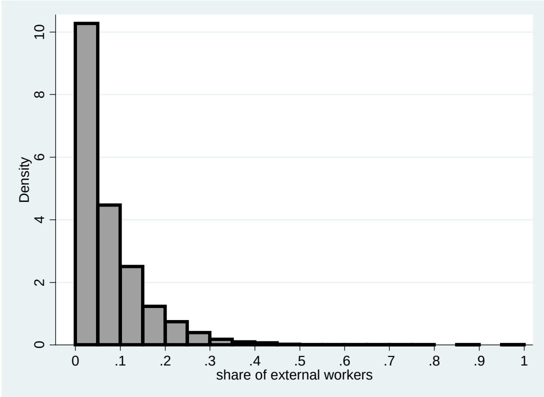

Figure 16: Distribution of the share of workers external to the work team

Fluctuation

The composition of work teams varies considerably over time. These fluctuations are due to variations in workload (demand side) and absence of workers (supply side). Absence of workers is due to vacation (12%), sickness (6%) and—as importantly—due to compensatory time off for extra hours worked (6%). Workers' contracts involve 7.5 hours per day while the assembly line runs 8 hours per work shift, so that workers accumulate one half hour overtime per day. Consequently, they can take every 16th day off. This means that in a team of 16 workers, on an average day, one worker is absent due to compensatory time off. In order to buffer these fluctuations, each work team has about 20% more members than are needed on a regular day. As a second means to level out these fluctuations, there is a pool of especially qualified workers who can fill in for absent workers. As a third possibility to accomodate fluctuations in workload and worker supply, regular workers may switch from their regular work team to another one. Fluctuation in the composition of work teams may be detrimental to productivity as communication is hindered if worker turnover is high. Figure 17 displays a variable of fluctuation that we use in our regression analysis. This variable is constructed as the number of consecutive days without change in the composition of the work team. As we are interested in the effects of fluctuation on productivity, we only consider changes that concern workers who do not regularly belong to the team. We do not consider changes within the regular team as we think that these daily changes should not affect communication or team work as within the regular work team, workers are used to working together in all possible constellations. In other words: We only consider fluctuation in the work team that is making use of what we called "possibilities 2 and 3" further up in this paragraph. A value of 4 of our variable means thus that in the four preceding days, no external worker joined or left the work team. Figure 12 shows that the team composition usually changes from day to day.

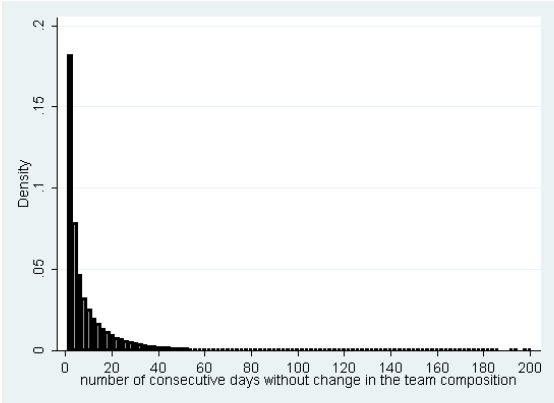

Figure 17: Distribution of the number of consecutive days without change in the team composition

3 Identification

Our aim in this paper is to relate the productivity of work teams to the age of the team members. The identification of this relation is potentially afflicted with the following problems:

- How can productivity be measured reliably? How can the productivity of different observation units be made comparable?
- How can we make sure that the variation in age that our estimates are based on is exogenous?
- How can we deal with potential sample selection bias in the presence of early retirement and career moves?
- How can we distinguish age effects from cohort effects?

In this section, we explain our identification strategy to deal with these potential problems.

3.1 Measuring productivity

We measure productivity inversely as the number and severity of production errors. These errors are observed at the level of work teams. Thus, our productivity measure takes into account the individual workers' contribution to their co-workers' productivity. This is important if older workers contributions are larger. In that case, looking at individual level productivity underestimates older workers' productivity. Examples for older workers' potential contribution to their team's productivity are the instruction of younger workers,[11] being relaxed in tense, hectic situations, contributing positively to the work climate, etc.

Are errors that occur at different places on the assembly line comparable? In some sense, they are because every error is given a severity weight. But still, a comparison across workplaces is probably unfair as some tasks are more error-prone than others. We take the fact that workplaces are different with respect to the susceptibility to errors into account by using work team fixed effects. In other words, to identify the relation between workers' age and the errors they make, we only use variation in errors and age over time within work teams where the tasks are homogeneous. As we have nearly 1,000 observations per work team (973 work days), we do not depend on comparisons across work teams. [12]

3.2 Exogeneity of variations in age

Studies on age and productivity that use plants as units of observation generally suffer from the problem that the age composition of plants is endogenous: The more productive firms are usually more profitable. They expand and increase their workforce. This leads to a rejuvenation of their workforce as new hires are more likely to be young. Relating productivity to the age of the workforce in this case results in spurios (negative) correlations between productivity and age. Studies on individuals or work teams within plants potentially suffer from a similar problem if age affects the assignment of workers to tasks: If, e.g., older workers are systematically assigned to easier tasks, their productivity may be overestimated.

We observe productivity (errors) on the level of work teams. The allocation of workers across work teams (i.e., workplaces) may be endogenous: Older (supposedly less productive) workers may be systematically allocated to workplaces that are less error-prone (where they can do less harm). But the variation *across* workplaces is not used in our estimation as we include work team fixed effects (see Subsection 3.1). The variation in the age composition *within* work teams over time—which we use to

[11] While an older worker helps a younger worker, the older worker's productivity is zero as he is not producing anything at that time. But the contribution to the work team's productivity is clearly positive.

[12] We also tried workplace fixed effects instead of work team fixed effects as the task within one workplace are equal for early and late shift. (Remember that on each day at each workplace, there are two work team, one in the early shift and one in the late shift.) However, because of potential endogeneity of the work team composition with respect to early vs. late shift, we use work team fixed effects (see Subsection 3.2).

identify our estimates—is exogenous. This may sound surprising, but—as the managers in the plant explained to us—the fluctuation within work teams over time (which results from sick leave, vacation, compensatory time off for extra hours worked and from workforce changes induced by changes in workload) does not leave any room for optimization. The fluctuation within teams has two components: The team composition changes from day to day as some workers call in sick and others return from vacation. This day-to-day fluctuation within the core team is random and not the result of any management decision. The second component of fluctuation is the employment of so-called "team hoppers". On days where the workers of the core team are too few to manage that day's workload, the vacancies are filled randomly with workers that do not belong to that team. At this stage, there is—in principle—room for optimization. But this would require knowledge about how the optimal age composition for that day's production program for the respective work team would look like.[13] Furthermore, if the optimal team composition on a certain work day requires a 37 year old worker it is quite unlikely that a 37 year old worker will be available on that very day. In addition, this optimization behavior would have to make a real difference in terms of production errors in order to justify the effort. The managers in the plant are convinced that neither requirement is met.[14]

Thus, we conclude, that the variation in the age composition that we use in our estimation is exogenous.

3.3 Sample selection bias

One recurrent criticism of studies on age-productivity profiles is that—due to sample selection bias—they overestimate the productivity of older workers. In fact, workers older than 55 years are underrepresented in the workforce. The obvious suspicion is that the remaining workers are a positive selection. The less motivated, less healthy workers probably retire earlier or are made redundant.

In our paper, we are able to correct for this potential bias in two ways: A Heckman-style selection correction model and worker fixed effects. The common problem with the correction for sample selection is that—by definition—we usually do not have information on those subjects who are not in the sample. As our sample covers sev-

[13] Table D.1 in Appendix D contains the production program for 2003 as an example. It shows that the combination of the 13 different types of trucks changes every two weeks.
[14] Rember that the point here is not that managers try to arrive at an optimal team composition for a given workplace. This is taken care of by the work team fixed effects. Endogeneity would be a

eral years, we do have information on those workers who enter or exit the sample within these three years. This enables us to estimate a Heckman-style selection correction model.

As our observation unit in the regression is a work team while selection into the sample is an individual phenomenon, we have to aggregate individual Mills ratios to team Mills ratios (see Appendix C).

In addition, selection into (and out of) our sample is more complex than selection out of the workforce of an entire plant or economy. Our sample consists only of workers on the assembly line. Even the foremen are not included. Workers who leave our sample before the age of 65 may retire early (possibly because they are less motivated or less healthy than those who remain) or they may be promoted to jobs off the assembly line (possibly because they performed better than those who remain). We therefore constructed different Mills ratios for younger workers (who are more likely to leave the sample for jobs off the assembly line) and for older workers (who are more likely to leave the sample for early retirement (see Appendix C for details).

A second possibility to correct for non-random sample selection is to include worker fixed effects. This is possible if worker-days are the observation unit and if—over time—workers move across work teams. This second possibility turns out to work very well (see Section 4.2).

3.4 Age vs. cohort effects

A common problem in estimating age effects (be it on productivity, on consumption, savings or other variables) is that in a cross section of individuals, age effects are indistinguishable from cohort effects (at least without strong assumptions on the functional form). In a panel, where each cohort is observed at different ages, the distinction becomes possible. However, as individuals are observed over time, the potential existence of time effects may be confounded with age effects. From discussions with the plant managers, we conclude that time effects do not play any role in our sample as there have not been any changes in technology or organization that could affect our productivity measure during the four years of our sample period. There was a major change before 2003 and no further change until 2006. The absence of any changes in technology or organization during these four years is the main reason for the choice of this time period.

problem only if managers would try to optimize the team composition from day to day based on the day-to-day variation in the production program.

4 Results

In this section, we present the results from our regressions. The regressions in Section 4.1 have work team-days as the observation unit. In Section 4.2, we present results from regressions where we use individual worker-days as unit of observation. Using work team-days as observation units is straightforward given that we observe errors on a daily basis on the work team level. Regressions on the individual level are worthwhile for the following reasons:

1. Dealing with the distinction between cohort effects and age effects becomes possible.
2. Dealing with sample selection is easier and more powerful.
3. We can look at higher ages.

Identifying age effects on the individual level is possible, because workers move a lot between work teams. Figure 14 displays the distribution of the number of work teams that workers work in during our observation period of four years. Roughly one third of the workers has worked in only one team, another third of the workers has worked in two teams and one third of the workers has worked in more than two teams. Figure 15 shows that those workers who work in more than one work team switch back and forth between their teams. The median number of team changes for them is 8. This movement of workers across work teams allows us to identify age effects on the individual level even though errors are observed on the level of work teams.

Worker fixed effects remove differences between workers that are constant over time. This implies that cohort effects (which are just one form of differences between workers (of different cohorts) that are constant over time) are removed. The remaining variation can be due to age effects and time effects. Time effects are unlikely as within the observation period, there have not been any changes in technology or organization in the plant. The tasks have not changed and are absolutely comparable across time. Thus, we conclude that the effects we find are age effects.

Worker fixed effects also help remove sample selection bias. If selection into (and especially out of) the sample is related to differences in health, motivation, etc. between workers that are constant over time, then the bias that results from this non-random selection is removed as we estimate our coefficients on the variation within workers over time which is not affected by this bias.[15]

On the work team level, the 5 percent oldest team observations (4,784 obs.) have an average age between 45 and 51 years (see also Figure 5). The 5 percent oldest individual observations (83,802 obs.) have an age between 54 and 65 years (see also Figure 4). On the individual level, we have thus the possibility to estimate an age productivity profile that ranges from 18 years to 65 years.[16]

These benefits of individual level regressions do not come at no cost. Regressions at the individual level make sense only if worker fixed effects are controlled for. In this case, age and job tenure are perfectly collinear and their effects cannot be studied seperately. At the team level, average age and average job tenure are only imperfectly correlated even within work teams as the team composition changes from day to day (see Figure 9). Therefore, the effects of age and job tenure can be studied in regressions on the team level. Another advantage of the team-level regressions is that the data set is smaller. This allows us to include many control variables and interactions of control variables with age without reaching the limits of computing power. Thus we can study which variables have an effect on the age-productivity profile (see Table 1).

Table 1: Roadmap Section 4

	Section 4.1	Section 4.2
observation unit	team day	worker day
work team fixed effects included	yes	yes
worker fixed effects included	no	yes
interactions with age included	yes	no
job tenure included	yes	no
number of controll variables	large	small

4.1 Regressions on the team level

This section reports regression results on the relation between workers' age and the number and severity of errors made in the truck assembly process. As explained in

[15] There remains a selection bias if selection is associated with differences in the age-productivity profile (rather than differences in the productivity level). See Appendix E where we perform some robustness checks to show that workers who leave the sample do not have an age productivity profile that is decreasing more steeply.

[16] See Figure 4. From age 60 on, standard errors become quite large, though.

Section 2.1, our inverse productivity measure is the weighted sum of errors per team per day where each error is given a weight according to its severity. The observation unit in the regressions is a team day. We observe 3824 workers in 100 work teams on 973 days. As—along the assembly line—workplaces differ quite substantially and the allocation of workers to these workplaces and to early vs. late shift may be endogenous, we control for work team fixed effects. Only the day-to-day variation within work teams is used to identify our estimates. This variation results from fluctuations in the work team composition due to vacation, sickness and—most importantly—due to compensatory time off for extra hours worked (see Section 2.3). From discussions with managers at the plant on how they replace missing workers, we conclude that this variation is truly exogenous. There is no optimization taking place at this level.

In order to allow for non-linear age effects, we use a piecewise linear specification (5-year linear splines). We also tried other specifications (polynomials, dummies for 5-year age groups. The results are robust with respect to these different specifications.

As explained in Section 2.2, we have artificially inflated the error data set in order to be able to uniquely match the error data with the personnel data. We have assigned observation weights to these additional observations such that for each team-day, the observation weights sum to 1. We use these weights in the regression. The reported numbers of observations refer to the non-inflated data set (and are equal to the sum of error weights).

The left column of Table 2 shows the results of our baseline regression. We also include a number of interactions with average team age. Therefore, the coefficients on the age splines cannot be interpreted as marginal effects. The estimated marginal effects including the interaction effects are reported in Table B.1 in Appendix B. The upper left chart in Figure 18 displays the (inverse) age productivity profile of work teams that results from the estimation. The weighted sum of errors per day increases up to an average age of 30 years and stabilizes thereafter. Column 2 of Table 2 and the upper right chart in Figure 13 display the results of a regression where we correct for possible sample selection bias using a Heckman-style approach (see Appendix C). Figure 18 shows that the slight decrease of errors with age between average age 30 and 45 turns into a slight increase. This finding indicates that workers who remain in the sample are indeed a positive selection. But the decrease in errors after average age 45 remains.

Table 2: Regression results: sum of error weights (team level)

dependent variable: sum of error weights

	baseline specification		correcting for sample selection		correcting for selection controlling for job tenure	
age splines						
20 – 25 years	5.75	(0.000)	6.10	(0.000)	8.71	(0.000)
25 – 30 years	5.15	(0.000)	5.50	(0.000)	6.08	(0.000)
30 – 35 years	4.87	(0.000)	5.22	(0.000)	5.66	(0.000)
35 – 40 years	4.88	(0.000)	5.21	(0.000)	5.60	(0.000)
40 – 45 years	4.88	(0.000)	5.19	(0.000)	5.64	(0.000)
45 – 50 years	4.45	(0.000)	4.77	(0.000)	5.48	(0.000)
50 – 55 years	5.61	(0.000)	5.83	(0.012)	8.27	(0.001)
job tenure splines						
0 – 4 years					-1.57	(0.000)
4 – 8 years					0.006450	(0.914)
8 – 12 years					-0.0965	(0.007)
12 – 16 years					0.0788	(0.034)
16 – 20 years					-1.73	(0.001)
20 – 24 years					-0.209	(0.011)
24 – 28 years					-0.971	(0.002)
28 – 32 years					-4.31	(0.227)
control variables						
schooling years	2.49	(0.000)	2.78	(0.000)	2.89	(0.000)
car specific educ	-5.42	(0.014)	-3.39	(0.123)	-2.46	(0.293)
tech spec. educ	0.517	(0.831)	2.07	(0.386)	3.76	(0.126)
female	23.2	(0.000)	21.8	(0.000)	23.9	(0.000)
external	-0.673	(0.803)	-2.13	(0.424)	-2.33	(0.386)
team size	0.357	(0.001)	0.125	(0.254)	0.094	(0.386)
(team size)^2	-0.00541	(0.000)	-0.00587	(0.000)	-0.00544	(0.000)
late shift	-0.0743	(0.874)	-0.0318	(0.938)	-0.150	(0.970)
days w/o change	0.0327	(0.044)	0.0342	(0.035)	0.0352	(0.031)
(days w/o change)^2	-0.000043	(0.184)	-0.0000485	(0.133)	-0.0000494	(0.125)
cycle time	0.312	(0.000)	0.283	(0.002)		(0.002)
workload	2.14	(0.057)	4.38	(0.053)	2.66	(0.241)
(workload)^2	-0.390	(0.075)	-1.68	(0.057)	-1.74	(0.050)
tryout Axor	-3.75	(0.001)	-3.83	(0.001)	-3.70	(0.002)
tryout Atego	2.51	(0.030)	2.44	(0.034)	2.48	(0.032)
French	-6.09	(0.158)	-6.88	(0.111)	-3.36	(0.449)
German	-10.9	(0.005)	-10.4	(0.008)	-6.01	(0.128)
Turkish	-9.53	(0.118)	-5.85	(0.340)	-0.719	(0.912)
temperature	0.191	(0.013)	0.164	(0.033)	0.160	(0.038)
temperature^2	-0.00545	(0.012)	-0.00559	(0.009)	-0.00547	(0.011)
humidity	-0.137	(0.000)	-0.139	(0.000)	-0.142	(0.000)
hours of sunshine	-0.0517	(0.510)	-0.363	(0.661)	-0.0387	(0.621)
rainfall	0.517	(0.000)	0.513	(0.000)	0.512	(0.000)
air pressure	0.0515	(0.099)	0.0502	(0.105)	0.0526	(0.090)
Monday	9.21	(0.000)	8.14	(0.000)	8.00	(0.000)
Tuesday	8.29	(0.000)	7.21	(0.000)	7.03	(0.000)
Wednesday	13.88	(0.000)	12.8	(0.000)	12.6	(0.000)
Thursday	8.99	(0.000)	7.90	(0.000)	7.73	(0.000)
Friday	12.8	(0.000)	11.7	(0.000)	11.5	(0.000)

In column 3 of Table 2, we control for average job tenure in the work team. The effect of age on productivity (i.e. errors) can now be decomposed in an "experience effect" and a "pure age effect".[17] As can be seen in the lower left chart of Figure 18, the "pure age effect" is positive. Older work teams make slightly more errors if job tenure is held constant.

Table 1 cont'd: Regression results: sum of error weights (team level)

dependent variable: sum of error weights

	baseline specification		correcting for sample selection		correcting for selection controlling for job tenure	
interactions of age with...						
schooling years	-0.0661	(0.000)	-0.0708	(0.000)	-0.0744	(0.000)
car specific educ	0.145	(0.015)	0.0986	(0.098)	0.0851	(0.173)
tech spec. educ	-0.0195	(0.756)	-0.0543	(0.391)	-0.0905	(0.167)
female	-0.614	(0.000)	-0.601	(0.000)	-0.658	(0.000)
external	0.0197	(0.794)	0.0274	(0.693)	0.0334	(0.640)
team size	0.0197	(0.955)	0.00533	(0.039)	0.00569	(0.028)
late shift	-0.000856	(0.948)	-0.00197	(0.864)	-0.00242	(0.814)
days w/o change	-0.000884	(0.038)	-0.000919	(0.031)	-0.000950	(0.025)
cycle time		(0.069)		(0.142)		(0.202)
workload	-0.0250	(0.389)	-0.0472	(0.413)	-0.00240	(0.956)
tryout Axor	0.0979	(0.009)	0.0817	(0.007)	0.0784	(0.010)
tryout Atego	-0.0553	(0.067)	-0.0552	(0.068)	-0.0562	(0.070)
French	0.111	(0.343)	0.151	(0.191)	0.0671	(0.569)
German	0.249	(0.019)	0.257	(0.015)	0.143	(0.184)
Turkish	0.231	(0.153)	0.164	(0.317)	0.0282	(0.870)
temperature	-0.00490	(0.016)	-0.00417	(0.040)	-0.00406	(0.050)
temperature^2	0.000133	(0.021)	0.000130	(0.023)	0.000126	(0.027)
humidity	0.00342	(0.000)	0.00349	(0.000)	0.00356	(0.000)
hours of sunshine	0.00281	(0.180)	0.00248	(0.238)	0.00257	(0.222)
rainfall	-0.0125	(0.000)	-0.0124	(0.000)	-0.0124	(0.000)
air pressure	-0.00167	(0.040)	-0.00166	(0.040)	-0.00173	(0.035)
Monday	-0.210	(0.000)	-0.192	(0.000)	-0.189	(0.000)
Tuesday	-0.186	(0.000)	-0.168	(0.000)	-0.163	(0.000)
Wednesday	-0.317	(0.000)	-0.299	(0.000)	-0.295	(0.000)
Thursday	-0.203	(0.000)	-0.185	(0.000)	-0.180	(0.000)
Friday	-0.287	(0.000)	-0.268	(0.000)	-0.264	(0.000)
constant	-7.96	(0.000)	-7.97	(0.000)	-7.97	(0.000)
Inverse Mills Ratio young			0.215	(0.000)	0.217	(0.000)
Inverse Mills Ratio old			0.313	(0.000)	0.302	(0.000)
adj. R^2 within	0.041		0.042		0.044	
adj. R^2 between	0.081		0.081		0.081	

observations: 95,684 (unbalanced panel of 100 work teams on 973 work days). p-values are in parentheses. All specifications control for work team fixed effects.

[17] What we call "pure age effect" here is of course again a composition of other effects that come along with age like deteriorating health, declining cognitive abilities, etc.

In the lower right chart, the weighted sum of errors is plotted against average job tenure. Holding average age constant, work teams with longer average job tenure (i.e. more experience) have a lower sum of error weights. For workers who grow old in the plant, the productivity enhancing effect of growing experience (job tenure) compensates the adverse “pure” age effect so that the overall age profile is rather flat.

Table 2 also displays the coefficients on a wealth of control variables. The second part of Table 2 (on page 23) contains coefficients on the interactions of these control variables with age. The marginal effects of these variables at ages 20, 40, 50, and 60 are given in Table B.2 in Appendix B.

As the focus of this paper is on the relation between productivity and age, we only comment on some interesting interaction effects of age with these variables.

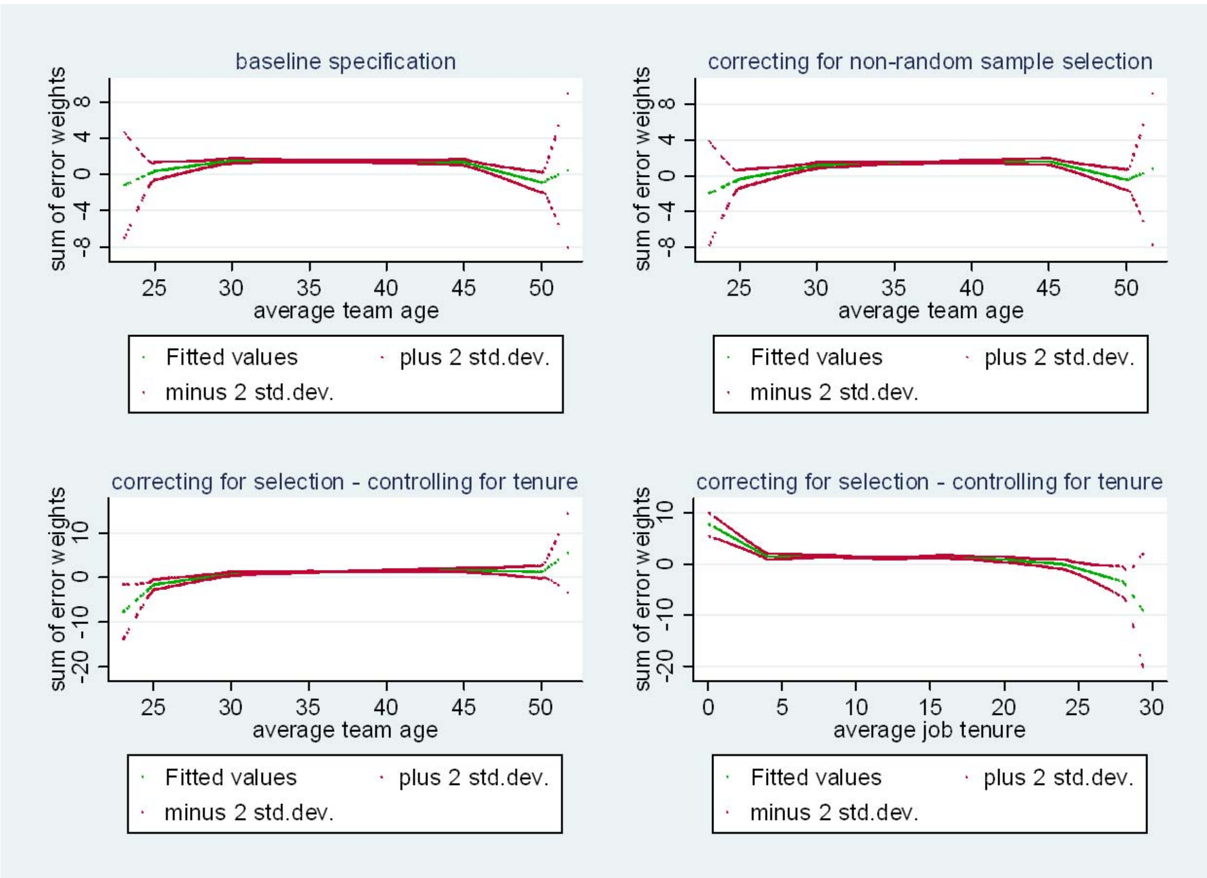

Figure 18: Some of weighted errors: Inverse age productivity profiles on the team level

Up to an age of 35, the average number of schooling years has a negative effect on productivity. Presumably, workers who have spent long time in school are overqualified for and bored of the tasks on the assembly line. This effect dies away (and even reverses) as workers grow older. A higher share of workers having a car specific education is good for productivity at least for young workers. This effect also wears out (and reverses) as worker age. It should be noted here, that the interaction with age might also reflect differences between cohorts rather than age-specific differences. School education and vocational training have changed over time and the interaction effects might reflect these differences. In Section 4.2, we discuss the problem of distinguishing between age and cohort effects.

A higher share of women in the work team is bad for productivity in young teams and good in old teams. One explanation for this finding is that women make fewer errors but young male workers get distracted (and make more errors) if women are in the team.

The presence of external workers improves productivity; presumably because these "team hoppers" are more experienced. This effect wears out in older teams where additional experience has less benefit.

In large work teams, the sum of error weights is larger. Reasons for this effect may be lower team cohesion and impeded communication. Older workers seem to have more problems with large work teams.

The weighted sum of errors is larger during the early shift. Working early in the morning seems even harder for older workers.

Fluctuation (inversely measured as "days without change in the team composition) seams to affect only older workers productivity adversely.

Excess workload leads to more errors. This effect is less pronounced in older work teams. It seems that experience helps in tense situations.

Another interesting explanatory variable is the cycle time, i.e. the time that workers have to perform their tasks on one car before the next car arrives at their workplace. The negative sign of the coeffient on the interaction with age implies that older workers find it hard to deal with short cycle times. However, the coefficient is very small and mostly insignificant. This finding rebuts concerns that our results might be specific to truck assembly where cycle times are rather long while in the production of passenger cars, where cycle times are shorter (1.5 – 5 minutes), age-productivity profiles might be less favorable for older workers (see Figure 19).

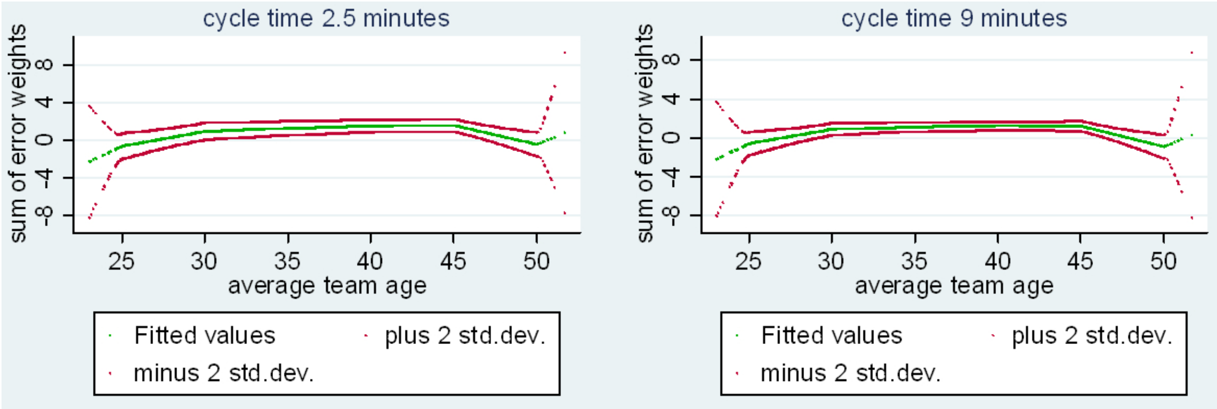

Figure 19: Age-error profiles for short and long cycle times

4.2 Regressions on the individual level

In this subsection, we present results from regressions where we use individual worker-days as unit of observation. Even though, we observe errors only on the team level, regressions on the individual level are worthwhile. More concretely, it has three advantages to look at the individual level:

4. Dealing with the distinction between cohort effects and age effects becomes possible.
5. Dealing with sample selection is easier and more powerful.
6. We can look at higher ages.

Identifying age effects on the individual level is possible, because workers move a lot between work teams. Figure 20 displays the distribution of the number of work teams that workers work in during our observation period of four years. This movement of workers across work teams allows us to identify worker fixed effects in addition to work team fixed effects.

Worker fixed effects remove differences between workers that are constant over time. This implies that cohort effects (which are just one form of differences between workers (of different cohorts) that are constant over time) are removed. The remaining variation can be due to age effects and time effects. Time effects are unlikely as within the observation period, there have not been any changes in technology or organization in the plant. The tasks have not changed and are absolutely comparable across time. Thus, we conclude that the effects we find are age effects.

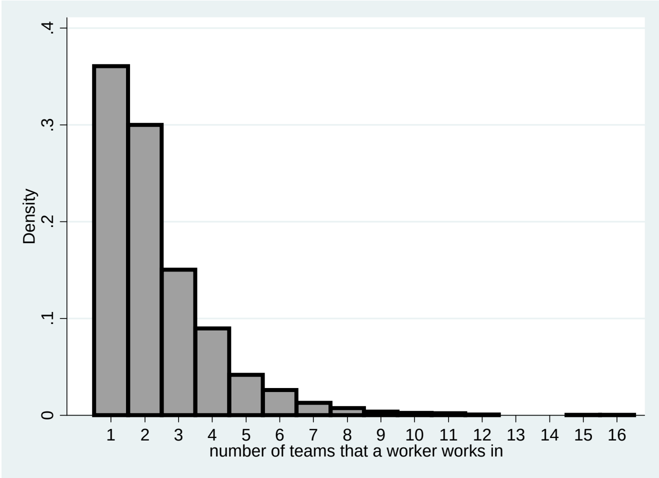

Figure 20: Distribution of the number of teams that a workers works in over time

Worker fixed effects also help remove sample selection bias. If selection into the sample is related to differences in productivity between workers that are constant over time, then the bias that results from this non-random selection is removed as we estimate our coefficients on the variation within workers over time which is not affected by this bias.

On the work team level, the 5 percent oldest team observations (4,784 obs.) have an average age between 45 and 51 years (see also Figure 5). The 5 percent oldest individual observations (83,802 obs.) have an age between 54 and 65 years (see also Figure 4). On the individual level, we have thus the possibility to estimate an age productivity profile that ranges from 18 years to 65 years.[18]

Table 3: Regression Results (individual level)

dependent variable:	sum of error weights				number of errors		error intensity	
	basic specification		correcting for sample selection				(given an error occurred)	
age splines								
15 – 20 years	0.858	(0.000)	0.857	(0.000)	0.0690	(0.000)	0.142	(0.613)
20 – 25 years	0.143	(0.000)	0.143	(0.000)	0.0134	(0.000)	-0.294	(0.000)
25 – 30 years	0.0193	(0.168)	0.0198	(0.225)	0.00517	(0.000)	-0.340	(0.000)
30 – 35 years	-0.0317	(0.017)	-0.0310	(0.025)	0.0000618	(0.983)	-0.324	(0.000)
35 – 40 years	-0.0416	(0.002)	-0.0452	(0.001)	0.000529	(0.666)	-0.385	(0.000)
40 – 45 years	0.0216	(0.097)	0.0189	(0.159)	0.00592	(0.000)	-0.389	(0.000)
45 – 50 years	-0.0462	(0.000)	-0.0464	(0.000)	0.000917	(0.416)	-0.433	(0.000)
50 – 55 years	0.0136	(0.323)	0.0116	(0.415)	0.00484	(0.000)	-0.422	(0.000)
55 – 60 years	-0.0314	(0.312)	-0.0557	(0.149)	-0.00168	(0.581)	-0.402	(0.000)
60 – 65 years	-0.0564	(0.653)	-0.0484	(0.693)	-0.000256	(0.974)	-0.195	(0.630)
control variables								
workload	0.371	(0.000)	0.371	(0.000)	0.0320	(0.000)	-0.388	(0.000)

[18] From age 60 on, standard errors become quite large.

workload^2	-0.270	(0.000)	-0.270	(0.000)	-0.0208	(0.000)	1.06	(0.000)
cycle time	0.0148	(0.000)						
team size	0.0601	(0.000)	0.0601	(0.000)	0.08388	(0.000)	0.0385	(0.151)
(team size)^2	-0.00126	(0.000)	-0.00126	(0.000)	-0.0000772	(0.000)	0.000451	(0.580)
external	0.0393	(0.110)	0.0388	(0.127)	0.00485	(0.017)	-0.133	(0.025)
late shift	-0.100	(0.000)	-0.100	(0.000)	-0.00838	(0.000)	0.107	(0.000)
days w/o change	-0.000854	(0.004)	-0.000851	(0.004)	-0.000951	(0.000)	0.00315	(0.000)
tryout Axor	-0.147	(0.000)	-0.147	(0.000)	-0.0142	(0.000)	0.111	(0.053)
tryout Atego	0.0334	(0.170)	0.0335	(0.171)	0.00610	(0.002)	-0.174	(0.001)
Monday	1.13	(0.000)	1.13	(0.000)	0.108	(0.000)	-1.41	(0.000)
Tuesday	1.10	(0.000)	1.10	(0.000)	0.103	(0.000)	-1.33	(0.000)
Wednesday	1.36	(0.000)	1.36	(0.000)	0.124	(0.000)	-1.38	(0.000)
Thursday	1.03	(0.000)	1.03	(0.000)	0.098	(0.000)	-1.40	(0.000)
Friday	1.09	(0.000)	1.09	(0.000)	0.105	(0.000)	-1.28	(0.000)
Inverse Mills Ratio young			0.0125	(0.931)	0.000823	(0.946)	-0.319	(0.360)
Inverse Mills Ratio old			0.157	(0.265)	0.00799	(0.484)	0.389	(0.249)
adj. R^2 within	0.003		0.003		0.004		0.005	
adj. R^2 between	0.438		0.438		0.519		0.031	
# observations:	1,676,030		1,676,030		1,676,030		150,772	

Unbalanced panel of 3,824 workers in 100 work teams on 973 work days. *p*-values are in parentheses. All specifications control for individual worker fixed effects and work team fixed effects.

Table 3 reports results from regressions on the individual level. Note that we have more than 1.5 million observations. Due to constraints regarding computing power and memory, we do not include interactions with age in this regression and only the most important controls. We cannot control for job tenure here as on the individual level (after controlling for fixed effects), age and job tenure are perfectly collinear.[19] Age effects are again specified as 5-year-age splines. Figure 21 shows the estimated (inverse) age productivity profiles where productivity is again measured as the sum of error weights per day. After a sharp increase up to the age of 25, the sum of error weights declines slightly but monotonously up to the age of 60. Thereafter, the mean prediction still declines but the standard errors become too large so that the decline is not significant.

The second column in Table 3 and the right panel of Figure 21 show the results for the regression where—in addition to worker fixed effects—we correct for non-random sample selection by including inverse Mills ratios from the selection equation (see Appendix C 3). The results are virtually unchanged indicating that the worker fixed effects essentially remove the selection bias.

In columns 3 and 4 of Table 3, we decompose our productivity measure in the frequency of errors (the number of errors per day) and the severity of errors (given that an error occurred). For the frequency of errors, we find a clearly increasing profile:

Older workers make significantly more errors. This can also be seen in Figure 22. On the other hand, the severity of errors is strongly decreasing with age (see Figure 23).

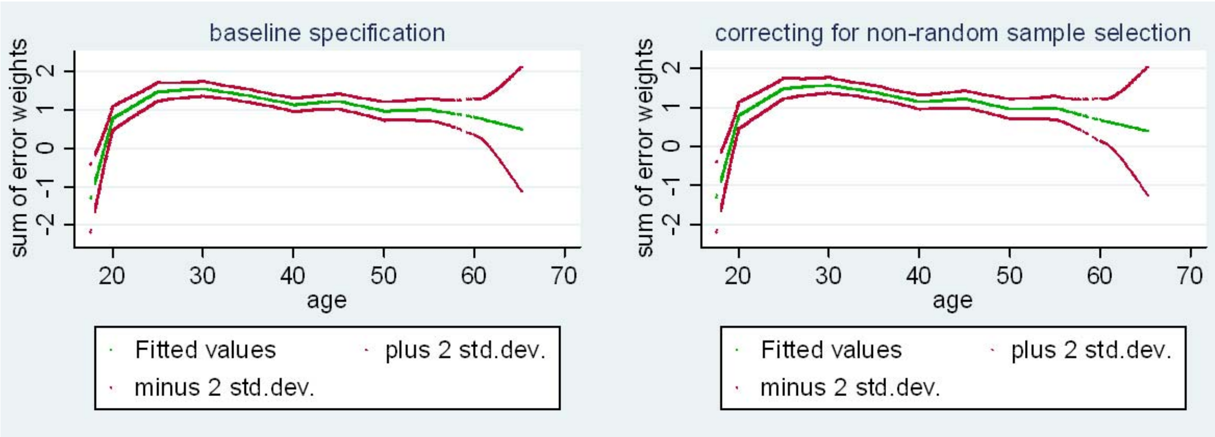

Figure 21: Age productivity profiles on the individual level (in terms of errors)

Our interpretation of these results is as follows: Errors are rare. They usually happen in special (tense) situations when maybe something went wrong and there is little time to fix it and do the regular tasks. In these situations, older more experienced workers seem to know better, which (severe) errors to avoid by all means. This concentration on the vital tasks secures that older workers perform better also in terms of our overall productivity measure sum of error weights.

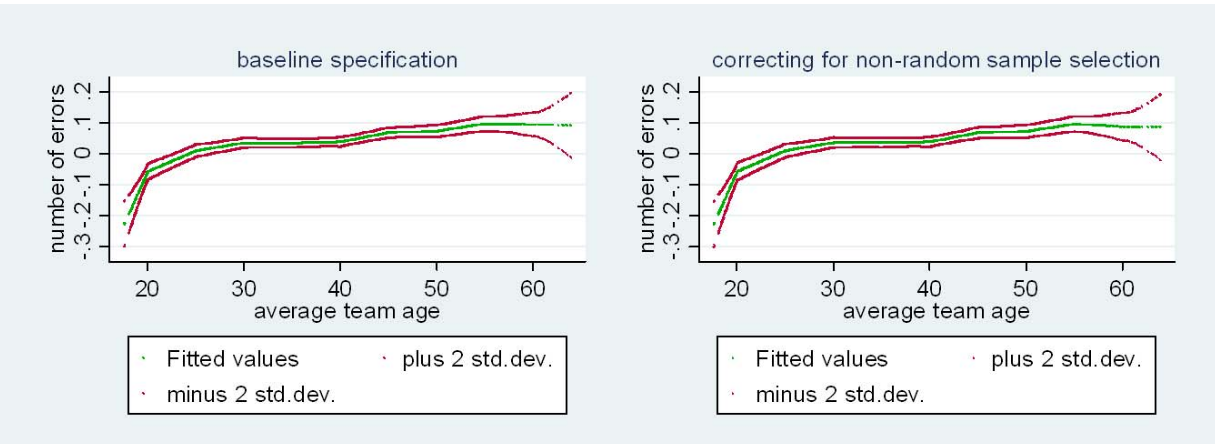

Figure 22: Age profiles for the frequency of errors

[19] Even with non-linear specifications the problem of multi-collinearity is too severe.

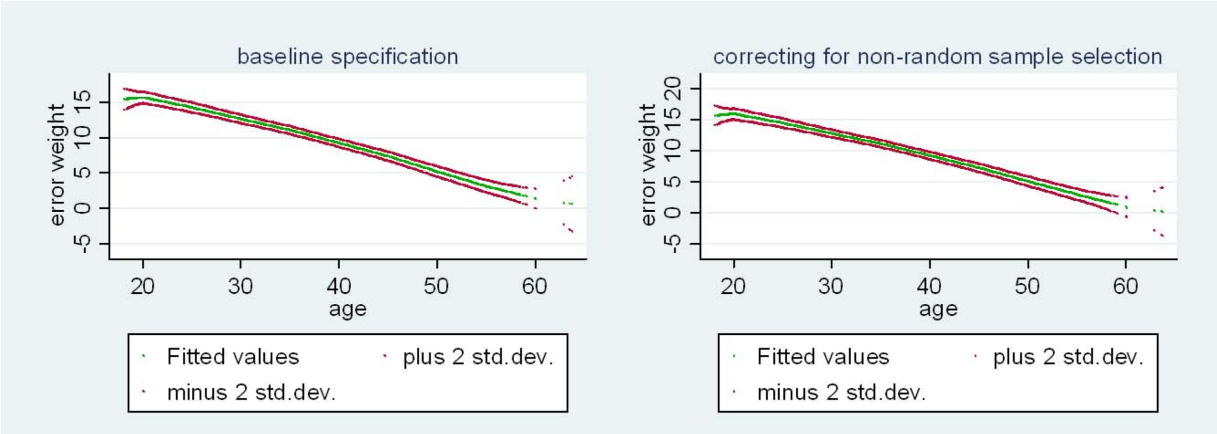

Figure 23: Age profiles for the severity of errors

5 Conclusion

We study the relation between the age structure of work teams and their performance in an assembly plant of a car manufacturer. We use data on errors made in the production process to construct our (inverse) productivity measure. As the production quantity is given by the speed of the assembly line, which is equal across all workplaces, work teams that are differently productive only differ in the errors they make. From these quality data, we know for all 100 work teams the number and severity of errors they made on any day in 2003 through 2006. We combine these data with data from the personnel department that gives us the daily composition of work teams with personal characteristics of the workers, especially age. In addition, we have information on the daily work load.

Controlling for individual worker fixed effects allows us to correct for sample selection bias. In addition, it prevents us from confounding age and cohort age effects. Controlling for work team fixed effects guarantees that the remaining variation we use to estimate the age productivity profile is exogenous: The fluctuation in the worker composition within work teams over time is random. This fluctuation is due to sick leave, vacation and the compensation for overtime and does not leave any room for optimization.

Our findings suggest that productivity is highest among the workers below 30 years of age. For workers older than 30 years, the age productivity profile is fairly stable (slightly increasing if anything). A decomposition into the effect of job tenure and a "pure age effect" reveals that it is indeed job tenure that keeps older workers productivity from falling. A decomposition of our productivity measure into the frequency of errors and error severity shows that the older workers' competence is their ability to avoid especially severe errors. While older workers are slightly more likely to make errors, they hardly make any severe errors. The results suggest that older workers are especially able to grasp difficult situations and then concentrate on the vital tasks.

The huge data set and the truly exogenous variation in team composition enable us to estimate age productivity profiles quite precisely. In addition, we are able to correct potential sample selection bias. On the other hand, our results refer to a single plant only. However, we believe that our results are of general interest. Regarding our estimates of the age-productivity profile, we find it interesting that even at the assembly line—where labor is physically demanding and experience should be comparatively unimportant as tasks are rather simple and do not require substantial training—productivity does not decline at older ages.

Appendix

A Descriptive statistics

Table A.1: Descriptive statistics of the variables used in the regressions

Variable	Mean	Median	Min	Max	Std. Dev.
Date	Jan 29th 2005	Dec 20th 2004	Jan 7th 2003	Dec 20th 2006	
# errors	0.0895	0	0	8	0.598
error intensity	11	10	0	95	5.7
weighted sum of errors	0.984	0	0	135	5.73
individual age	37.1	36.9	17.5	65.2	10.5
average team age	37.1	36.8	23	51.7	4.33
individual job tenure	11.7	10.7	0	39.4	9.92
average team tenure	11.7	11.2	0.0865	31.3	4.55
female dummy	0.0418	0	0	1	0.208
share of women	0.0418	0	0	0.554	0.0662
ind. years of schooling	11.3	11	9	20	2.16
av. years of schooling	11.3	11.3	9	16.6	0.826
dummy for technical training	0.367	0	0	1	0.482
share of workers with technical training	0.367	0.364	0	1	0.16
dummy for car specific training	0.254	0	0	1	0.434
share of workers with car specific training	0.254	0.222	0	1	0.183
team size	14.4	14	4	36	4.44
German dummy	0.653	1	0	1	0.473
share of Germans	0.653	0.662	0	1	0.163
French dummy	0.26	0	0	1	0.434
share of French	0.26	0.25	0	1	0.155
Turkish dummy	0.0410	0	0	1	0.192
share of Turkish	0.0410	0	0	0.418	0.0532
dummy for external workers	0.0718	0	0	1	0.235
share of external workers	0.0718	0.0594	0	1	0.0859
individual inverse Mills ratio young	0.297	0.204	0	3.19	0.354
team inverse Mills ratio young	5.34	4.63	0	20.1	3.33
individual inverse Mills ratio old	0.098	0	0	3.21	0.172
team inverse Mills ratio old	1.76	1.37	0	12.3	1.32
# days without change in team composition	10.2	4	1	200	15.4
dummy for late shift	0.489	0	0	1	0.5
dummy for Axor tryout	0.0634	0	0	1	0.244
dummy for Atego tryout	0.0651	0	0	1	0.247
excess work load	0.0163	0.0291	-0.458	0.826	0.134
max temperature (C°)	17.6	18	-6.8	40.2	9.5
air humidity (%)	75.2	76	31	99	14.3
hours of sunshine	5.75	5.3	0	15.1	4.57
precipitation (mm)	2.15	0	0	76.5	6.12
air pressure (hPa)	1000.4	1000.5	968.5	1022.3	7.35

B Significance tests for marginal effects

In this appendix, we report tests on the significance of the gradient of the age variables as well as some control variables. Table B.1 reports the test results for the regressions presented in Table 1. These regressions include a set of interactions with age. The regression equation is given by

$$\text{sum of errorweights} = \beta_0 + \sum_{a=20}^{50} \beta_{a,a+5} \cdot AgeSpline_{a,a+5} + \sum_k \beta_k \cdot x_k + \sum_k \beta_{age,k} \cdot x_k \cdot AverageAge + \varepsilon \quad (1)$$

where subscripts for work teams and days are omitted for the sake of clarity. x_k are the control variables. and $x_k \cdot AverageAge$ are interactions of these control variables with average team age.

The gradient the errors-age-profile (evaluated at sample means) is thus the linear combination of the coefficients on the interactions (where coefficients on interaction terms are multiplied by the sample means of the respective variables) and the coefficient on the respective age spline. The gradient of the error-age profile at age 37 is for example given by:

$$\left.\frac{\partial(\text{sum of errorweights})}{\partial AverageAge}\right|_{AverageAge\ =\ 37} = \beta_{35-40} + \sum_k \beta_{age,k} \cdot \overline{x_k} \quad (2)$$

where $\overline{x_k} = \frac{1}{I \cdot T} \cdot \sum_i^I \sum_t^T x_{kit}$ is the sample mean of variable x_k. These gradients and their significance levels are reported in Table B.1.

Table B.1: Age gradients

20 - 25 years	0.915	(0.225)
25 - 30 years	0.316	(0.000)
30 - 35 years	0.0380	(0.061)
35 - 40 years	0.0221	(0.130)
40 - 45 years	0.00791	(0.699)
45 - 50 years	-0.415	(0.000)
50 - 55 years	0.640	(0.613)

Gradients are calculated from coefficients in the second column of Table 1. *p*-values are in parentheses.

Table B.2 reports the marginal effects of selected control variables at different ages.

Table B.2: Marginal effects of control variables at different ages

	years of schooling	car specific training	share of women	share of team hoppers	team size	late shift	days without change	excess workload
20 years	1.36 (0.000)	-1.42 (0.023)	9.79 (0.000)	-1.58 (0.039)	0.103 (0.001)	-0.0712 (0.549)	0.0149 (0.001)	3.44 (0.000)
30 years	0.656 (0.000)	-0.435 (0.167)	3.78 (0.000)	-1.31 (0.001)	0.156 (0.000)	-0.0909 (0.109)	0.00570 (0.021)	2.97 (0.000)
40 years	-0.0519 (0.270)	0.551 (0.021)	-2.24 (0.000)	-1.04 (0.000)	0.209 (0.000)	-0.111 (0.001)	-0.00349 (0.048)	2.50 (0.000)
50 years	-0.760 (0.000)	1.53 (0.003)	-8.25 (0.000)	-0.763 (0.171)	0.262 (0.000)	-0.130 (0.136)	-0.0127 (0.000)	2.02 (0.000)
60 years	-1.47 (0.000)	2.52 (0.003)	-14.3 (0.000)	-0.490 (0.610)	0.316 (0.000)	-0.150 (0.322)	-0.0219 (0.000)	1.55 (0.036)

Gradients are calculated from coefficients in the second column of Table 1. *p*-values are in parentheses.

C Sample selection

C.1 The problem

Older workers are underrepresented in our sample. This might lead to a bias in the estimation of the age productivity profile if the selection into the sample is non-random with respect to productivity and age. There are two possible mechanisms of sample selection that are related to productivity:

- Early retirement
- Stepping up the career ladder

If those workers who are less motivated, less healthy, and less productive are more likely to retire early then those workers who remain in the sample are a positive selection. Early retirement thus potentially leads to an overestimation of the productivity of older workers relative to younger workers. If workers who are more productive are more likely to be promoted to jobs off the assembly line then those who remain in the sample are a negative selection. Selection due to careers thus potentially leads to an underestimation of the relative productivity of older workers.

We try to correct this sample selection bias in two ways:

1. Worker fixed effects
2. Correction of selection bias à la Heckman (1979)

C.2 Worker fixed effects

Workers differ in productivity. If sample selection is related to these differences (and to age), the estimation of the age productivity profile in a cross section is biased. Controlling for worker fixed effects in the estimation removes the bias that results from differences between workers that are constant over time.

C.3 Correction of selection bias à la Heckman (1979)

We have non-random selection and the selection is different for old and young. Workers at the assembly line are not a random sample of the working age population. There is selection based on age (which is not a problem) but there is probably also selection based on something correlated with productivity (motivation, etc.). Younger workers may exit the sample if they are good enough to get a job outside the assembly line. Older workers may exit the sample if they are *not* good enough to keep working.

C.3.1 Different selection for young and old

We observe a person i at date t if he is still working at the assembly line. Suppose that younger workers i remain in the sample ($s^y = 1$) if some latent variable $z_{it}' \cdot \gamma^y + \varepsilon_{it}$ is positive:

$$s_{it}^y = 1\left[z_{it}' \cdot \gamma^y + \varepsilon_{it} > 0\right], \quad \varepsilon_{it} \sim N(0,1) \text{ i.i.d.} \tag{3}$$

Accordingly, selection for older workers is

$$s_{it}^o = 1\left[z_{it}' \cdot \gamma^o + \varepsilon_{it} > 0\right], \quad \varepsilon_{it} \sim N(0,1) \text{ i.i.d.} \tag{4}$$

For given z_{it}, the workers with high ε_{it} are observed. The probability that person i is observed is

$$P\left(z_{it}' \cdot \gamma + \varepsilon_{it} > 0\right) = P\left(\varepsilon_{it} > -z_{it}' \cdot \gamma\right) \overset{\text{symetry}}{=} P\left(\varepsilon_{it} < z_{it}' \cdot \gamma\right) = \Phi\left(z_{it}' \cdot \gamma\right) \tag{5}$$

If a person is observed, the number of errors y_{it} is given by

$$y_{it} = x_{it}' \cdot \beta + u_{it} \tag{6}$$

For given x_{it}, individuals with high u_{it} make more errors. Now, we need an assumption regarding the relation between u_{it} and the ε_{it}. We assume that

$$E\left(u_{it}\left|\varepsilon_{it}\right.\right)=\begin{cases}\xi^{y}\cdot\varepsilon_{it} & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot\varepsilon_{it} & \Leftrightarrow \quad i \quad \text{old}\end{cases} \tag{7}$$

Now, what about the conditional means of u_{it} with respect to the ε_{it}?

$$E\left(u_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\ \ \forall i\right.\right)=\begin{cases}\xi^{y}\cdot E\left(\varepsilon_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma^{y}\right.\right) & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot E\left(\varepsilon_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma^{y}\right.\right) & \Leftrightarrow \quad i \quad \text{old}\end{cases}=\begin{cases}\xi^{y}\cdot\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)} & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)} & \Leftrightarrow \quad i \quad \text{old}\end{cases} \tag{8}$$

What's the expectation of y_{it} given x_{it} and z_{it} such that we observe the worker?

$$E\left(y_{it}\mid x_{it},s_{it}=1\ \ \forall i\right)=E\left(y_{it}\left|x_{it},\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\right.\right)=E\left(x_{it}^{'}\cdot\beta+u_{it}\left|x_{it},\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\right.\right) \tag{9}$$

$$E\left(y_{it}\mid x_{it},s_{it}=1\right)=x_{it}^{'}\cdot\beta+E\left(u_{it}\left|x_{it},\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\right.\right) \tag{10}$$

The expected value of y_{it} given that worker i is observed is:

$$E\left(y_{it}\mid x_{it},s_{it}=1\ \ \forall i\right)=x_{it}^{'}\cdot\beta+\begin{cases}\xi^{y}\cdot\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)} & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)} & \Leftrightarrow \quad i \quad \text{old}\end{cases} \tag{11}$$

Equation (11) is estimated where the inverse Mills ratios $\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}$ and $\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}$ are predictions from estimating equation (5) using a probit specification. Results from estimating equation (11) are reported in the central column of Table 2. Results from estimating equation (5) are in Table C.1.

C.3.2 Errors on the team level

So far, we considered the case where errors and selection are both observed at the individual level. In our data, however, the errors are observed at the team level. This makes correction of the selection bias a bit more complicated. If the team j is observed, the number of errors y_{it} is given by

$$y_{jt}=x_{jt}^{'}\cdot\beta+u_{jt} \tag{12}$$

where x_{it} are team characteristics like average age or share of women. For given x_{it}, teams with high u_{it} make more errors. Selection of workers into the sample is given by (equation: selection young) and (equation: selection young). Now, we need an assumption regarding the relation between u_{it} and the $\left\{\varepsilon_{it}\right\}_{i=1}^{N_j}$. We assume that

$$E\left(u_{jt}\left|\varepsilon_{it}\right.\right)=\begin{cases}\xi^{y}\cdot\varepsilon_{it} & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot\varepsilon_{it} & \Leftrightarrow \quad i \quad \text{old}\end{cases} \tag{13}$$

and

$$E\left(u_{jt}\left|\left\{\varepsilon_{it}\right\}_{i=1}^{N_j}\right.\right)=\xi^{y}\cdot\sum_{i\ \text{young}}\varepsilon_{it}+\xi^{o}\cdot\sum_{i\ \text{old}}\varepsilon_{it} \tag{14}$$

This implies that within the young and within the old, each individual ε_{it} of any worker i has the same effect on the teams performance. The individual ε_{it} are i.i.d. The individual effects just add up.

Now, what about the conditional means of u_{it} with respect to the ε_{it}'s?

$$\begin{aligned}E\left(u_{jt}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\ \ \forall i\right.\right)&=\xi^{y}\cdot\sum_{i\ \text{young}}E\left(\varepsilon_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma^{y}\ \ \forall i\right.\right)+\xi^{o}\cdot\sum_{i\ \text{old}}E\left(\varepsilon_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma^{y}\ \ \forall i\right.\right)\\&=\xi^{y}\cdot\sum_{i\ \text{young}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}+\xi^{o}\cdot\sum_{i\ \text{old}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}\end{aligned} \tag{15}$$

What's the expectation of y_{it} given x_{it} and $\left\{z_{it}\right\}_{i=1}^{N}$ such that we observe the team?

$$E\left(y_{jt}\mid x_{jt},s_{it}=1\ \ \forall i\right)=E\left(y_{jt}\left|x_{jt},\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\ \ \forall i\right.\right)=E\left(x_{jt}^{'}\cdot\beta+u_{jt}\left|x_{it},\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\ \ \forall i\right.\right) \tag{16}$$

$$E\left(y_{jt}\mid x_{jt},s_{it}=1\ \ \forall i\right)=x_{jt}^{'}\cdot\beta+E\left(u_{jt}\left|x_{jt},\varepsilon_{it}>-z_{it}^{'}\cdot\gamma\ \ \forall i\right.\right) \tag{17}$$

The expected value of y_{it} given that team j is observed is:

$$E\left(y_{jt}\mid x_{jt},s_{it}=1\ \ \forall i\right)=x_{jt}^{'}\cdot\beta+\xi^{y}\cdot\sum_{i\ \text{young}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}+\xi^{o}\cdot\sum_{i\ \text{old}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)} \tag{18}$$

Equation (18) is estimated where the inverse Mills ratios $\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}$ and $\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}$ are predictions from estimating equation (5) using a probit specification. Results from estimating equation (18) are reported in the second and third column of Table 1. Results from estimating equation (5) are in Table C.1.

C.3.3 Estimating the selection equation

Table C.1 reports results from estimating the selection equation (5) using a probit specification: $P\left(z_{it}^{'}\cdot\gamma+\varepsilon_{it}>0\right)=\Phi\left(z_{it}^{'}\cdot\gamma\right)$. The left hand column reports the results from the selection equation for the younger workers (<40 years) while the right hand column contains the results for the older workers ($\geq$ 40 years). An important variable that affects the probability of being in the sample but not the number and severity of errors is the individual sickness rate. For every worker, we calculate the average absence rate due to sickness and include it in the selection equation but not in the error regressions.

Table C.1: Regression results: sample selection

dependent variable: dummy for being in the sample				
	workers younger than 40 years		workers older than 40 years	
age dummies				
age < 20	0.143	(0.239)		
20 < age < 22	0.130	(0.285)		
22 < age < 24	0.267	(0.028)		
24 < age < 26	0.321	(0.008)		
26 < age < 28	0.503	(0.000)		
28 < age < 30	0.686	(0.000)		
30 < age < 32	0.823	(0.000)		
32 < age < 34	0.924	(0.000)		
34 < age < 36	1.13	(0.000)		
36 < age < 38	1.16	(0.000)		
38 < age < 40	1.10	(0.000)		
40 < age < 42			6.58	(0.000)
42 < age < 44			6.54	(0.000)
44 < age < 46			6.65	(0.000)
46 < age < 48			6.56	(0.000)
48 < age < 50			6.50	(0.000)
50 < age < 52			6.59	(0.000)
52 < age < 54			6.63	(0.000)
54 < age < 56			6.19	(0.000)
56 < age < 58			5.65	(0.000)
58 < age < 60			4.54	(0.000)
60 < age < 62			4.84	(0.000)
62 < age < 64			6.04	(0.000)
64 < age			4.87	(0.000)
average team age	0.0308	(0.000)	-0.0516	(0.000)
sickness rate	0.00641	(0.000)	-0.00683	(0.000)
years of schooling	0.0971	(0.000)	-0.0294	(0.000)
av. team schooling years	0.0605	(0.000)	-0.289	(0.000)
German dummy	1.03	(0.000)	0.196	(0.000)
share of Germans	-3.93	(0.000)	-2.38	(0.000)
French dummy	1.06	(0.000)	0.418	(0.000)
share of French	-5.26	(0.000)	-3.96	(0.000)
Turkish dummy	1.45	(0.000)	0.259	(0.000)
share of Turkish	-3.61	(0.000)	-5.13	(0.000)
female dummy	-0.556	(0.000)	-0.339	(0.000)
share of women	-0.821	(0.000)	-0.121	(0.012)
late shift	0.379	(0.000)	-0.395	(0.000)
team hopper dummy	-0.740	(0.000)	-1.75	(0.000)
share of team hoppers	-3.30	(0.000)	-1.56	(0.000)
team size	-0.00659	(0.000)	0.0255	(0.000)
av. team tenure	0.00438	(0.000)	0.0547	(0.000)
PFI	0.0741	(0.000)	-0.640	(0.000)
R^2	0.247		0.285	
# observations	2030939		1164115	
p-values in parentheses				

D Sample of a Production Progam

Table D.1 Production Program for 2003

truck type	66	67	68	70	71	72	74	80	81	75	76	77	79
driver's cab	LKN high roof	LKN short	LKN long	SKN short	SKN medium	SKN long	SKN LH	SKN-C long	SKN-C high roof	MPII short	MPII medium	MPII long	MPII LH
07.01. - 17.01.03	5	56	7	13	17	21	4	3	7	1	4	11	18
20.01. - 31.01.03	3	49	6	14	22	25	7	4	4	1	2	10	20
03.02. - 14.02.03	4	50	6	17	18	26	9	3	6		1	9	15
17.02. - 28.02.03	5	47	8	21	14	27	9	4	7	1	2	9	12
05.03. - 14.03.03	3	45	9	15	15	32	7	3	9	1	2	11	13
17.03. - 28.03.03	5	46	7	13	16	32	5	3	8	3	3	11	14
31.03. - 11.04.03	4	48	7	15	17	27	7	5	6		4	12	14
14.04. - 17.04.03	3	53	7	11	18	28	7	5	5	1	3	12	15
22.04. - 30.04.03	4	51	7	12	15	31	5	3	5	2	4	13	17
05.05. - 16.05.03	4	51	5	16	14	28	3	3	5	3	7	14	15
19.05. - 28.05.03	4	53	6	12	15	23	3	3	5	7	10	13	15
02.06. - 13.06.03	4	51	6	9	7	20	4	4	6	5	11	23	16
16.06. - 27.06.03	4	50	6	12	8	16	4	4	6	5	12	28	15
30.06. - 11.07.03	4	50	6	11	7	16	1	5	6	9	12	26	17
14.07. - 18.07.03	3	55	5					10	4	13	20	30	21
21.07. - 01.08.03	3	50	4					7	6	11	16	26	19
04.08. - 15.08.03	4	49	6					5	6	10	15	27	21
18.08. - 29.08.03	3	49	7					4	5	12	14	30	21
01.09. - 12.09.03	3	50	6					4	7	11	12	35	18
15.09. - 26.09.03	4	58	8					5	9	12	17	48	27
29.09. - 10.10.03	5	52	14					5	7	12	21	51	24
13.10. - 24.10.03	6	51	14					5	5	12	21	49	28
27.10. - 31.10.03	4	62	8					3	8	13	26	45	23
03.11. - 14.11.03	7	52	12					5	4	13	15	59	23
17.11. - 28.11.03	5	54	14					4	5	15	18	52	25
01.12. - 12.12.03	4	51	15					3	6	15	21	42	29
15.12. - 09.01.04	6	48	13					4	5	18	18	46	23

This table shows the numbers of trucks of the 13 different types that are produced on each day. The program changes every two weeks. The program is the same for all work teams on the assembly line.

References

Aubert, P. (2003): “Productivity, wage and demand for elder workers; an examination on French matched employer-employee data,” Working Paper INSEE, Paris.

Aubert, P., and B. Crépon (2007): “Are Older Workers Less Productive? Firm-Level Evidence on Age-Productivity and Age-Wage Profiles,” mimeo.

Bayer, T., J. Page, Y Raviv, and J. Rosett (2009): “Age, Human Capital and the Quality of Work: New Evidence from Old Masters,” Robert Day School Working Paper No. 2009-04.

Börsch-Supan, A., I. Düzgün, and M. Weiss (2005): “Altern und Produktivität: Zum Stand der Forschung,” MEA Discussion Paper, 73-05, Mannheim Research Institute for the Economics of Aging, http://www.mea.uni-mannheim.de.

Castellucci, F., M. Padula, and G. Pica (2009): “The Age-Productivity Gradient: Evidence from a Sample of F1 drivers,” CSEF Working Paper No. 226, Center for Studies in Economics and Finance.

Crépon, B., N. Deniau, and S. Pérez-Duarte (2003): “Wages, Productivity, and Workers Characteristics: A French Perspective,” CREST Document de Travail No. 2003-04, Centre de Recherche en Économie et Statistique (CREST), Paris, France.

Daveri, F. and M. Maliranta (2007): “Age, seniority and labour costs: lessons from the Finnish IT revolution,” Economic Policy 22 (49), 117 – 175.

Dostie, B. (2006): “Wages, Productivity, and Aging,” IZA Discussion Paper No. 2496, Institut zur Zukunft der Arbeit.

Fair, R. (1994): “How Fast Do Old Men Slow Down?,” The Review of Economics and Statistics, 76(1), 103 — 118.

—— (2005): “Estimated Age Effects in Baseball,” mimeo, Cowles Foundation, Yale University, http://fairmodel.econ.yale.edu/.

—— (2007): “Estimated Age Effects in Athletic Events and Chess,” Experimental Aging Research 33, 37 – 57.

Frank, R. H., and R. M. Hutchens (1993): “Wage, Seniority, and the Demand for Rising Consumption Profiles,” Journal of Economic Behavior and Organization, 21(3), 251 — 276.

Galenson, D. W. (2005): “The Greatest Artists of the Twentieth Century,” NBER Working Papers, 11899.

Galenson, D. W., and B. A. Weinberg (2000): “Age and the Quality of Work: The Case of Modern American Painters,” Journal of Political Economy, 108(4), 671 — 777.

—— (2001): “Creating Modern Art: The Changing Careers of Painters in France from Impressionism to Cubism,” The American Economic Review, 91(4), 1063 — 1071.

Göbel, C. and T. Zwick (2009): “Age and Productivity – Evidence from Linked Employer Employee Data,” ZEW Discussion Paper No. 09-020, Zentrum für Europäische Wirtschaftsforschung.

Grund, C., and N. Westergård-Nielsen (2005): “Age Structure of the Workforce and Firm Performance,” IZA Discussion Paper, 1816.

Haltiwanger, J. C., J. I. Lane, and J. R. Spletzer (1999): “Productivity Differences Across Employers: The Roles of Employer Size, Age, and Human Capital,” The American Economic Review, 89(2), 94 — 98.

—— (2007): “Wages, Productivity, and the Dynamic Interaction of Business and Workers,” Labour Economics, forthcoming.

Hellerstein, J. K., and D. Neumark (1995): “Are Earnings Profiles Steeper Than Productivity Profiles? Evidence from Israeli Firm-Level Data,” Journal of Human Resources, 30(1), 89 — 112.

—— (2004): “Production Function and Wage Equation Estimation with Heterogeneous Labor: Evidence From a New Matched Employer-Employee Data Set,” NBER Working Paper, 10325.

Hellerstein, J. K., D. Neumark, and K. R. Troske (1999): “Wages, Productivity, and Worker Characteristics: Evidence from Plant Level Production Function and Wage Equations,” Journal of Labor Economics, 17(3), 409 — 446.

Hunter, J. E., and R. F. Hunter (1984): “Validity and Utility of Alternative Predictors of Job Performance,” Psychological Bulletin, 96, 72 — 98.

Ilmakunnas, P., and M. Maliranta (2005): “Technology, Labour Characteristics and Wage-Productivity Gaps,” Oxford Bulletin of Economics and Statistics, 67(5), 623 — 645.

Ilmakunnas, P., and M. Maliranta (2007): “Aging, Labor Turnover, and Firm Performance,” ETLA Discussion Paper No. 1092 http://www.etla.fi/.

Jones, B. F. (2005): “Age and Great Invention,” NBER Working Paper, 11359.

Kotlikoff, L., and D. Wise (1989): “Employee Retirement and a Firm’s Pension Plan,” in The Economics of Aging, ed. by D. Wise, pp. 279 — 334. University of Chicago Press, Chicago.

Kotlikoff, L. J., and J. Gokhale (1992): “Estimating a Firm’s Age-Productivity Profile Using the Present Value of Workers’ Earnings,” Quarterly Journal of Economics, 107(4), 1215 — 1242.

Laitner, J., and D. Stolyarov (2005): “Technological Progress and Worker Productivity at Different Ages,” mimeo, University of Michigan.

Lallemand, T., and F. Rycx (2009): “Are Young and Old Workers Harmful for Firm Productivity?” IZA Discussion Paper 3938, Institut zur Zukunft der Arbeit.

Lazear, E. (1979): “Why is There Mandatory Retirement?,” Journal of Political Economy, 87(6), 1261 — 1284.

—— (1981): “Agency, Earnings Profiles, Productivity, and Hours Restrictions,” American Economic Review, 71(4), 606 — 620.

Loewenstein, G., and N. Sicherman (1991): “Do Workers Prefer Increasing Wage Profiles?,” Journal of Labor Economics, 9(1), 67 — 84.

Malmberg, B., T. Lindh, and M. Halvarsson (2005): “Productivity Consequences at the Plant Level of Work-Force Ageing: Stagnation or a Horndal Effect?,” mimeo, Institute for Future Studies, Stockholm, Sweden.

McEvoy, G. M., and W. F. Cascio (1989): “Cumulative Evidence of the Relationship Between Employee Age and Job Performance,” Journal of Applied Psychology, 74(1), 11 — 17.

Medoff, J. L., and K. G. Abraham (1980): “Experience, Performance, and Earnings,” Quarterly Journal of Economics, 95(4), 703 — 736.

Prskawetz, A., T. Lindh, B. Mahlberg, V. Skirbekk, I. Freund, M. W. Dworak, T. Lindh, B. Malmberg, A.-C. Jans, O. S. Nordström, and F. Andersson (2006): "The Impact of Population Ageing on Innovation and Productivity Growth in Europe," Research Report 28, Vienna Institute of Demography, Austrian Academy of Sciences, Wien, Austria.

Salthouse, T. A., and T. J. Maurer (1996): "Aging, Job Performance, and Career Development," in Handbook of the Psychology of Aging, 4th Edition, ed. by J. E. Birren, and K. W. Schaie, pp. 353 — 364. Kluwer Academic Publishers, New York.

Schneider, L. (2007): "Mit 55 zum alten Eisen? Eine Analyse des Alterseinflusses auf die Produktivität anhand des LIAB", Zeitschrift für Arbeitsmarktforschung, 40 (1),77- 97.

Schneider, H., and D. Stein (2006): "Personalpolitische Strategien Deutscher Unternehmen zur Bewältigung demografisch bedingter Rekrutierungsengpässe bei Führungskräften," Research Report 6, IZA, Bonn, Germany.

Skirbekk, V. (2004): "Age and Individual Productivity: A Literature Survey," in Vienna Yearbook of Population Research, ed. by G. Feichtinger, Vienna. Verlag der Österreichischen Akademie der Wissenschaften, Vienna.

van Ours, J. C. (2009): "Will you still need me – when I'm 64?," IZA Discussion Paper No. 4246, Institut zur Zukunft der Arbeit.

Weinberg, B. A., and D. W. Galenson (2005): "Creative Careers: The Life Cycles of Nobel Laureates in Economics," NBER Working Paper, 11799.