Maastricht University

Axel Börsch-Supan,
Matthias Weiss

Productivity and age: Evidence from work teams at the assembly line

RM/13/029

GSBE

Maastricht University School of Business and Economics
Graduate School of Business and Economics

P.O Box 616
NL- 6200 MD Maastricht
The Netherlands

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

AXEL BÖRSCH-SUPAN† AND MATTHIAS WEISS‡

Abstract

We study the relation between workers' age and their productivity in work teams, based on a new and 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 related to the errors. We correct for non-random sample selection and the potential endogeneity of the age-composition in work teams. Our results suggest that productivity in this plant which is typical for large-scale manufacturing does not decline at least up to age 60.

JEL codes: J24, J14, D24

Keywords: Age-Productivity Profiles

* We are grateful to Patrick Aubert, Paul Baltes, Lisa Berkman, Dana Goldman, Silke Januszewski, Arthur Kennickell, Ed Lazear, Melanie Lührmann, Jürgen Maurer, Carsten Ochsen, Jay Olshansky, Andrew Oswald, Steffen Reinhold, Jack Rowe, Joachim Winter and David Wise for valuable comments and suggestions on earlier versions of this paper. 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. Financial support from the Hans-Böckler-Stiftung for this project is as gratefully acknowledged as the funding of MEA through the State of Baden-Württemberg and the German Insurance Association. We are especially grateful to the Mercedes-Benz 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.

† Munich Center for the Economics of Aging (MEA) at the Max Planck Institute for Social Law and Social Policy, Munich, Germany and NBER, Cambridge, MA, USA. E-Mail: axel@boersch-supan.de.

‡ Munich Center for the Economics of Aging (MEA) at the Max Planck Institute for Social Law and Social Policy, Munich, Germany and ROA, Maastricht, NL. E-Mail: m.weiss@mea.mpisoc.mpg.de.

1 Introduction

Productivity is at the core of economics. Growth in modern economies is dominated by productivity growth (Maddison, 2001). Employers are keen to increase productivity of their workers as this is an essential element in competition across firms. Population aging now gives productivity a new dimension: older workers are often thought to be less productive. This impression is widespread and implicit in many discussions about aging, even in our economic textbooks.[1] If the impression were true, population aging would have negative effects on overall productivity as the share of older workers is increasing, and would thus directly reduce economic growth. This would amplify the economic strains on aging societies already exerted by increasing Social Security and Medicare bills. The impression of declining productivity by age has also implications on the micro level, e.g. on personnel policies by employers and on retirement choices made by employees. In many countries, the assertion that productivity declines with age is used as a motivation for early retirement policies. This paper provides new and maybe surprising evidence on the age-productivity relation, shedding doubt on the wide-spread assertion at least for the age range which is commonly referred to as the "early retirement window".

Labor economists have a long-standing tradition to study the productivity of workers. Recent papers have focused on the relation between incentives, cooperation, peer effects, supervision, and productivity in work teams (Bandiera et. al. (2005, 2007, 2009, 2010), Lazear (2000), Lazear et. al. (2011), Mas and Moretti (2007)). Also estimating age-productivity profiles has been on the agenda of labor economists for a long time, see the surveys by Skirbekk (2004), Gelderblom (2006), and *Labour Economics'* recent "Special Issue: Ageing and Productivity" (Vol. 22, June 2013)[2]. Such studies, however, have encountered three fundamental challenges: measurement, selectivity/endogeneity, and aggregation. These methodological challenges have made it hard to distinguish fact from fiction. This paper and the underlying large data collection effort are tailored to overcome these difficulties.

First, productivity is hard to measure directly. While it is well documented by occupational medicine, cognitive psychology, and gerontology that muscle strength, sight, lung, kidney, and heart functioning, and many other biometric indicators deteriorate from early age onwards, experience and the ability to deal

[1] E.g., Lazear (1995, p. 40, figure 4.1)

[2] Bloom and Sousa-Poza (2013), Göbel and Zwick (2013), Lovàsz and Rigó (2013), Mahlberg et. al. (2013), Romeu Gordo and Skirbekk (2013), and Vandenberghe (2013).

with human nature appear to increase with age. Since the latter characteristics are hard to measure, there is a bias towards direct measures that decline early in life. This may have contributed to the above-mentioned impression.

A second challenge is the potential endogeneity of the age composition through various selection processes. Being in the labor force is endogenous since employers are more likely to hold on to productive than unproductive workers. Hence plant closures and early retirement tend to create a positive selection of productive workers. A related endogeneity problem exists for the age-structure on the company level. Since more productive firms are usually more profitable, they expand and increase their workforce. This leads to a rejuvenation of their workforce because new hires are more likely to be young. Relating productivity to the age of the workforce in this case results in a spurious negative correlation between productivity and age.

Finding the right level of aggregation is the third challenge. An individualistic view fails to take into account that workers often work in teams and thereby affect one another's productivity. Older workers may devote some of their working time to helping or teaching younger workers. In this case, an individualistic approach will underestimate older workers' and overestimate younger workers' productivity.[3] Related aspects are workers' contributions to their team's work climate and how teams deal with emergency situations. A plant or company view, on the other hand, obscures job heterogeneity and its interaction with motivation and thus productivity. One would expect, e.g., that the productivity effect of older workers on the shop floor whose careers have peaked is quite different from the productivity effect of equally old managers who still might have ambitions for a position at the company's top or a realistic chance to move to another company. Plant view regressions that average over different non-linear age-productivity profiles might therefore create misinterpretations.

We have assembled a unique combination of company data that permit us to overcome these problems in an unprecedented way. The data have three innovative elements. First, we measure productivity in an assembly line environment in which the time to produce a unit of output is as standardized as the quality of the final product. As the assembly line has the same speed for all work teams and the design of the trucks is pre-defined, more productive work teams are not able to produce more or better output than less productive work teams. Workers, however, make errors which are detected at end control. More productive work teams differ from less productive work teams only in the errors they make. We therefore

[3] In principle, it can also be the other way round: Young workers helping the old. In either case, the individualistic approach fails.

use the number and severity of production errors during the assembly process as a precise and well-observed measure of productivity. We exploit the daily variation in the team composition of work teams over four years to identify the age-productivity profiles both on the work team and the individual workers' level.

Second, we have merged the production error data with longitudinal personnel data. This permits us to hold a broad range of workers' characteristics constant and, most importantly, to correct for the selection effects marring so many earlier studies due to the endogeneity of early retirement and redundancy decisions by employees and employers to productivity.

Third, we measure the joint productivity of workers in a work team. This takes into account the individual worker's contribution to his co-workers' productivity. Particularly the contribution of older workers' may be underestimated if productivity is measured at an individual level. Examples for such potential contributions to a team's productivity are the instruction of younger workers,[4] being relaxed in tense or hectic situations, and contributing positively to the work climate. We think that our approach solves the major aggregation problems in earlier studies.

Earlier studies can be broadly divided into four groups. There are many studies relating plant level productivity to the age of the plants' employees.[5] Plant level productivity can be measured easily and reliably but the level of aggregation is quite high when the goal is to study the relation between productivity and age. Furthermore, the age structure of companies is probably not exogenous as pointed out before.

A second group of studies uses individual's wages as a productivity measure.[6] Wages, however, often increase with age and/or seniority independently of productivity, and wage decreases are extremely rare.[7]

---

[4] If an older worker helps a younger worker, the older worker's productivity, narrowly defined by individual, is zero as the older worker is not producing anything at that time. The contribution to the work team's productivity, however, is positive.

[5] E.g., Hellerstein and Neumark (2004), Hellerstein et al. (1999), Haltiwanger, Lane, and Spletzer (1999, 2007) for the U.S., Hægeland and Klette (1999) for Norway, Aubert (2003), Crépon, Deniau, and Pérez-Duarte (2003), Aubert and Crépon (2007) for France, Hellerstein and Neumark (1995) for Israel, Grund and Westergård-Nielsen (2008) for Denmark, Ilmakunnas and Maliranta (2005, 2007), Daveri and Maliranta (2007) for Finland, Malmberg et al. (2008) for Sweden, Dostie (2011) for Canada, Prskawetz et al. (2006) for Austria and Sweden, Lallemand and Ryckx (2009) for Belgium, van Ours (2009) for the Netherlands, Schneider (2007), Göbel and Zwick (2009) for Germany.

[6] E.g., Kotlikoff and Wise (1989), Kotlikoff and Gokhale (1992), and Laitner and Stolyarov (2005).

[7] 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

Third, a group of studies relies on managers' subjective evaluations of their employees' performance.[8] These supervisors' assessments are problematic because they may reflect prejudices about age-productivity profiles.

There are finally many studies which employ direct measures of individual productivity like, e.g., the number and quality of publications or Nobel prizes in academic research,[9] the number and quality of completed court cases,[10] the value of artists' paintings,[11] or performance in sports and chess.[12] These studies are able to measure productivity quite precisely but the range of occupations, where this approach is feasible, is small. Moreover, these studies usually refer to top performance. In everyday work life, however, the workflow is customized to average rather than top performance.

Our study is most closely related to this fourth group but relates to average performance. We have compiled our data from a truck assembly plant owned by a large German car manufacturer with plants in Asia, Europe and the U.S. At this plant, trucks are assembled by work teams on an assembly line. We have selected this plant because it features a taylorized production process typical for the manufacturing industry, and because it stacks our cards against finding flat or increasing productivity with age. Compared to many service-sector jobs, productivity in this plant requires more physical strength, dexterity, agility etc. (which tend to decline with age) than experience and knowledge of the human nature (which tend to increase with age).

Seen in this light, our results are striking. Due to the very large number of observations and our identification strategy, we are able to estimate rather precise age-productivity profiles at the individual level and at the level of a work team. These profiles do not show a decline in the relevant age range between 25 and 65 years of age. On the individual workers' level, our average productivity measure actually increases monotonically up to age 65.

We conclude that even in a work environment requiring substantial physical strength, its decline with age is compensated by characteristics that appear to increase with age and are hard to measure directly, such as experience and the

---

increasing wage profiles and explain this with loss aversion and problems of self-control.

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

[9] Jones (2010), Jones and Weinberg (2011), Weinberg and Galenson (2005), van Ours (2009).

[10] Backes-Gellner, et al. (2011).

[11] Galenson and Weinberg (2000) and (2001), Galenson (2009) and Bayer et al. (2009).

[12] Fair (1994), (2005), and (2007), van Ours (2009), and Castellucci et al. (2011).

ability to operate well in a team when tense situations occur, typically when things go wrong and there is little time to fix them.

Our study is also loosely related to a series of papers that study the relation between incentives, cooperation, peer effects, supervision and productivity in work teams (Bandiera et. al. (2005, 2007, 2009, 2010), Lazear (2000), Lazear et. al. (2011), Mas and Moretti (2007)). These papers study work teams in environments where variations in payment schemes, social networks, and supervisors can be used to identify effects on productivity.

In the assembly plant of our study, wages are paid according to a collective agreement between employers and a labor union. Workers are classified into wage groups according to the tasks they perform. Workers cannot affect their wage income through higher or lower effort (i) for contractual reasons and (ii) because their employer is not allowed and not able to monitor workers' productivity.[13] This is rooted in the German labor law and happens to be of great advantage for our study as we can be confident that the age productivity profiles we estimate are not confounded by incentive effects.

The paper is structured as follows. In the next section, we present our data set, our productivity measure, and descriptive statistics. Section 3 is devoted to our identification strategy, in particular our solution to the endogeneity and selection problems and the separation of age from cohort and time effects. Section 4 presents our results at the individual and the work team level. Section 5 provides conclusions. An extensive appendix with methodological details is available on request.

2 The data

The truck assembly plant is located in the South of Germany and owned by Mercedes-Benz. Similar plants are located elsewhere in Europe, in the U.S. ("Freightliner"), and Asia. It has two stages: the assembly line and a quality control unit.

The assembly line is divided into 50 work stations that are located one after another. The number of work teams is double the number of work stations be-

[13] When the recording of errors was introduced in the plant, the works council vetoed the linkage of the error data to the personnel data. We are grateful to the works council that they allowed us to establish an anonymized version of such a link for our study. While we cannot attribute a specific error to a specific person, we can attribute the number and weight of errors to a set of personnel characteristics such as age, sex, etc.

cause on every day, there is an early and a late shift.[14] Work on the assembly line is completely standardized: at each work station, the workers assemble a certain element of a standardized type of truck in a fixed time that is given by the speed of the conveyor belt. Hence, at first glance, labor productivity at these 50 work stations (defined as the ratio of output units to units of labor input) appears to be independent of the work teams' actions since quantity and quality of output is given as well as the time units of labor input. At second glance, however, production quality differs across work teams as they can make errors. Variation in productivity thus becomes manifest *only* in the variation of production errors. This is a central feature of our approach.

At the end of the assembly line, quality inspectors check the assembled trucks, record the errors and assign them to the work station where they occurred. We observe 8564 errors in 100 teams on 973 days. The error rate is thus 8.8 percent. Every error is given a weight that specifies the severity of the error. This severity weight depends on the costs to make up for the error. The severity weights range from 5 to 95 with mean 11 (median 10) and standard deviation 5.7.

Our inverse productivity measure is then defined as the sum of errors per work team per day where the errors are rated with their respective weights. E.g., if a team makes two errors on a day with weights 5 and 30, our inverse productivity measure for this team for this day takes the value 35. Errors are relatively rare and most errors are small.

The daily record of errors is our first core data set. Errors by work team are matched with personnel data, our second core data set.[15] We observe 3824 workers in 100 work teams at 50 workplaces on 973 work days during the years 2003 through 2006. The personnel data are very rich and longitudinal. They 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. They also record exits from jobs due to early retirement, transfer to a different unit and promotion to group leader or a managerial position.

The age composition in the plant is fairly representative for the production line of a manufacturing plant. 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.

[14] To be precise: At every workplace there is a „team A" and a „team B". A-teams work early in even work weeks and late in uneven weeks. B-teams work early in uneven weeks and late in even weeks.

[15] Appendix A details how errors are matched with work teams and personnel data. Appendix B provides descriptive statistics of all variables used in the paper. The distribution of the key variables is described in an extended version of this paper available upon request.

About 5% of the workers are older than 55 years. While this percentage is low, the absolute number is large: workers older than 55 years represent some 89,000 worker-day-observations (out of a total of 1,767,030 observations). This large number permits a relatively precise measurement of age-productivity profiles up to the statutory retirement age of 65.

Another key feature of our data is that errors are assigned to work teams. Hence, our primary unit of observation is the work team. The individual worker responsible for the error is not identified. This is a legal requirement. We think that this is an advantage rather than a disadvantage since it provides the adequate solution to the aggregation problem in productivity studies. Workers operating in a team are supposed to help each other, and, consequently, also make errors in teams. If worker A "makes" an error, it might not be his fault because worker B did the preliminary work improperly or worker C assisted inadequately. Work teams are thus the appropriate aggregation unit in order to study age-related productivity differences among workers at the assembly line. Due to personnel fluctuations across work teams, we are also able to identify the contribution of each individual worker to a work team's productivity. In our results section, we will show age-productivity profiles for both entire work teams and individual workers.

In addition to the two core data sets (errors and personnel) there are several supplemental data sets such as the daily production plan which gives us information on the work load and information linking personnel ID, work station, shift and date which have allowed us to merge the data sets.

3 Identification

Aim of this paper is to relate labor productivity to age. The identification of this relation poses several methodological challenges:

- Which variation across workers and teams should be exploited to identify the age-productivity relation?
- How can we make sure that the observable variation in age is exogenous?
- How can we deal with sample selection bias in the presence of early retirement and career moves?
- How can we distinguish age from cohort effects?

They are dealt with in this section.

3.1 Identifying variation

To a first degree, errors occurring at different work stations on the assembly line are strictly comparable because every error is given a severity weight that accounts for the costs of fixing that error. The severity weights thus translate every error into a common metric independent of the work stations' specific tasks.

Comparisons across teams at different work stations may nevertheless be a source of biases as work stations may be different with respect to the susceptibility to errors. Since we have nearly 1,000 observations per work team (973 work days spread over four years), we do not depend on comparisons across work teams. We prefer being on the safe side and take such potential differences into account by using work-team fixed effects.[16] We therefore identify the relation between workers' age and the errors they make by only using that variation in errors and age over time which occurs within work teams where tasks are homogeneous, and avoid comparisons across potentially different work stations. In a similar fashion, we insert a fixed effect for each worker in our individual-level regressions. The age effect is thus identified by the variation in weighted errors over the four-year observation period.

3.2 Exogeneity of variations in age

As pointed out in the introduction, the age composition of a plant tends to be endogenous to labor productivity since, e.g., fast growing start-ups have freshly hired and thus typically younger staff than established companies. This problem is well known for studies of age-productivity profiles based on comparisons across plants or entire companies.

Studies on individuals or work teams within plants may suffer from a similar problem if age affects the assignment of workers to tasks, e.g., if older workers are systematically assigned to work stations which have easier tasks. Hence, if productivity is measured by comparing workers at different work stations, the productivity of older workers may be overestimated. We avoid this problem by not using the variation *across work stations* in our estimation because we include work team fixed effects as described in the previous subsection.

Correct identification in our case rests on the assumption that the variation in

[16] Work-team fixed effects are not equal to work-station fixed effects because each work station has an early and a late shift. Because we want to avoid any potential endogeneity of the work team composition with respect to early vs. late shift, we choose fixed effects for work teams rather than a fixed effect for each work-station.

the age composition *within work teams over time* is exogenous and not subject to optimizing management decisions. The fluctuation within teams has two components: The team composition changes from day to day as some workers call in sick (6%) and others are on vacation (12%) or have compensatory time off for extra hours workers (6%).[17] This day-to-day fluctuation within the core team is random and not the result of any management decision.[18] 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 with workers from other work teams. At this stage, there is in principle room for optimization, potentially creating an endogeneity problem. In practice, however, replacement decisions are made on very short notice with a rather limited pool of team hoppers. Optimization would require knowledge about the optimal age composition for the day's production program and the respective work team. Furthermore, a team hopper with the optimal characteristics has to be available. The managers in the plant convinced us that neither requirement is met in daily practice and that optimizing behavior would not pay off to justify their efforts.

It is important to stress that the line managers do optimize the composition of the long-run team composition at every work station for a normal day. What is not optimized, however, are the short-run deviations from that composition. It is precisely this exogenous variation that we exploit in our estimation.

3.3 Sample selectivity

Workers older than 55 years are underrepresented in the work force. The same holds for our sample. 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. We correct for this selectivity bias in two ways: first by employing a Heckman-style selectivity-correction model and second by adding worker fixed effects in addition to the work team fixed effects described earlier. Worker fixed effects can be separately identified from work-team fixed effects due to personnel fluctuations across work teams described earlier.

Since we have four years of personnel data, we have information on those

[17] This means that in a team of 16 workers, on an average day, one worker is absent due to compensatory time off, another one due to sickness and two more workers are on vacation. In order to buffer these fluctuations, each work team has about 25% more members than are needed on a regular day.

[18] Workers who call in sick frequently may be less productive. However, this potential endogeneity problem is taken care of by our workers fixed effect.

workers who enter or exit the sample within the four-year observation period. This enables us to estimate a Heckman-style selection correction model. We use individual sickness rates as an instrument as they affect selection but have no effect on the number and severity of errors. Identification is further improved by the fact that selection is observed on the individual level while errors are observed on the work team level. Since our observation unit in the regression analyses is a work team while selection into the sample is an individual phenomenon, we aggregate individual Mills ratios to team Mills ratios (see Appendix C for details).

There is a second selection process specific to our data. Our sample consists of workers on the assembly line only. Even the foremen are not included. Workers who leave our sample before the age of 65 may also be promoted to jobs off the assembly line, most likely because they performed better than those who remain. This creates a selectivity bias in the opposite direction of the above-mentioned one generated by early retirement of less productive workers.

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. Both first stage regressions fit well with R^2s of 0.25 and 0.28, see Appendix C.

In addition to these Heckman-style corrections, we included worker fixed effects in order to correct for non-random sample selection. This is possible for the regressions where worker-days are the observation unit because workers move across work teams over time. This correction is particularly effective, see Section 4.1, as it does not depend on instruments or functional assumptions.

3.4 Age vs. cohort effects

The well-known identification problem of age, cohort and time effects appears also in our data. Since we have panel data such that each cohort is observed at different ages, a distinction between age and cohort effects at the individual level is possible in principle if time effects are absent. Such time effects could be generated by changes in technology or organization that affect our productivity measure during the sample period.

Indeed, there was a major technological and organizational change before 2003. Thereafter, no further changes happened until 2006. The absence of changes in technology or organization during these four years was the main reason for our choice of this time period. Given the absence of such time effects and the inclusion of cohort effects in the worker-specific constants of our fixed-effects re-

gressions, age effects are identified at the individual worker's level.

While the relatively short observation period of 4 years guarantees the absence of technological change and thus time effects, the statistical basis for separating age and cohort effects is weak. However, as we will see, it suffices to create significant differential effects.

There is one more degree of freedom at the team level because average age of the team, average cohort of the team, and time are not perfectly correlated as they are at the individual worker's level. We therefore include a time trend in the team level regressions to control for possible time effects and interact this trend with age, see Section 4.2.

4 Results

We present our regression results in two subsections. In Section 4.1, we exploit the fluctuation across work teams documented in Section 2 and display our age productivity profiles based on individual worker-days. Section 4.2 is dedicated to the separation of job tenure from residual age effects and to interaction effects. The regressions presented in this section are based on work team-days. Table 1 illustrates our two-level approach:

Table 1: Individual vs. work-team regressions

	Section 4.1	Section 4.2
Observation unit	Worker day	Team day
Object	Age productivity profile	Interaction effects Effect of job tenure
Work team fixed effects included	Yes	yes
Worker fixed effects included	Yes	no
Interactions with age included	No	yes
Job tenure included	No	yes
Number of control variables	small	large

Regressions on the individual level are attractive because they permit a distinction between cohort effects and age effects, because dealing with sample selection is easier and more powerful, and because we can look at higher ages. Worker fixed effects remove differences between workers that are constant over time, including cohort effects. Since time effects are absent during our observation period,

as explained in Section 3.4, we identify 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.

Regressions on the individual level also allow us to identify effects at higher ages. Age-productivity profiles on the individual level range from 18 years to 65 years while average work-team age ranges between 25 and 50 years.

The main disadvantage of regressions at the individual level is collinearity: we cannot separate age from job tenure effects because on the individual level, after controlling for fixed effects, age and job tenure are perfectly collinear.[19]

Therefore, we also ran regressions at the team level where average age and average job tenure are only imperfectly correlated even within work teams as the team composition changes from day to day. This allows us to identify the effect of job tenure given age. 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. Thus we can study in great detail and functional flexibility which variables have an effect on the age-productivity profile.

4.1 Regressions on the individual level

We observe 3824 workers in 100 work teams on 973 days. While productivity is measured at the team level, the fluctuation of workers across work teams permits us to attribute errors (and thus productivity) to individuals. The regressions on the individual level are based on more than 1.5 million worker-days. Our regressions include age and a large set of control variables. We report on two specifications, baseline and selectivity controlled.

In order to allow for non-linear age effects, we use a piecewise linear specification of average team age (5-year linear splines). Regression results are presented in Table 2, columns 1 and 2. Figures 1 and 2 display the resulting age profile of our inverse productivity measure (sum of severity-weighted errors per day).

The results are striking. The average sum of severity-weighted errors declines until age 65. It is measured rather precisely up to age 60. The decline becomes insignificant at ages between 60 and 65 years where we have too few observations for precise estimation. This holds for both specifications, with and without selectivity correction, which are virtually identical, indicating that the worker

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

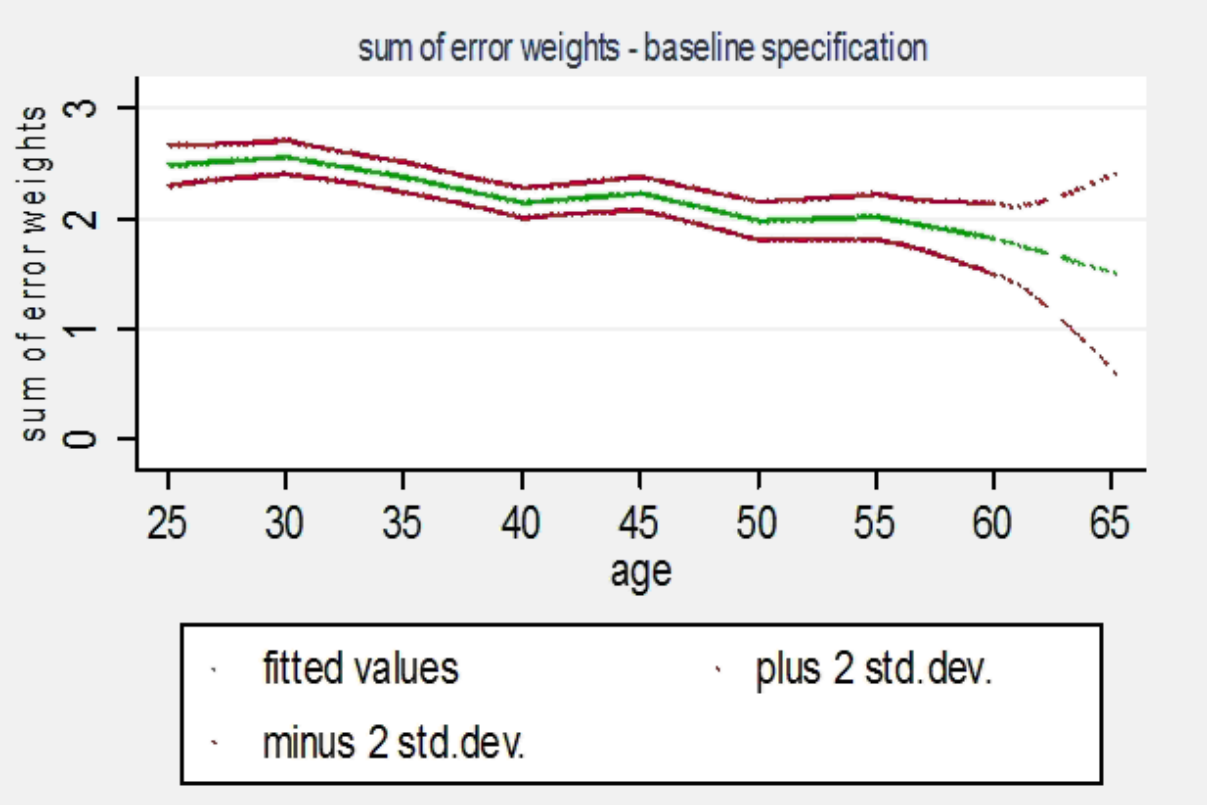

Figure 1: Age productivity profiles (in terms of errors)

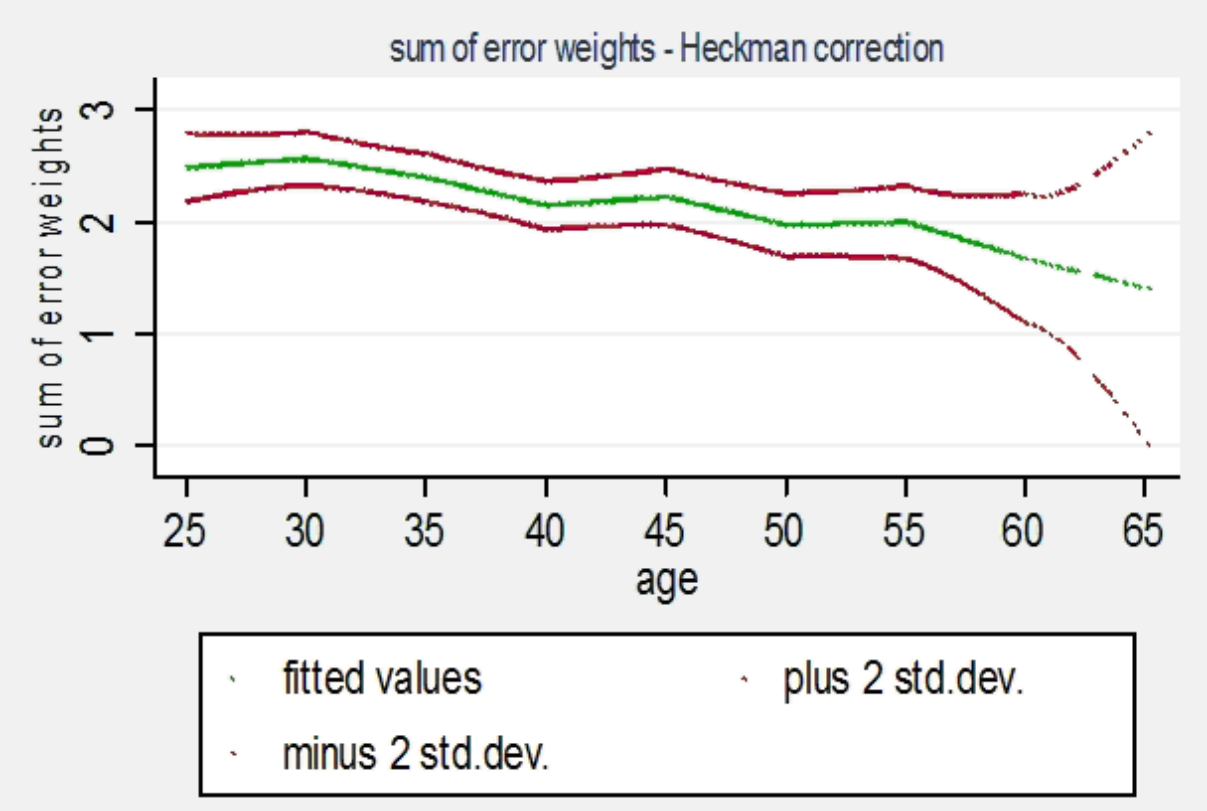

Figure 2: Age productivity profiles (in terms of errors), Heckman correction

fixed effects have already removed the selection bias. Our main result is no evidence for a productivity decline in this assembly plant at least until age 60.

In Figures 3 and 4 (columns 3 and 4 in Table 2), 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 (Figure 2), we find a clearly increasing profile: Older workers make significantly more errors. On the other hand, the severity of errors is strongly decreasing with age (Figure 3).

Table 2: Regression Results (individual level)

Dependent variable:	sum of error weights				number of errors		error severity (given an error occurred)	
	basic specification		correcting for sample selection		correcting for sample selection		correcting for sample selection	
Age splines								
17 – 25 years	0.155	(0.0174)	0.157	(0.0176)	0.0143	(0.00152)	-0.288	(0.0793)
25 – 30 years	0.0143	(0.0132)	0.0162	(0.0139)	0.00482	(0.00117)	-0.339	(0.0857)
30 – 35 years	-0.0361	(0.0111)	-0.0344	(0.0116)	-0.000285	(0.000955)	-0.328	(0.0805)
35 – 40 years	-0.0464	(0.0104)	-0.0493	(0.0107)	0.000130	(0.000884)	-0.385	(0.0821)
40 – 45 years	0.0171	(0.00996)	0.0152	(0.0104)	0.00555	(0.000869)	-0.389	(0.0736)
45 – 50 years	-0.0499	(0.00962)	-0.0503	(0.00964)	0.000555	(0.000788)	-0.433	(0.0814)
50 – 55 years	0.00730	(0.00996)	0.00515	(0.0100)	0.00424	(0.000799)	-0.428	(0.0911)
55 – 60 years	-0.0391	(0.0254)	-0.0653	(0.0314)	-0.00253	(0.00247)	-0.412	(0.272)
60 – 65 years	-0.0611	(0.0822)	-0.0542	(0.0826)	-0.000638	(0.00618)	-0.214	(0.986)
control variables								
workload	1.12	(0.00644)	1.12	(0.0644)	0.0996	(0.00559)	0.0174	(0.237)
workload^2	-2.17	(0.0173)	-2.17	(0.173)	-0.184	(0.0146)	0.422	(0.465)
cycle time	0.0147	(0.00150)	0.0147	(0.00150)	0.00138	(0.000124)	0.0526	(0.0188
team size	0.0672	(0.00689)	0.0672	(0.00688)	0.00458	(0.000561)	0.0540	(0.0556)
(team size) ^{2}	-0.00125	(0.000173)	-0.00125	(0.000173)	-0.000077	(0.0000137)	0.000342	(0.00162)
external	0.0440	(0.0234)	0.0445	(0.0235)	0.00530	(0.00202)	-0.136	(0.125
Late shift	-0.100	(0.00821)	-0.100	(0.00821)	-0.00838	(0.000691)	0.104	(0.0473
days w/o change	-0.000877	(0.000275)	-0.000873	(0.000275)	-0.000097	(0.0000227)	0.00313	(0.00152
tryout Axor	-0.150	(0.0178)	-0.150	(0.0198)	-0.0145	(0.00178)	0.112	(0.117
tryout Atego	0.0317	(0.0198)	0.0328	(0.0193)	0.00594	(0.00184)	-0.174	(0.101
Monday	1.13	(0.0193)	1.13	(0.0193)	0.108	(0.00167)	-1.40	(0.630
Tuesday	1.10	(0.0191)	1.10	(0.0191)	0.103	(0.00164)	-1.33	(0.630
Wednesday	1.36	(0.0190)	1.36	(0.0190)	0.124	(0.00164)	-1.38	(0.629
Thursday	1.03	(0.0191)	1.03	(0.0191)	0.0984	(0.00164)	-1.40	(0.630
Friday	1.10	(0.0190)	1.10	(0.0190)	0.106	(0.00163)	-1.28	(0.629
inverse Mills Ratio young			0.0471	(0.124)	0.00351	(0.0102)	-0.287	(0.814)
inverse Mills Ratio old			0.170	(0.128)	0.00882	(0.0105)	0.392	(0.807)
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. Huber-White robust standard errors are in parentheses. All specifications control for individual worker fixed effects and work team fixed effects.

We interpret these results as follows. Errors are rare. They usually happen in especially tense situations, typically when things go wrong and there is little time to fix them. In these situations of improvisation, older more experienced

workers seem to know better which severe errors to avoid by all means. This concentration on the vital tasks—potentially at the cost of some minor errors—implies that older workers perform better in terms of our overall productivity measure, the severity-weighted sum of errors. A detailed study of the regression coefficients in Table 2 reveals that the *severity* of errors is mostly explained by the age variables while the control variables are mostly insignificant. Hence, external conditions seem not to matter much. It is experience that prevents severe errors. For the *number* of errors, the opposite is true: Only three of the nine age splines are significant but almost all control variables. Higher age leads to more errors but other factors seem more important.

Figure 3: Age profile for the frequency of errors

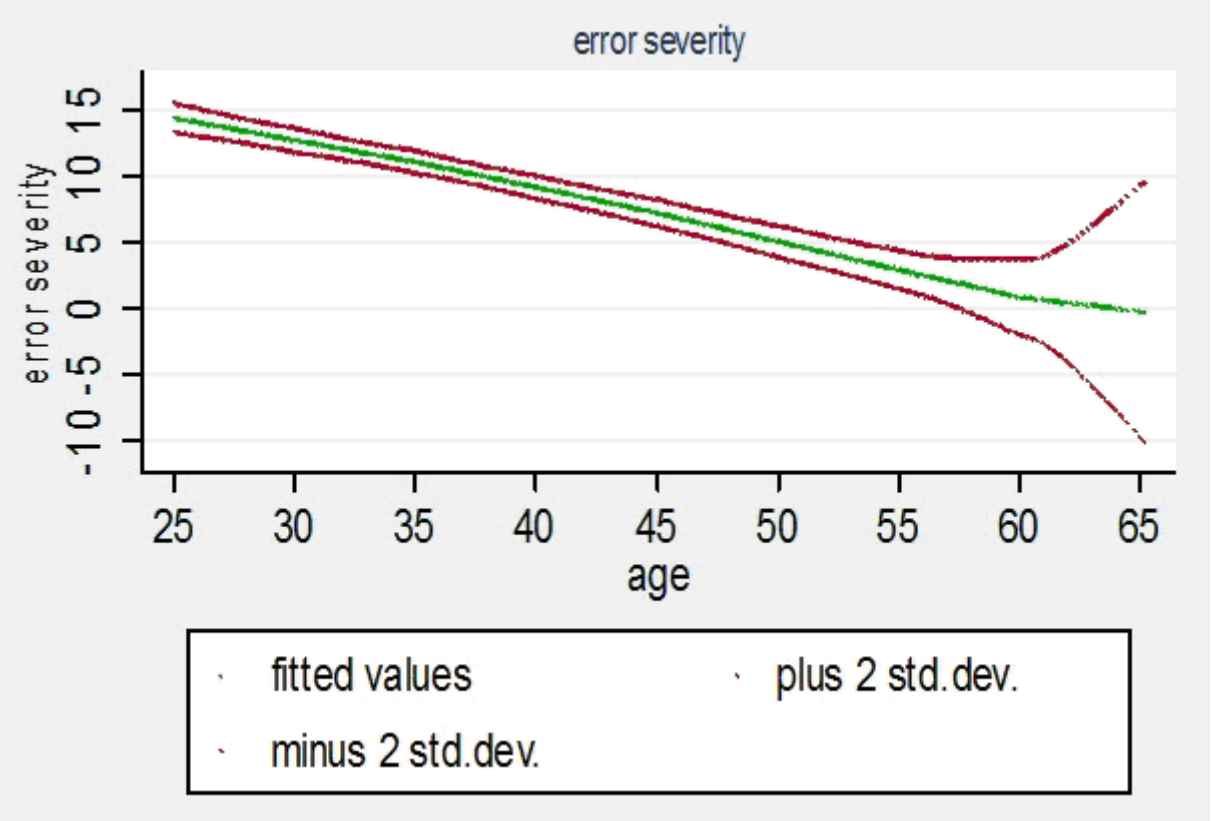

Figure 4: Age profile for the severity of errors

4.2 Regressions on the team level

We now change the observation unit to a *team* day where average team age and average job tenure are not perfectly correlated. This allows us to identify the effect of job tenure given age. In addition, we are able to include more control variables and interactions of these with age.

Age effects are again specified as piecewise linear 5-year-age splines.[20] In addition, average team age is interacted with a large set of control variables. Average job tenure of the team is described by 4-year linear splines. Since the focus of this paper is on the relation between productivity and age, and the age effects are a function of the parameters of the age splines and the interaction terms, we relegate the detailed regression coefficients to Appendix D and summarize our results in Figures 4 and 5 where the effect of age on inverse productivity (i.e., severity-weighted errors) is decomposed in an "experience effect" and a "residual age effect".[21]

[20] We also tried other specifications, e.g., polynomials and dummies for 5-year age groups. The results are robust with respect to these different specifications.

[21] What we call "residual age effect" here is a composition of effects that come along with age like deteriorating health, declining cognitive abilities, etc.

As can be seen in Figure 5, the "residual age effect" is positive. The sum of error weights is larger in older work teams if job tenure is held constant.

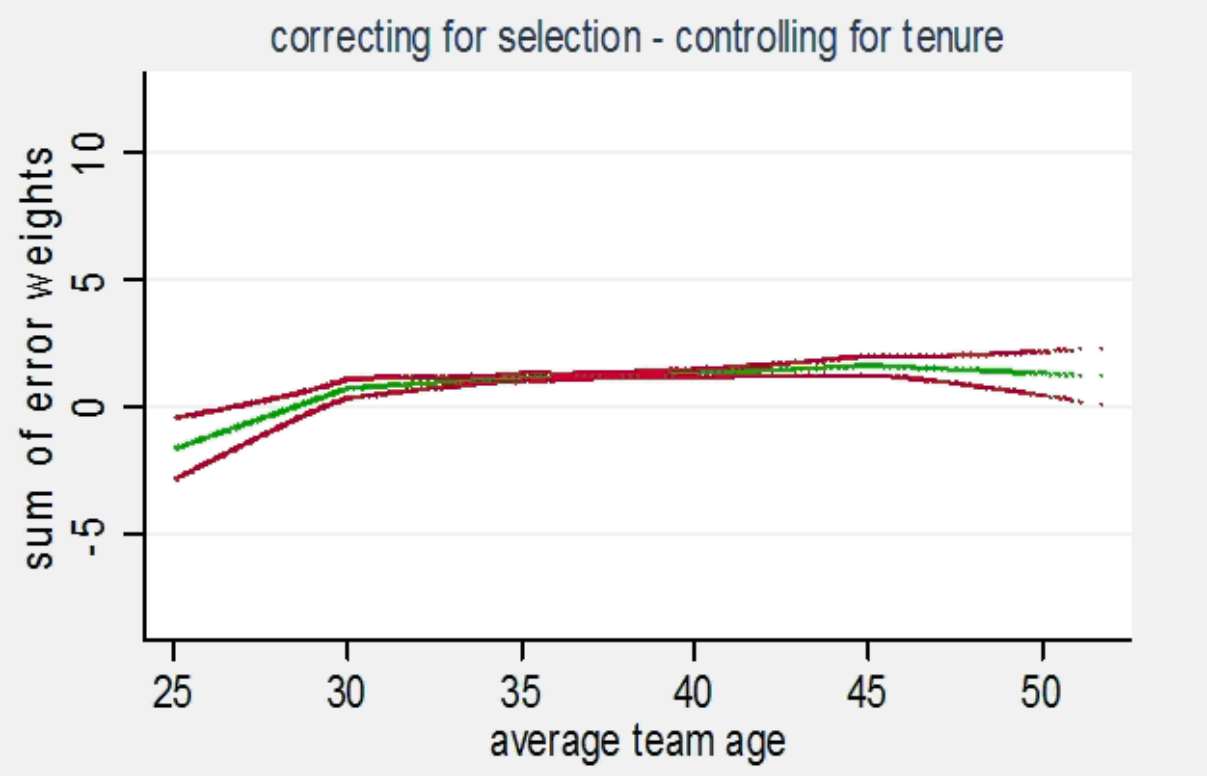

Figure 5: Sum of weighted errors: Inverse productivity by age

Figure 6: Sum of weighted errors: Inverse productivity by job tenure

In Figure 6, the weighted sum of errors is plotted against average job tenure. Holding average age constant, work teams with longer average job tenure and supposedly more experience have a lower sum of error weights. Hence, for workers who grow old in the plant, the productivity enhancing effect of accumulating more experience compensates the adverse "residual" age effect so that the overall age profile is rather flat. On the other hand, workers who are not able to accumulate experience in the plant face decreasing productivity as they age.

Our regressions contain a large set of control variables. Since 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. The corresponding figures are in Table D.3 in Appendix D.2. These interactions might also reflect differences between cohorts in addition to age-specific differences.

Up to an age of 40, 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. School education has changed over time and the interaction effects might reflect these differences.

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.

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.

One potential objection is that the assembly line may have been set at a low speed to accommodate older workers and could run faster (and thus increase productivity) if there were fewer older workers in the plant. Our evidence, however, suggests that this is not the case. We have included an interaction term between workload and age in our regression. Its coefficient is negative and insignificant, implying that, if anything, older workers are at least as good as younger workers in dealing with the higher workload generated by a faster running assembly line.

Cycle time, i.e., the time that workers have to perform their tasks on one car at the conveyer belt before the next car arrives, is an important variable to assess the external validity of our results. Cycle times are generally quite long in truck assembly (6 – 12 minutes) and much shorter in the production of passenger cars (1.5 – 5 minutes) and other assembly-line jobs. This difference may raise concerns that our results might be specific to truck assembly where older workers

are not stressed by short cycle times. Since one part of the assembly plant (where drivers' cabs are finished) has very short cycle times (2 – 3 minutes) variation in cycle times in our sample is between 2 and 12 minutes. While the interaction of cycle time with age has negative sign, implying that older workers find it harder to deal with short cycle times, the coefficient is tiny so that the age productivity profile is virtually unaffected. The coefficient of -0.003 means that a change in cycle time by 10 minutes changes the age gradient by -0.03 error points per year of age.

We also included dummy variables for days on which new types trucks were assembled on a prototypical basis. These "tryouts" require workers to adapt to (slightly) new procedures. The negative sign of these variables indicates that the difficulties associated with a new type where overcompensated by the production plan (more workers or fewer cars on that day) or effective training. This effect wears out in older work teams implying that older workers find it harder to adapt to new procedures. This effect, however, is not significant.

The richness of the production records in our data allowed us to include two weather variables. Weather might negatively affect productivity by reducing physical fitness (heat, humidity), by lowering workers' motivation or enhancing workers' spirits (nice weather). As it turns out, higher air pressure (implying nice weather) is good for productivity with no difference between age groups. On hot and humid days (>30°C and >70% relative humidity), older workers' productivity goes down while, maybe somewhat surprisingly, younger workers' productivity is higher on these days.[22]

Fluctuation is bad for productivity. This effect is stronger among older workers.

Young workers are more productive on weekends while older workers are more productive during the week. However, we cannot exclude that the selection into and out of weekend shifts is non-random.

Finally, the less than perfect correlation of average team age with average job tenure and time permits an estimation of a time trend which we also interact with average team age. Our results suggest that young workers' productivity decreased over the four years of the observation period while the productivity of the older workers increased. The age productivity profile has thus over time become more favorable for the old. Over all work teams, the number and severity of errors remained at the same level over these four years, corroborating our assumption that there are no time effects generated by technological or technical change

[22] After 2006, the assembly hall was equipped with air condition.

in the assembly plant that would bias our results.

5 Conclusions

Based on our study of the relation between workers' age and their productivity in an assembly plant of a truck manufacturer, we cannot confirm the widespread opinion that older workers are less productive than younger workers. We use data on errors made in the production process in which quality and quantity is as standardized as the time to produce a unit of output. This laboratory-like environment and a very large number of daily observations allows us to construct very precise inverse age-productivity profiles in the age range from 25 to 60 years.

Controlling for individual worker fixed effects and the availability of personnel data with an appropriate instrument for a Heckman-style selectivity correction allows us to adjust for sample selection bias. In addition, fixed effects prevent us from confounding age and cohort age effects. Controlling for work team fixed effects ensures that the remaining variation we use to estimate the age productivity profile is exogenous since the identifying fluctuation due to sick leave, vacation and the compensation for overtime does not leave any room for optimization by the management.

Our findings show that the average age-productivity profile of individual workers is increasing until age 65. Decomposition into the effect of job tenure and age reveals that it is indeed experience 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.

Our results refer to a single plant only. We believe, however, that our results are typical for large-scale manufacturing and thus of general interest. Moreover, by choosing a truck assembly plant in which physical strength and agility is still quite important, we have stacked our cards in favor of finding declining age-productivity, since these characteristics are well known to decline with age. Experience is likely to count even more, e.g., in the service industry. In addition, while the shape of the age productivity profile is certainly task-specific, the decomposition into (i) experience effects and residual age effects, and into (ii) effects

on frequency and severity of errors, seems to be more fundamental. Finally, we hope that this kind of work encourages researchers to replicate this study in other work environments.

Appendix

A 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. 71 percent of all errors in our data set can be related to one single work team. For the remaining 2490 of 8564 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. The resulting heteroskedasticity is taken into account by computing Huber-White robust standard errors.

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.

B Descriptive statistics

Table B.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
# consecutive 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
hot and humid (dummy, >30°C, >70% rel. humidity)	0.041	0	0	1	0.197
air pressure (hPa)	1000.4	1000.5	968.5	1022.3	7.35

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 popu-

lation. 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:

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

Accordingly, selection for older workers is

$$s_{it}^o = 1\left[z_{it}' \cdot \gamma^o + \varepsilon_{it} > 0 \right], \quad \varepsilon_{it} \sim N(0,1) \text{ i.i.d.} \tag{2}$$

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

$$P\left(z_{it}' \cdot \gamma + \varepsilon_{it} > 0 \right) = P\left(\varepsilon_{it} > -z_{it}' \cdot \gamma \right) \overset{symetry}{=} P\left(\varepsilon_{it} < z_{it}' \cdot \gamma \right) = \Phi\left(z_{it}' \cdot \gamma \right) \tag{3}$$

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

$$y_{it} = x_{it}' \cdot \beta + u_{it} \tag{4}$$

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

$$E\left(u_{it} \middle| \varepsilon_{it} \right) = \begin{cases} \xi^y \cdot \varepsilon_{it} & \Leftrightarrow \quad i \quad \text{young} \\ \xi^o \cdot \varepsilon_{it} & \Leftrightarrow \quad i \quad \text{old} \end{cases} \tag{5}$$

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

$$E\left(u_{it} \middle| \varepsilon_{it} > -z_{it}' \cdot \gamma \ \ \forall i \right) = \begin{cases} \xi^y \cdot E\left(\varepsilon_{it} \middle| \varepsilon_{it} > -z_{it}' \cdot \gamma^y \right) & \Leftrightarrow \quad i \quad \text{young} \\ \xi^o \cdot E\left(\varepsilon_{it} \middle| \varepsilon_{it} > -z_{it}' \cdot \gamma^y \right) & \Leftrightarrow \quad i \quad \text{old} \end{cases} = \begin{cases} \xi^y \cdot \frac{\phi\left(z_{it}' \cdot \gamma^y \right)}{\Phi\left(z_{it}' \cdot \gamma^y \right)} & \Leftrightarrow \quad i \quad \text{young} \\ \xi^o \cdot \frac{\phi\left(z_{it}' \cdot \gamma^o \right)}{\Phi\left(z_{it}' \cdot \gamma^o \right)} & \Leftrightarrow \quad i \quad \text{old} \end{cases} \tag{6}$$

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

$$E\left(y_{it} \mid x_{it}, s_{it} = 1 \ \ \forall i \right) = E\left(y_{it} \middle| x_{it}, \varepsilon_{it} > -z_{it}' \cdot \gamma \right) = E\left(x_{it}' \cdot \beta + u_{it} \middle| x_{it}, \varepsilon_{it} > -z_{it}' \cdot \gamma \right) \tag{7}$$

$$E\left(y_{it} \mid x_{it}, s_{it} = 1 \right) = x_{it}' \cdot \beta + E\left(u_{it} \middle| x_{it}, \varepsilon_{it} > -z_{it}' \cdot \gamma \right) \tag{8}$$

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

$$E\left(y_{it} \mid x_{it}, s_{it}=1 \ \ \forall i\right)=x_{it}^{'}\cdot\beta+\begin{cases}\xi^{y}\cdot\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)} & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)} & \Leftrightarrow \quad i \quad \text{old}\end{cases} \tag{9}$$

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

$$y_{jt}=x_{jt}^{'}\cdot\beta+u_{jt} \tag{10}$$

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

$$E\left(u_{jt}\left|\varepsilon_{it}\right.\right)=\begin{cases}\xi^{y}\cdot\varepsilon_{it} & \Leftrightarrow \quad i \quad \text{young}\\ \xi^{o}\cdot\varepsilon_{it} & \Leftrightarrow \quad i \quad \text{old}\end{cases} \tag{11}$$

and

$$E\left(u_{jt}\left|\left\{\varepsilon_{it}\right\}_{i=1}^{N_j}\right.\right)=\xi^{y}\cdot\sum_{i\ \text{young}}\varepsilon_{it}+\xi^{o}\cdot\sum_{i\ \text{old}}\varepsilon_{it} \tag{12}$$

This implies that within the young and within the old, each individual ε_{it} of any worker i has the same effect on the team's 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?

$$\begin{aligned}E\left(u_{jt}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma \ \ \forall i\right.\right)&=\xi^{y}\cdot\sum_{i\ \text{young}}E\left(\varepsilon_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma^{y} \ \ \forall i\right.\right)+\xi^{o}\cdot\sum_{i\ \text{old}}E\left(\varepsilon_{it}\left|\varepsilon_{it}>-z_{it}^{'}\cdot\gamma^{y} \ \ \forall i\right.\right)\\ &=\xi^{y}\cdot\sum_{i\ \text{young}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}+\xi^{o}\cdot\sum_{i\ \text{old}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}\end{aligned} \tag{13}$$

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

$$E\left(y_{jt} \mid x_{jt}, s_{it}=1 \ \ \forall i\right)=E\left(y_{jt} \left| x_{jt}, \varepsilon_{it}>-z_{it}^{'}\cdot\gamma \ \ \forall i\right.\right)=E\left(x_{jt}^{'}\cdot\beta+u_{jt} \left| x_{it}, \varepsilon_{it}>-z_{it}^{'}\cdot\gamma \ \ \forall i\right.\right) \quad (14)$$

$$E\left(y_{jt} \mid x_{jt}, s_{it}=1 \ \ \forall i\right)=x_{jt}^{'}\cdot\beta+E\left(u_{jt} \left| x_{jt}, \varepsilon_{it}>-z_{it}^{'}\cdot\gamma \ \ \forall i\right.\right) \quad (15)$$

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

$$E\left(y_{jt} \mid x_{jt}, s_{it}=1 \ \ \forall i\right)=x_{jt}^{'}\cdot\beta+\xi^{y}\cdot\sum_{i \text{ young}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{y}\right)}+\xi^{o}\cdot\sum_{i \text{ old}}\frac{\phi\left(z_{it}^{'}\cdot\gamma^{o}\right)}{\Phi\left(z_{it}^{'}\cdot\gamma^{o}\right)} \quad (17)$$

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 (≥ 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. A higher sickness rate increases the probability of remaining in the sample for the young and it decreases the probability of remaining in the sample for the old. This is in line with the notion the young workers who leave the sample are a positive selection while old workers who leave the sample are a negative selection. In addition, we include for most variables the individual values as well as team averages which further improves identification of the inverse Mills ratios in the second stage.

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	reference category			
20 < age < 22	-0.0130	(0.260)		
22 < age < 24	0.125	(0.000)		
24 < age < 26	0.178	(0.000)		
26 < age < 28	0.360	(0.000)		
28 < age < 30	0.543	(0.000)		
30 < age < 32	0.680	(0.000)		
32 < age < 34	0.781	(0.000)		
34 < age < 36	0.984	(0.000)		
36 < age < 38	1.01	(0.000)		
38 < age < 40	0.953	(0.000)		
40 < age < 42			reference category	
42 < age < 44			-0.351	(0.005)
44 < age < 46			0.0770	(0.000)
46 < age < 48			-0.0114	(0.376)
48 < age < 50			-0.0694	(0.000)
50 < age < 52			0.0236	(0.084)
52 < age < 54			0.0659	(0.000)
54 < age < 56			-0.375	(0.000)
56 < age < 58			-0.916	(0.000)
58 < age < 60			-2.03	(0.000)
60 < age < 62			-1.73	(0.000)
62 < age < 64			-0.526	(0.000)
64 < age			-1.70	(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 in team	-3.93	(0.000)	-2.38	(0.000)
French dummy	1.06	(0.000)	0.418	(0.000)
share of French in team	-5.26	(0.000)	-3.96	(0.000)
Turkish dummy	1.45	(0.000)	0.259	(0.000)
share of Turkish in team	-3.61	(0.000)	-5.13	(0.000)
female dummy	-0.556	(0.000)	-0.339	(0.000)
share of women in team	-0.821	(0.000)	-0.121	(0.250)
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 in team	-3.30	(0.000)	-1.56	(0.000)
team size	-0.00659	(0.000)	0.0255	(0.000)
av. Team tenure	0.00438	(0.009)	0.0547	(0.000)
Finishing of driver's cabin	0.0741	(0.000)	-0.640	(0.000)
Constant	0.190	(0.065)	5.92	(0.000)
R^2	0.247		0.285	
# observations	2030939		1164115	
p-values in parentheses				

D Detailed regression results (team level)

Table D.1: Regression results (team level)

Dependent variable: sum of error weights					
correcting for sample selection, controlling for job tenure					
age splines			job tenure splines		
20 – 25 years	7.51	(2.02)	0 – 4 years	-1.62	(0.360)
25 – 30 years	4.52	(0.789)	4 – 8 years	0.0155	(0.0668)
30 – 35 years	4.13	(0.782)	8 – 12 years	-0.0908	(0.0353)
35 – 40 years	4.08	(0.781)	12 – 16 years	0.0814	(0.0319)
40 – 45 years	4.11	(0.780)	16 – 20 years	-0.148	(0.0309)
45 – 55 years	3.98	(0.786)	20 – 24 years	-0.237	(0.0649)
			24 – 32 years	-0.849	(0.205)
control variables			interactions of these control variables with average team age		
schooling years	3.05	(0.449)	-0.0764	(0.0122)	
car specific educ	-1.35	(2.34)	0.0497	(0.0608)	
tech spec. educ	3.92	(2.40)	-0.0963	(0.0621)	
Female	26.0	(5.29)	-0.693	(0.145)	
French	-7.62	(4.31)	0.180	(0.112)	
German	-10.5	(3.95)	0.259	(0.104)	
Turkish	-6.82	(5.69)	0.186	(0.149)	
External	-3.45	(2.56)	0.0694	(0.0659)	
team size	0.138	(0.104)	0.00793	(0.00244)	
(team size)^2	-0.00644	(0.00110)			
late shift	-0.000165	(0.413)	-0.00261	(0.0106)	
days w/o change	0.0325	(0.166)	-0.000872	(0.000421)	
(days w/o change)^2	-0.0000488	(0.0000310)			
cycle time	0.276	(0.0544)	-0.00298	(0.00136)	
Workload	8.24	(2.43)	-0.0250	(0.0603)	
(workload)^2	-12.5	(0.936)			
tryout Axor	-2.01	(1.08)	0.0375	(0.0276)	
tryout Atego	0.969	(1.08)	-0.0223	(0.0276)	
air pressure	-0.0122	(0.0276)	0.0000249	(0.000699)	
hot and humid	-4.62	(1.00)	0.113	(0.0264)	
Monday	8.37	(0.994)	-0.198	(0.0257)	
Tuesday	6.78	(0.983)	-0.158	(0.0254)	
Wednesday	12.4	(0.981)	-0.290	(0.0253)	
Thursday	7.84	(0.986)	-0.184	(0.0255)	
Friday	12.2	(0.987)	-0.282	(0.0255)	
time trend	2.91	(0.249)	-0.0717	(0.00635)	
Inverse Mills Ratio young	0.246	(0.0251)			
Inverse Mills Ratio old	0.366	(0.0513)			
adj. R^2 within	0.044				
adj. R^2 between	0.081				

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

Table D.1 includes age as a function of age splines and interactions of age with a large set of control variables given by:

$$sum\ of\ \ error\ \ weights = \beta_0 + \sum_{a=20}^{50} \beta_{a,a+5} \cdot AgeSpline_{a,a+5} + \sum_k \beta_k \cdot x_k + \sum_k \beta_{age,k} \cdot x_k \cdot AverageAge + \varepsilon \quad (18)$$

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 of the errors-age-profile (evaluated at sample means and displayed in Figure 3) 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:

$$\left.\frac{\partial\left(sum\ of\ \ errorweights\right)}{\partial AverageAge}\right|_{AverageAge\ =\ 37} = \beta_{35-40} + \sum_k \beta_{age,k} \cdot \overline{x_k} \quad (19)$$

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 D.2.

Table D.2: Residual age gradients and job tenure gradients

age	residual age gradient	job tenure	tenure gradient
20 – 25 years	3.48 (1.86)	0 – 4 years	-1.62 (0.360)
25 – 30 years	0.493 (0.118)	4 – 8 years	0.0155 (0.0668)
30 – 35 years	0.101 (0.0423)	8 – 12 years	-0.0908 (0.0353)
35 – 40 years	0.0451 (0.0280)	12 – 16 years	0.0814 (0.0319)
40 – 45 years	0.0796 (0.0348)	16 – 20 years	-0.148 (0.0399)
45 – 55 years	-0.0467 (0.0737)	20 – 24 years	-0.237 (0.0649)
		24 – 32 years	-0.849 (0.205)

Gradients are calculated from coefficients in Table 2 at mean values of all control variables. Huber-White robust standard errors are in parentheses.

Table D.3 reports the marginal effects of selected control variables at different ages.

Table D.3: Marginal effects of some control variables at different ages

	years of schooling	share of women	French	German	Turkish	team size	days w/o change
30 years	0.762 (0.0973)	5.19 (1.12)	-2.22 (1.05)	-2.71 (0.924)	-1.24 (1.36)	0.233 (0.0293)	0.00537 (0.00435)
40 years	-0.00346 (0.0691)	-1.74 (0.873)	-0.398 (0.550)	-0.119 (0.492)	0.610 (0.732)	0.312 (0.0194)	-0.00331 (0.00262)
50 years	-0.769 (0.173)	-8.67 (2.12)	1.43 (1.42)	2.47 (1.34)	2.46 (1.92)	0.391 (0.0329)	-0.0120 (0.00550)

	excess workload	cycle time	tryout Axor	air pressure	hot and humid
30 years	7.43 (0.689)	0.185 (0.0152)	-0.882 (0.268)	-0.0113 (0.00696)	-1.24 (0.232)
40 years	7.21 (0.351)	0.156 (0.00785)	-0.507 (0.109)	-0.0113 (0.00263)	-0.104 (0.117)
50 years	6.98 (0.706)	0.127 (0.0162)	-0.131 (0.323)	-0.0112 (0.00794)	1.03 (0.336)

	Monday	Tuesday	Wednesday	Thursday	Friday	time trend
30 years	2.44 (0.244)	2.05 (0.242)	3.70 (0.241)	2.33 (0.243)	3.79 (0.242)	0.761 (0.0643)
40 years	0.468 (0.117)	0.477 (0.116)	0.809 (0.116)	0.498 (0.117)	0.971 (0.116)	0.0464 (0.0301)
50 years	-1.51 (0.316)	-1.10 (0.312)	-2.08 (0.312)	-1.34 (0.313)	-1.85 (0.314)	-0.668 (0.0758)

Gradients are calculated from coefficients in the second column of Table 2. Huber-White robust standard errors are in parentheses. The reference category for weekday dummies is saturdays and sundays on some of which the assembly line was running.

E Sample of a Production Program

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	0	0	0	0	10	4	13	20	30	21
21.07. - 01.08.03	3	50	4	0	0	0	0	7	6	11	16	26	19
04.08. - 15.08.03	4	49	6	0	0	0	0	5	6	10	15	27	21
18.08. - 29.08.03	3	49	7	0	0	0	0	4	5	12	14	30	21
01.09. - 12.09.03	3	50	6	0	0	0	0	4	7	11	12	35	18
15.09. - 26.09.03	4	58	8	0	0	0	0	5	9	12	17	48	27
29.09. - 10.10.03	5	52	14	0	0	0	0	5	7	12	21	51	24
13.10. - 24.10.03	6	51	14	0	0	0	0	5	5	12	21	49	28
27.10. - 31.10.03	4	62	8	0	0	0	0	3	8	13	26	45	23
03.11. - 14.11.03	7	52	12	0	0	0	0	5	4	13	15	59	23
17.11. - 28.11.03	5	54	14	0	0	0	0	4	5	15	18	52	25
01.12. - 12.12.03	4	51	15	0	0	0	0	3	6	15	21	42	29
15.12. - 09.01.04	6	48	13	0	0	0	0	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, Patrick (2003): "Productivity, wage and demand for elder workers; an examination on French matched employer-employee data," Working Paper INSEE, Paris.

Aubert, Patrick, and Bruno Crépon (2007): "Are Older Workers Less Productive? Firm-Level Evidence on Age-Productivity and Age-Wage Profiles," mimeo.

Backes-Gellner, Uschi, Martin R. Schneider, and Stephan Veen (2011): "Effect of Workforce Age on Quantitative and Qualitative Organizational Performance: Conceptual Framework and Case Study Evidence", *Organization Studies* 32 (8), 1103 - 1121.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul (2005): "Social Preferences and the Response to Incentives: Evidence from Personnel Data", *Quarterly Journal of Economics* 120 (3), 917 – 962.

—— (2007): "Incentives for Managers and Inequality among Workers: Evidence from a Firm-Level Experiment", *Quarterly Journal of Economics* 122 (2), 729 – 773.

—— (2009): "Social Connections and Incentives in the Workplace: Evidence from Personnel Data", *Econometrica* 77 (4), 1047 - 1094.

—— (2010): "Social Incentives in the Workplace", *Review of Economic Studies* 77 (2), 417 – 458.

Bayer, Thomas, John Page, Yaron Raviv, and Joshua Rosett (2009): "Age, Human Capital and the Quality of Work: New Evidence from Old Masters," Robert Day School Working Paper No. 2009-04.

Bloom, David, Alfonso Sousa-Poza (2013): "Aging and Productivity: Introduction," *Labour Economics*, 22, 1 – 4.

Castellucci, Fabrizio, Mario Padula, and Giovanni Pica (2011): "The Age-Productivity Gradient: Evidence from a Sample of F1 drivers," *Labour Economics* 18 (4), 464 – 473.

Crépon, Bruno, Nicolas Deniau, and Sébastien 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, Francesco and Mika Maliranta (2007): "Age, seniority and labour costs: lessons from the Finnish IT revolution," *Economic Policy* 22 (49), 117 – 175.

Dostie, Benoit (2011): "Wages, Productivity, and Aging," *De Economist* 159 (2), 139 – 158.

Fair, Ray (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, Robert H., and Robert M. Hutchens (1993): *Journal of Economic Behavior and Organization*, 21(3), 251 — 276.

Galenson, David W. (2009): "The Greatest Artists of the Twentieth Century," Chapter 2 in: Galenson, David W. (ed) *Conceptual Revolutions in Twentieth-Century Art,* 30 – 61.

Galenson, David W., and Bruce 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.

Gelderblom, Arie (2006): "The Relationship of Age with Productivity and Wages", Chapter 6 in: *Ageing and Employment: Identification of Good Practice to Increase Job Opportunities and Maintain Older Workers in Employment*, Final Report to the European Commission, DG Employment, http://ec.europa.eu/social/BlobServlet?docId=2073&langId=en

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

—— (2013): "Are personnel measures effective in increasing productivity of old workers?," *Labour Economics*, 22, 80 – 93.

Grund, Christian, and Niels Westergård-Nielsen (2008): "Age Structure of the Workforce and Firm Performance," *International Journal of Manpower*, 29(5), 410 – 422.

Hægeland, T Torbjørn and Tor J. Klette (1999): "Do higher wages reflect higher productivity? Education, gender and experience premiums in a matched plant-worker data set" In: John C. Haltiwanger, Julia I. Lane, James R. Spletzer, J. Theeuwes, and Kenneth Troske (Eds.) *The Creation and Analysis of Employer-Employee Matched Data*. Holland, Elsevier Science.

Haltiwanger, John C., Julia I. Lane, and James 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*, 14, 575 – 602.

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

—— (2007: "Production Function and Wage Equation Estimation with Heterogeneous Labor: Evidence From a New Matched Employer-Employee Data Set," in *Hard-to-Measure Goods and Services: Essays in Memory of Zvi Griliches*, ed. Charles Hulten and Ernst Berndt, 31-71, Chicago: University of Chicago Press.

Hellerstein, Judith K., David Neumark, and Kenneth 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, John E., and Ronda F. Hunter (1984): "Validity and Utility of Alternative Predictors of Job Performance," *Psychological Bulletin*, 96, 72 — 98.

Ilmakunnas, Pekka, and Mika Maliranta (2005): "Technology, Labour Characteristics and Wage-Productivity Gaps," *Oxford Bulletin of Economics and Statistics*, 67(5), 623 — 645.

—— (2007): "Aging, Labor Turnover, and Firm Performance," ETLA Discussion Paper No. 1092 http://www.etla.fi/.

Jones, Benjamin F. (2010): "Age and Great Invention," *Review of Economics and Statistics* 92(1), 1 – 14.

Jones, Benjamin F. and Bruce Weinberg (2011): "Age Dynamics in Scientific Creativity," *Proceedings of the National Academy of Sciences* 108 (47), 18855 – 19096.

Kotlikoff, Laurence J., and David 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, Laurence J., and Jagadeesh 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, John, and Dmitriy Stolyarov (2005): "Technological Progress and Worker Productivity at Different Ages," mimeo, University of Michigan.

Lallemand, Thierry, and François Rycx (2009): "Are Older Workers Harmful for Firm Productivity?" *De Economist* 157 (3), 273 – 292.

Lazear, Edward P. (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.

—— (1995): *Personnel Economics*, The MIT Press, Cambridge, Mass.

—— (2000): "Performance Pay and Productivity", *American Economic Review* 90 (5), 1346 – 1361.

Lazear, Edward P., Kathryn L. Shaw, and Christopher T. Stanton (2012): "The Value of Bosses", NBER Working Paper No. 18317.

Loewenstein, George, and Nachum Sicherman (1991): "Do Workers Prefer Increasing Wage Profiles?," *Journal of Labor Economics*, 9(1), 67 — 84.

Lovász, Anna, Mariann Rigó (2013): "Vintage effects, aging and productivity," *Labour Economics*, 22, 47 – 60.

Mas, Alexandre, and Enrico Moretti (2009): "Peers at Work", *American Economic Review* 99 (1), 112 – 145.

Malmberg, Bo, Thomas Lindh, and Max Halvarsson (2008): "Productivity Consequences at the Plant Level of Work-Force Ageing: Stagnation or a Horndal Effect?," *Population and Development Review*, 34, 238 – 256.

Maddison, Angus (2001): *The World Economy: A Millennial Perspective*, OECD Development Center Studies, Paris.

Mahlberg, Bernhard, Inga Freund, Jesús Crespo Cuaresma, Alexia Prskawetz (2013): "Ageing, productivity and wages in Austria," *Labour Economics*, 22, 5 – 15.

McEvoy, Glenn M., and Wayne F. Cascio (1989): "Cumulative Evidence of the Relationship Between Employee Age and Job Performance," *Journal of Applied Psychology*, 74(1), 11 — 17.

Medoff, James L., and Katharine G. Abraham (1980): "Experience, Performance, and Earnings," *Quarterly Journal of Economics*, 95(4), 703 — 736.

Prskawetz, Alexia, Bernd Mahlberg, Vegard Skirbekk, Inga Freund, Maria Winkler-Dworak, Thomas Lindh, Bo Malmberg, Ann-Christin Jans, Oskar Skans Nordström, and Fredrik 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.

Romeu Gordo, Laura, Vegard Skirbekk (2013): "Skill demand and the comparative advantage of age: Jobs tasks and earnings from the 1980s to the 2000s in Germany," *Labour Economics*, 22, 61 – 69.

Salthouse, Timothy A., and Todd 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, Lutz (2007): "Mit 55 zum alten Eisen? Eine Analyse des Alterseinflusses auf die Produktivität anhand des LIAB", *Zeitschrift für ArbeitsmarktForschung*, 40 (1),77- 97.

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

Skirbekk, Vegard (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.

Vandenberghe, Vincent (2013): "Are firms willing to employ a greying and feminizing workforce?," *Labour Economics*, 22, 30 – 46.

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

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