

Age, Wage and Productivity in Dutch Manufacturing

Jan C. van Ours · Lenny Stoeldraijer

Published online: 24 April 2011

© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract Previous empirical studies on the effect of age on productivity and wages find contradicting results. Some studies find that if workers grow older there is an increasing gap between productivity and wages, i.e. wages increase with age while

The authors thank the Dutch Ministry of Social Affairs and Employment for financial support; the estimates in this paper are based on own calculations using various micro datasets made available through a remote access facility by Statistics Netherlands. A preliminary version of the paper was presented in a labor-health seminar at Tilburg University and in seminars at Australian National University (RSSS—Canberra) and University of New South Wales (School of Economics—Sydney). We thank participant at these seminars as well as Sjef Ederveen, Jos Jansen, Henk Nijboer, Matteo Picchio, Daniel van Vuuren and two anonymous referees for comments on a previous version of the paper.

J. C. van Ours (✉)

Department of Economics, Tilburg University, Tilburg, The Netherlands
e-mail: vanours@uvt.nl; ecomist@uvt.nl

J. C. van Ours

Department of Economics, University of Melbourne, Melbourne, Australia

J. C. van Ours

CentER, Tilburg, The Netherlands

J. C. van Ours

CEPR, London, UK

J. C. van Ours

CESifo, Munich, Germany

J. C. van Ours

IZA, Bonn, Germany

L. Stoeldraijer

Statistics Netherlands, The Hague, The Netherlands
e-mail: lenny_stoeldraijer@hotmail.com

productivity does not or does not increase at the same pace. However, other studies find no evidence of such an age related pay-productivity gap. We perform an analysis of the relationship between age, wage and productivity using a matched worker-firm panel dataset from Dutch manufacturing covering the period 2000–2005. We find little evidence of an age related pay-productivity gap.

Keywords Age · Wage · Productivity

JEL Classification J23 · J31

1 Introduction

In many countries, including the Netherlands, population is aging. Therefore, there is increasing attention to the labor market position of older workers. The current situation is that older workers are not very likely to lose their job but once they have lost their job they need a long time to find a new job. This situation is often attributed to the gap between wages and productivity, i.e. older workers may have a wage that is higher than their productivity. At their current employer they are protected by seniority rules and employment protection legislation. But once unemployed, employers are reluctant to hire an older worker because of the pay-productivity gap. In this paper we explore empirically to what extent there is indeed evidence of a pay-productivity gap in Dutch manufacturing firms.

The main theoretical framework on the relationship between age and productivity is based on [Lazear \(1979\)](#). Age-earnings profiles are thought to be upward sloping because this will discourage workers from shirking. Workers and firms engage in long-term relationships in which the worker is initially underpaid—the wage is lower than the value of the marginal product—but later on in life the worker is overpaid. Such delayed-compensation contracts will discourage the worker from shirking, but at the same time require mandatory retirement to avoid firms paying more than the value of the marginal product—averaged over the working life i.e. over the duration of the contract between workers and firms. The theory of Lazear requires that workers and firms want to be engaged in long-term relationships and assumes that rising earnings do not fully reflect increased productivity. There are other explanations for the existence of a potential pay-productivity gap at higher ages (see [Van Vuuren and de Hek \(2009\)](#) for an overview). Such a gap may also arise from union bargaining. If unions care more for the preferences of senior workers, wages increase according to seniority. Incumbent workers controlling the union exploit newcomers. Furthermore, employment protection legislation in particular the last-in first-out rule may protect older workers more than younger workers. Due to this rule firms cannot simply replace high wage older workers for low wage young workers. Finally, the age related pay-productivity gap may be due to workers preferring increasing wage profiles over flat or decreasing wage profiles of greater monetary value.

Aging may affect productivity levels for various reasons. Older workers are thought to be more reliable and to have better skills than younger workers. Older workers are

generally considered to be more consistent, cautious, and conscientious. Furthermore, older workers have fewer accidents and they are less likely to quit, thus reducing hiring costs (Garibaldi et al. 2010). However, older workers also have higher health care costs, less flexibility in accepting new assignments and fewer incentives for training (Barth et al. 1993).

It is difficult to establish how age itself affects labor productivity not only because productivity is highly individual and sector-specific but also because of a convolution of age, cohort and selection effects. Individual productivity is complex and multi-dimensional. Among characteristics that relate to productivity are communication skills, information processing speed, strength and endurance, health, self-discipline, flexibility, administrative and strategic capacities, math proficiency, vocabulary size, education, motivation, energy and job experience. Productivity of individual workers is hardly ever observed. Nevertheless, if individuals are aggregated to the firm level, the relationship between productivity and age should still hold. To establish the relationship between age and productivity preferably matched worker-firm panel data are needed.

In this paper we present an empirical analysis of the relationship between age and productivity using matched worker-firm panel data from Dutch manufacturing. The set-up of our paper is as follows. In Sect. 2 we provide a general discussion on the relationship between population aging and productivity. In Sect. 3 we give a systematic overview of previous studies on age and productivity. In Sect. 4 we describe our matched-worker firm data. In Sect. 5 we present the set-up of our analysis. Section 6 discusses the parameter estimates, Sect. 7 concludes. Appendix A provides details on the data we use.

2 Workforce Aging and Productivity

In a perfectly competitive labor market there is no reason for an age-related pay-productivity gap to occur because firms pay workers according to (marginal) productivity. If a firm experiences a negative productivity shock and wages do not adjust the least productive workers are fired until the equality of productivity and wages is restored. If a firm experiences a positive productivity shock and wages do not adjust new workers will be hired until again the equality of productivity and wages is restored. With the existence of labor market institutions, imperfect information and/or costly monitoring of productivity the direct relationship between age and productivity disappears and an age-related pay-productivity gap may occur.

According to Johnson (1993), most employers and probably most employees seem to believe in a rule of thumb that average labor productivity declines after some age between 40 and 50. This assumption is so common that few attempts have been made to gather supporting evidence; why bother to prove the obvious? It is not easy to establish the relationship between age and productivity for a variety of reasons: productivity is a complex phenomenon, the age-productivity profile changes over time, is potentially endogenous and individual profiles are not easily aggregated (Garibaldi et al. 2010).

Productivity is multidimensional with the effects of aging depending on the extent to which age-induced changes in work experience, physical strength and cognitive abilities are relevant for work performance. Accumulated experience benefits employees performance throughout the working life. However, physical strength and health are reduced as workers grow older. Concerning cognitive abilities the age effect is not uniform. On the basis of a literature survey, *Skirbekk (2003)* concludes that job performance decreases after age 50, in particular for jobs where problem solving, learning and speed are important and less so for jobs where experience and verbal abilities are important. Some cognitive abilities such as vocabulary size and verbal ability increase to a relatively late age or remain stable throughout the working life. Other cognitive abilities such as speed and memory deteriorate with age. Different types of work require different cognitive abilities and physical strength. The net effect of the age-specific productivity determinants depend on how individual skills are used in the work process, how the work is organized and how the individual interacts with other workers and firm level factors such as technology. The variation in the type of skill required in the workplace is likely to cause differences in the age-productivity pattern across occupations.

A further complication in assessing the importance of the age effect on productivity are calendar time changes in the age-productivity profile. This is due to changes in the demand for skills as well as changes in individual characteristics across generations. New technologies and new working techniques imply a decreased need for manual labor and an increase in the demand for high skilled workers. Physical strength and bodily co-ordination have become less important in the workplace, while cognitive abilities are increasingly important. Over time, health levels have improved at adult and older ages. Improved age-specific mental and physical health levels are likely to create a strong positive effect on the work potential of older workers. In other words, even if it would be possible to establish the current relationship between age and productivity, drawing conclusions about the future age-productivity profile is not straightforward.

Finally, there is the issue of aggregation. Aging has an unambiguous effect on the age structure of the population: the number of old-age individuals increases relative to young and prime age individuals. Aging will affect the size and composition of the workforce but as yet the consequences at aggregate level are unclear. If workers of different ages are imperfect substitutes in production aging will affect relative wages of younger and older workers. When young workers become scarcer the relative wage of young workers should rise. Whether they actually do also depends on the labor market institutions. Therefore, it is difficult to draw straightforward conclusions except for the obvious conclusion that even if at the individual level productivity would decline with age there may be compensating effects at the aggregate level.

3 Previous Matched Worker-Firm Data Studies on Age and Productivity

Job performance or productivity of workers can be investigated using worker-firm matched data sets. We first discuss studies based on cross-sectional data. Then we discuss panel data studies.

3.1 Cross-Sectional Data

Initially, studies using matched worker-firm data used cross-sectional information. [Hellerstein et al. \(1999\)](#) is one of the first studies of this kind. Plant-level production function and wage equations are estimated simultaneously to compare productivity and wages for various groups of workers.¹ The results indicate that prime-age workers are as productive as younger workers. For prime age workers and older workers productivity and earnings rise at the same rate over the life cycle. The authors conclude that their evidence is most consistent with models in which wages rise in accordance with productivity, such as the general human capital model. [Hellerstein and Neumark \(2004\)](#) use a much larger data set to perform similar estimates. They find that both wage and productivity profiles are rising but concave; the estimated wage profile being steeper than the productivity profile, supporting Lazear's idea that young workers are underpaid and older workers are overpaid.

[Lallemand and Rycx \(2009\)](#) investigate the effects of the workforce age structure on the productivity of large Belgian firms in two subsequent cross-sections. They find that a higher share of young workers within a firm is favorable to firms' productivity while a higher share of older workers is harmful.

3.2 Panel Data

Many longitudinal studies on matched worker-firm data draw on the work of [Hellerstein et al. \(1999\)](#). The method used by [Hellerstein et al. \(1999\)](#) is expanded by implementing panel data estimation which tries to remove some of the bias associated to OLS estimation of production functions. One of the first studies which uses matched employer-employee panel data is [Haegeland and Klette \(1999\)](#) who estimate plant level production functions using Norwegian data for the period 1986–1993. Their results indicate that the higher wages earned by workers with higher education largely correspond to their higher productivity. Experienced workers are more productive than inexperienced workers and they also earn higher wages. [Crépon et al. \(2003\)](#) use French matched worker-firm data providing a new method using costs for the employer. They find that older workers are relatively overpaid. The age profile of wages is concave while the age profile of productivity stops rising—and even decreases—after some experience level. [Crépon et al. \(2003\)](#) conclude that for workers over 35 increases in wages are not a reflection of human capital accumulation.² [Aubert \(2003\)](#) using a large French data set shows that the endogeneity bias of inputs and unobserved shocks is large. After correcting for this, the age-productivity profile is increasing and concave, with highly skilled workers having the steepest

¹ For the production function, they use a output specification, which they instrument with lagged materials because these are likely to be endogenous, and a value-added specification. Also for the wage equations they use three compensation measures.

² From this the authors conclude that the early retirement policy in France is consistent with this evidence and they conclude that a policy of raising the mandatory retirement age may be problematic because of the poor performance of older workers in the labor market. They do not consider the possibility that the age-productivity profile may change if the attachment between workers and firm lasts longer.

age-productivity profile. Furthermore, there seems to be a decrease of the productivity of unskilled workers after 55. [Ilmakunnas and Maliranta \(2005\)](#) estimate the usual production functions and wage equations, but also examines how the results are influenced by using imperfect employment indicators like number of employees instead of hours of work. The increasing pattern of the wage-productivity gap by age is among the most robust results, although the productivity effect of age varies somewhat from model to model. They attribute the positive correlation between age and the wage-productivity gap to strong seniority effects in wage setting. [Dostie \(2011\)](#) uses a method similar to [Hellerstein et al. \(1999\)](#), but improves by controlling for unobserved productivity shocks. On the basis of an analysis of Canadian linked worker firm data, [Dostie \(2011\)](#) concludes that both wage and productivity profiles are concave, but productivity is diminishing faster than wages for workers aged 55 and over.³

[Aubert and Crépon \(2007\)](#) find that productivity, defined as the average contribution of particular age groups to the productivity of firms, increases with age until age 40–45 and then remains stable after this age. The results are stable across industries. They also show that the age-productivity profile is similar to the age-labor cost profile which is contradicting the overpayment of older workers. The evidence for what happens after 55 remains inconclusive due to data and precision issues. [Ilmakunnas and Maliranta \(2007\)](#) examine the connection of aging work force to firm performance by using information on the hiring and separation of employees. They show that firm-level labor productivity change can be decomposed to the effects of the hiring and separation rates of the age groups and to the effect of productivity growth of those workers in different age groups who are staying in the firm. The evidence shows that separations of older workers are profitable to firms, especially in the manufacturing ICT-industries, because there are indeed differences between the age groups in their relative productivity and wage levels.

[Malmberg et al. \(2008\)](#) aim to answer the empirical question of how labor productivity at the plant level is related to the age composition of the labor force without imposing any given theoretical structure. The data comes from the Swedish Manufacturing and Mining Surveys and Regional Labor Market Statistics and consists of 8,000–9,000 establishments each year over the period 1985–1996. According to their estimates, an accumulation of high shares of older adults in manufacturing plant does not seem to have a negative effect on plant-level productivity. On the contrary, when plant-level effects are controlled for, high shares of older adults are associated with higher productivity than high shares of young adults.

[Göbel and Zwick \(2009\)](#) analyzing German data find considerable differences in parameter estimates depending on the estimation strategy. The results of their preferred estimates depict an increase in productivity until the age 40–45, but no meaningful decline until the age of 60. This is different from existing studies that find inverse U-shaped age productivity profiles. In [Garibaldi et al. \(2010\)](#) a study on Finnish manufacturing is presented which compares the productivity and wage profiles by age using average age (and its square) as the measure of plant age structure. There is weak

³ He indicates that it is only possible to distinguish workers by age, gender and education. An important missing variable is occupation so it is not possible to distinguish workers in managerial positions from workers in production positions.

evidence that in the early 1990s, the industry-wide age-productivity profile was hump-shaped. On the other hand, in the late 1990s and early 2000s the profile was declining with age. Irrespective of what is the shape of the age-productivity profile, wages tend to increase with age or tenure. They also present results on the role of age diversity in Finnish industrial plants; age diversity is positively related to productivity.

3.3 Conclusions from Previous Studies

Studies on the relationship between age and productivity using matched worker-firm data find varying results. Some studies find evidence of a productivity-wage gap that is increasing with age, while other studies find no evidence for this phenomenon. Cross-sectional studies are limited also because they cannot take into account that there may be unobserved differences between firm productivity and firm workforce composition. Longitudinal matched-worker firm data provide a rich source of information about potential relationships between age and productivity. But even in this type of studies conclusions are not clear-cut. Although most studies find that when age increases productivity does not increase as much as wages do, there are also some studies that find no such an effect.

4 Data

4.1 Combining Datasets

We use microdata collected by Statistics Netherlands (CBS) to match information about individual workers with information about the firm at which they are employed. The matched dataset combines information from four different sources and six different years.⁴ Our dataset provides information at the firm level about value-added, employment, wage costs, firm and industry characteristics and about the composition of the workforce. The dataset contains firms from manufacturing and covers the period 2000–2005.

After cleaning, the dataset consists of 13,941 firms with 44,371 observations. The composition of the workforce in each firm is computed by using personal information about age and gender. For each firm eight age groups (less than 25, 25–29, 30–34, 35–39, 40–44, 45–49, 50–56 and 57 and older) and two gender groups (males and females) are considered. Employment is measured in two ways: first in full-time equivalents determined through the number of social insurance days,⁵ and second by the number of workers. In our baseline estimates we use the number of full-time equivalents, in a sensitivity analysis we use the number of persons working during the year.

⁴ The files from different sources can be combined using unique identifier for each firm and each worker. Appendix A presents a more detailed description of our data.

⁵ The full-time equivalence scale is determined by the number of days a person has social insurance; if an individual works 50% of the maximal number of days for which an individual can be social insured this individual counts for 0.5 full-time equivalent.

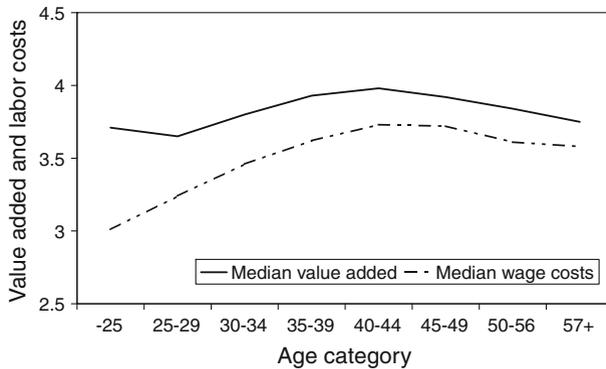


Fig. 1 Median value-added and median labor costs by age group; log-scale. Note that this graph is based on observations of the mean age of the workforce in each firm in each year of our dataset

Table 1 presents summary statistics for the final set of plant-level data. The average value added is equal to €5.9 million. The labor costs are on average €3.5 million and the depreciation on fixed assets are on average equal to €827,000. The average employment consists of 97 persons and which is equal to 77 full-time equivalent workers. According to the middle part of the table, most firms fall in the category 20–49 employees. The lower part of Table 1 presents the mean and standard deviation of the shares of age groups on the basis of number of employees and number of full-time equivalent workers. The youngest category, less than 25 years old, has a larger share in the number of employees than in the number of full-time equivalent workers, which indicates that younger persons work a shorter time at a firm compared to the other categories.

Figure 1 gives a graphical representation of the relationship between age, wage and productivity in our sample. Firms are characterized by the average age of their workforce in a particular year and allocated to one of the eight age categories. Then for every age category median value added per worker and median labor costs per worker are calculated and plotted. As shown in Fig. 1 value added is not much affected by the average age of the work force. Labor costs increase at lower ages and are roughly constant from age 40 onwards. Apparently, as the average age of the work force increases the difference between value added and wage costs is smaller. Of course, we cannot derive any causal conclusion from Fig. 1 as across age groups different firms are compared.⁶

⁶ Also note that for some age intervals only few observations are available; for the –25 interval there are 261 observations, for the 50–56 interval there are 189 observations and for the 57+ interval there are only 9 observations. The highest number of observations is for the 35–39 interval for which we have 19,100 observations.

on firm-level micro survey data. The estimated effect of capital stock on production in [Dostie \(2011\)](#) is close to zero in all specifications. Also [Aubert and Crépon \(2007\)](#) and [Hellerstein et al. \(1999\)](#) report small productivity effects from the capital stocks. In our baseline estimate we ignore capital input. Nevertheless, to get some idea about the effects of ignoring capital input we perform a sensitivity analysis in which as a rough indicator for the capital stock we use depreciation on fixed assets.

Another weakness of the CBS data is the lack of information on educational attainment of workers. To provide an idea about how this might affect our parameter estimates we discuss the results of similar studies which did have information about education, skills or experience of their employees. According to [Haegeland and Klette \(1999\)](#) there is a positive premium for all education levels for both wage and productivity. Wage returns to education correspond quite well to productivity differences. Also [Ilmakunnas and Maliranta \(2005\)](#) find that higher education has in general a positive influence on productivity with a stronger effect for non-technical education than for technical education. However, the wage and productivity effects of education are not monotonous when education increases, i.e. productivity increases are ‘under-compensated’. [Hellerstein and Neumark \(2004\)](#) find a similar results that productivity premium for education somewhat exceeds the wage premium. [Dostie \(2011\)](#) on the other hand finds that older men with a degree earn higher wages relative to their productivity, while the productivity of workers aged 55 or more without a degree seems to be slightly higher than their wage. For women with a degree the hypothesis that wages are equal to productivity can not be rejected. [Crépon et al. \(2003\)](#) and [Aubert \(2003\)](#) include skill categories to distinguish between different types of workers. According to [Crépon et al. \(2003\)](#) productivity is much higher for older workers if skill categories are excluded from the analysis. [Aubert \(2003\)](#), with the same data but over a longer period, shows that skilled workers have the steepest age-productivity profile and their productivity increases at all ages. Because we cannot distinguish between skill categories we might underestimate the age-productivity profile, though not necessarily the age related productivity-wage gap since by the same line of reasoning we might also underestimate the age-wage profile. In order to get some idea about the importance of educational attainment of workers by way of sensitivity analysis in addition to manufacturing we also present estimates of other industries; construction, retail trade, wholesale trade and commercial services.

5 Set-Up of the Analysis

5.1 Model

We assume that the production per worker can be specified using the following Cobb–Douglas production function⁷

$$\ln(Q_{it}/L_{it}) = \ln(q_{it}) = \gamma \ln(\lambda_{it} L_{it}) + \delta F_{it} + \varepsilon_{it} \quad (1)$$

⁷ As indicated before, in our baseline estimate we ignore capital input. By way of sensitivity analysis we include a proxy for capital input.

where Q_{it} defines the production, measured as firm i 's value added in year t . Furthermore, λ is the productivity of an individual unit of labor and L_{it} the quantity of labor. Thus, $\lambda_{it}L_{it}$ denotes the level of productive labor. F_{it} contains workplace characteristics like industry, age (of the firm) as well as calendar year dummies.⁸

With the assumption of perfect substitution among workers of different ages, the level of productive labor λL can be written as

$$\lambda L = \sum_j \lambda_j L_j = \lambda_0 L \left(1 + \sum_{j-\{0\}} \left(\frac{\lambda_j}{\lambda_0} - 1 \right) \frac{L_j}{L} \right) \tag{2}$$

i.e. the sum of productive labor of each category of workers, where λ_j is age group j 's marginal productivity. So, we assume that workers of different age groups are perfect substitutes but they may have a different marginal productivity. The production function can now be written as

$$\ln(q_{it}) = \gamma \ln(\lambda_0) + \gamma \ln(L_{it}) + \gamma \ln \left(1 + \sum_{j-\{0\}} \left(\frac{\lambda_j}{\lambda_0} - 1 \right) \left(\frac{L_j}{L} \right)_{it} \right) + \delta F_{it} + \varepsilon_{it}$$

or, after linearization

$$\begin{aligned} \ln(q_{it}) &\cong \gamma \ln(\lambda_0) + \gamma \ln(L_{it}) + \sum_{j-\{0\}} \gamma \left(\frac{\lambda_j}{\lambda_0} - 1 \right) \left(\frac{L_j}{L} \right)_{it} + \delta F_{it} + \varepsilon_{it} \\ &= \alpha + \gamma \ln(L_{it}) + \sum_j \gamma_j l_{j,it} + \delta F_{it} + \varepsilon_{it} \end{aligned} \tag{3}$$

where $\alpha = \gamma \ln(\lambda_0)$, $\gamma_j = \gamma \left(\frac{\lambda_j}{\lambda_0} - 1 \right)$, and $l_{j,it} = \left(\frac{L_j}{L} \right)_{it}$. Relative labor costs of categories of workers can be estimated by applying the same approach as in the case of marginal productivities. The mean hourly labor cost can be rewritten as

$$\bar{w} = \frac{\sum_j w_j L_j}{\sum_j L_j} = w_0 \left(1 + \sum_{j-\{0\}} \left(\frac{w_j}{w_0} - 1 \right) \frac{L_j}{L} \right) \tag{4}$$

where we assume constant relative labor costs of groups across firms. To allow for a direct comparison with the parameter estimates of the production equation we introduce firm size and other workplace characteristics as explanatory variables of wage

⁸ Since many of the workplace characteristics are not time varying, in the baseline estimates we restrict the workplace characteristics to the share of women in the workforce. In all the estimates calendar year dummies are included as well. Note that without further restrictions one cannot disentangle age, calendar year and birth cohort year effects because of the identity: calendar year–birth year = age. Nevertheless, following Deaton and Paxson (1994), De Ree and Alessie (2010) show that under some assumptions birth cohort effects may be identified.

costs. Thus, similar to the production function the relative labor costs can be estimated by regressing the following equation

$$\ln(\bar{w}_{it}) = \gamma_w \ln(w_0) + \gamma_w \ln(L_{it}) + \ln \left(1 + \sum_{j \neq \{0\}} \left(\frac{w_j}{w_0} - 1 \right) \left(\frac{L_j}{L} \right)_{it} \right) + \delta_w F_{it} + \varepsilon_{w,it} \quad (5)$$

or, after linearization

$$\ln(\bar{w}_{it}) \cong \ln(w_0) + \gamma_w \ln(L_{it}) + \sum_{j \neq \{0\}} \left(\frac{w_j}{w_0} - 1 \right) \left(\frac{L_j}{L} \right)_{it} + \delta_w F_{it} + \varepsilon_{w,it} \quad (6)$$

$$= \alpha_w + \gamma_w \ln(L_{it}) + \sum_j \xi_j l_{j,it} + \delta_w F_{it} + \varepsilon_{w,it} \quad (7)$$

where $\alpha_w = \ln(w_0)$ and $\xi_j = \left(\frac{w_j}{w_0} - 1 \right)$. Equality of the age effects on labor cost and productivity can be tested by comparing the estimated coefficients of ξ_j and γ_j .

5.2 Methodology

Different methods are used to estimate the equations. First we present pooled cross-section estimates ignoring firm fixed effects ($\alpha = \alpha_i$ and $w = w_i$). If firm fixed effects are not accounted for the relationship between age composition and productivity is identified on cross-sectional variation. The interpretation of the between estimation results is that an age group is estimated to be more productive than another group if a firm with a higher share of this age group in its labor force on average produces more than a comparable firm with a lower share for this age group.

Second, we introduce firm fixed effects. Then, the relationship is identified as an average over within firm variation. The results from the fixed effects estimation can be interpreted as follows: a group is estimated to be more productive than another group if, in comparable firms, production increases more (or less) on average when the share of labor of the group increases more (or less) than the share of the other group. We remove the fixed effects by estimating the equations in first differences.⁹

Although the introduction of firm fixed effect removes potential spurious correlation between age composition and productivity it does not solve all problems. The problem is that changes in age composition may not be exogenous to changes in productivity. It could be that there is a negative productivity shock which induces firm to fire young workers, causing the average age of the workforce to increase. Then, the negative productivity shock seems to be due to the increase in average age of the workforce while in fact there is an exogenous explanation for this correlation. To address the potential endogeneity bias, we use an instrumental variable approach. It is difficult to find variables that can serve as valid instruments, i.e. variables that have

⁹ Note that to the extent that F_{it} contains fixed firm effects, these drop out of the equation in first differences.

an effect on the endogenous variable but do not directly affect productivity. We estimate the production and labor cost functions in first differences using lagged values of the age structure as instruments for the change in age structure. The underlying assumption is that shocks occurring between $(t - 1)$ and t are uncorrelated with levels of inputs earlier than $(t - 2)$. Thus, when accounting for potential endogeneity of the age structure, the following equation is estimated as our baseline model:

$$\Delta \ln(q_{it}) = \gamma \Delta \ln(L_{it}) + \sum_j \gamma_j \Delta l_{jit} + \delta \Delta F_{it} + \Delta \varepsilon_{it} \quad (8)$$

where we use l_{jit-s} ($s \geq 2$) as instrumental variables for Δl_{jit} . The wage equation is estimated similarly:

$$\Delta \ln(\bar{w}_{it}) = \gamma_w \Delta \ln(L_{it}) + \sum_j \xi_j \Delta l_{jit} + \delta_w \Delta F_{it} + \Delta \varepsilon_{w,it} \quad (9)$$

To account for the presence of heteroskedasticity we estimate the relevant parameters of our model using the General Method of Moments (GMM).¹⁰ We present two sets of parameter estimates. First, we instrument the potential endogenous first differenced age structure variables with the two periods lagged levels of the age structure. Second, we also use the three periods lagged levels of the age structure as additional instrumental variables. Over the order of the estimation methods the number of observations gets increasingly smaller. For the pooled time series—cross section estimates we have 44,371 observations, for the first differences estimates there are 28,775 observations. Introducing lagged instrumental variables reduces the sample size even more, to 24,509 for GMM-1 estimates and 16,689 for GMM-2 estimates.

To assess the credibility of our approach we perform a range of diagnostic tests. If the correlation between the instrumental variables and the endogenous variables is poor, i.e. if we have “weak” instruments our parameter estimates may be biased. To investigate the quality of our instruments we use two Kleibergen-Paap statistics, an underidentification test statistic and a weak instruments test statistic. The first test statistic deals with the issue of whether the equation is identified, i.e. the excluded instruments are all relevant. The second statistic is a generalization of the first stage F -statistic. The null hypothesis is that the instruments are weak. Rejection of the null hypothesis implies that the instruments pass the weak instruments test, i.e. they are highly correlated with the endogenous variables.¹¹ To investigate the validity of our instrument we use the Hansen–Sargan test—also called Hansen’s J test—of overidentifying restrictions. The null hypothesis is that the instruments are valid instruments,

¹⁰ Initially we performed 2SLS estimates. However, using the Pagan-Hall statistic we also found evidence of the presence of heterogeneity. Therefore we use GMM, as in the presence of heteroskedasticity the GMM estimator is more efficient than an IV estimator. Nevertheless, the differences between the 2SLS and GMM estimates are very small. Our estimates are based on Stata’s “ivreg2” suite, with the “gmm2” option (see Baum et al. 2003, 2007).

¹¹ An exact rejection rule for weak identification is not yet established. Baum et al. (2007) suggest to use the old “rule of thumb” rule that the F -statistic should be at least 10 for weak identification not to be considered a problem.

i.e., uncorrelated with the error term, and that the instruments are correctly excluded from the estimated equation. The null hypothesis of the test is that the overidentifying restrictions are valid. Under the null, the test statistic is distributed as chi-squared in the number of overidentifying restrictions. A failure to reject the null hypothesis implies that the instruments are exogenous.¹² In virtually all our GMM estimates our instruments pass the weak instruments test while in those estimates where we can test overidentifying restrictions we cannot reject the null hypothesis that these restrictions are valid. Finally, we test whether indeed the age structure is endogenous and we need an instrumental variable approach. For this we use a Hausman test in which OLS and IV estimates are compared. The test statistic is χ^2 distributed with degrees of freedom equal to the number of regressors tested. The null-hypothesis is that the specified age structure can actually be treated as exogenous. Rejection of the null-hypothesis indicates that the age structure is endogenous.

In most of the estimates we find support for the age structure being endogenous, thus supporting the instrumental variable approach. This all being the case unless there is a reason to do so we do not comment on the diagnostic tests but focus on the relevant parameter estimates.

6 Parameter Estimates

6.1 Productivity and Labor Costs

This section presents parameter estimates of the baseline model, which has no capital stock included, and has employment as one of the right-hand side variables. The dependent variable in the productivity estimates is the value added per worker at the firm level. As indicated before, we use unbalanced panel data of 13,941 firms with 44,371 observations.

The top part of Table 2 presents the parameter estimates of the productivity and labor cost functions estimated using Ordinary Least Squares (OLS) and first differences to remove plant-level fixed effects. The first two columns present the estimates of the production function. The OLS parameter estimates indicate that larger firms are more productive. Furthermore, productivity is (inverted) U-shaped with age, having a maximum level between age 35 and 39: all the estimated coefficients of the age share groups are negative. To give some idea about the magnitude of the estimates, consider Table 2, first column, age 57+: the parameter has a value of -0.30 . This means that if the share of 57+ workers within a firm would increase with 1%-point, productivity declines with $0.01 \cdot 0.30 = 0.003 = 0.3\%$. So roughly the estimates can be interpreted as age-productivity elasticities.¹³ Of course, the increase in the share of one age group influences the shares of the other age groups. A higher share of females

¹² In the first approach the number of instruments equals the number of potentially endogenous variables so we cannot test overidentifying restrictions.

¹³ Note that for example an estimate of -1 for a particular age category would imply that an increase in employment of workers from that age category with 1% point would reduce average productivity per worker with 1% so that there is no net gain in production.

has a negative effect on the productivity of the firm. This is related to many women working part-time.

Introducing firm fixed effects changes many parameter estimates. Within firms, an expansion of employment reduces productivity and labor costs. Apparently, new workers are less productive and earn lower wages than incumbent workers. Similarly, reducing employment increases productivity and labor costs, which is consistent with the last in–first out principle. Furthermore, the fixed effects estimates imply a flat age-productivity profile. Apparently, the positive age-wage and age-productivity profiles only hold in a cross-section. Firms with a younger workforce may pay lower wages and have a lower productivity than firms with an older workforce. Once differences between firms are taken into account, these profiles vanish. A change in the male-female composition of a firms' workforce would have no effect.

The GMM parameter estimates are shown in the bottom part of Table 2. Due to the use of instrumental variables the precision of the parameter estimates is much lower. Nevertheless the results are very different from the fixed effects results. Whereas in the fixed effects parameter estimates the productivity of older employees was estimated to decrease after a particular age, we now find that productivity increases with age. Only the oldest categories—57 years or older in the GMM-1 estimate and 50 years and older in the GMM-2 estimate are insignificantly different from zero, but still positive.¹⁴ Apparently, we cannot rule out that the age structure is endogenous causing a negative correlation between age structure and error terms of the estimated equation. Once we take this into account and instrument the age structure, this negative correlation is removed leading to a positive age-productivity profile. Apparently, a shock to productivity causes an endogenous change of the age composition. If firms are hit by negative shocks their productivity goes down and as a consequence, perhaps because of a last in–first out principle, they layoff younger workers thus increasing the average age of the workforce. Once this potential endogeneity of the age composition of the workforce is taken into account there is an upward sloping age-productivity profile.

The last two columns of Table 2 present the estimates of the labor cost equation. The dependent variable is the total labor cost divided by the total number of employees at the firm using the number of full-time equivalents. In the OLS estimates the labor costs are lower for young workers up to age 35. After that age does not affect wage costs. In the fixed effects estimates the age-wage cost profile is basically flat. A large female share of the labor force has a negative effect on the labor cost in the between estimation which almost disappears in the fixed effects estimation. The results from the GMM estimates for the labor cost function are similar to those for the production function in the sense that age has a positive effect. Nevertheless, in the GMM-estimates only the lowest age group has an effect that is different from zero at conventional levels of significance.

The OLS and fixed effects estimates show that the age effect on productivity and wage costs at young ages is very similar. The evidence from the GMM estimates is

¹⁴ Note that because of the lagged levels of explanatory variables, the number of observations is much smaller.

Table 2 Baseline parameter estimates of the production and labor cost function

	Production function		Labor cost function	
	OLS	First differences	OLS	First differences
Employment	0.02 (0.00)***	-0.72 (0.01)***	0.02 (0.00)***	-0.70 (0.01)***
Age share				
-25	-0.85 (0.04)***	0.03 (0.05)	-0.94 (0.03)***	-0.06 (0.04)
25-29	-0.35 (0.05)***	0.10 (0.06)	-0.28 (0.03)***	0.06 (0.04)
30-34	-0.12 (0.05)***	0.03 (0.05)	-0.09 (0.03)***	0.05 (0.04)
35-39	-	-	-	-
40-44	-0.09 (0.05)*	-0.06 (0.05)	-0.02 (0.04)	-0.07 (0.04)*
45-49	-0.27 (0.05)***	0.01 (0.06)	-0.04 (0.04)	-0.05 (0.04)
50-56	-0.37 (0.04)***	-0.08 (0.06)	0.04 (0.04)	-0.07 (0.04)
57+	-0.30 (0.05)***	-0.07 (0.07)	-0.05 (0.04)	-0.07 (0.05)
Female	-0.46 (0.02)***	0.00 (0.01)	-0.54 (0.01)***	-0.04 (0.03)
Observations	44371	28775	44371	28775
	GMM-1	GMM-2	GMM-1	GMM-2
Employment	-0.72 (0.02)***	-0.72 (0.03)***	-0.68 (0.02)***	-0.68 (0.02)***
Age share				
-25	-0.58 (0.33)*	-0.29 (0.38)	-0.50 (0.24)**	-0.76 (0.28)***
25-29	-0.15 (0.25)	-0.09 (0.32)	-0.02 (0.19)	-0.15 (0.25)
30-34	-0.10 (0.19)	0.13 (0.22)	-0.06 (0.15)	0.12 (0.17)
35-39	-	-	-	-
40-44	0.16 (0.21)	0.26 (0.23)	0.25 (0.16)	0.17 (0.18)
45-49	0.49 (0.28)*	0.72 (0.32)**	0.40 (0.20)**	0.28 (0.23)
50-56	0.60 (0.36)*	0.57 (0.40)	0.46 (0.26)*	0.23 (0.30)
57+	0.09 (0.49)	0.27 (0.52)	0.41 (0.34)	-0.18 (0.40)
Female	-0.01 (0.04)	0.09 (0.06)	-0.06 (0.03)*	-0.03 (0.05)
Observations	24509	16689	24509	16689
Diagnostic tests				
Overidentification	-	2.77	-	3.14
Underidentification	162.2***	144.2***	162.2***	144.2***
Weak identification	23.6	10.7	23.6	10.7
Endogeneity	14.5**	9.7	13.5*	19.2***

All estimates include year dummies. The GMM estimates are in first differences with lagged levels as instruments; GMM-1 has two period lagged levels of variables as instruments; GMM-2 has two and three period lagged levels of variables as instruments; The following diagnostic tests are presented: overidentification = Hansen-Sargan statistic; underidentification = Kleibergen-Paap rk LM statistic; weak identification = Kleibergen-Paap Wald F statistic; endogeneity = Durbin-Wu-Hausman statistic; in parentheses robust standard errors; *** (**, *) indicates significance at a 1% (5%, 10%) level

Table 3 Parameter estimates of the productivity-wage gap

	First differences	GMM-1	GMM-2
Employment	-0.03 (0.01)***	-0.04 (0.02)**	-0.04 (0.02)**
Age share			
-25	0.09 (0.04)**	-0.07 (0.26)	0.37 (0.30)
25-29	0.04 (0.04)	0.13 (0.20)	-0.01 (0.24)
30-34	-0.03 (0.04)	-0.04 (0.16)	-0.00 (0.17)
35-39	-	-	-
40-44	0.01 (0.04)	-0.09 (0.16)	0.09 (0.17)
45-49	0.06 (0.04)	0.09 (0.22)	0.37 (0.24)
50-56	-0.00 (0.1)	0.14 (0.27)	0.26 (0.9)
57+	0.00 (0.05)	-0.31 (0.37)	0.40 (1.1)
Female	0.04 (0.03)	0.05 (0.04)	0.10 (0.04)**
Observations	28775	24509	16689
Diagnostic tests			
Overidentification	-	-	5.29
Underidentification	-	162.2	144.2
Weak identification	-	23.6	10.7
Endogeneity	-	5.5	3.8

The productivity-wage gap is defined as $\ln(q_{it}) - \ln(\bar{w}_{it})$; see also footnote Table 2

somewhat mixed. For the GMM-1 estimates the pattern of age dependence is similar for age and productivity except for the share of workers beyond age 57 where productivity goes down much faster than wage costs. For the GMM-2 estimates relative to the reference age group, productivity beyond age 40 is higher than wage costs. However, due to the imprecision of the parameter estimates the two profiles are not significantly different from each other.

Our main conclusions from these baseline parameter estimates are the following. The pooled time series—cross section estimates show that there are upwards sloping age-productivity and age-wage costs profiles. However, these upward slopes disappear when we introduce firm fixed effects. This indicates that the pooled estimates are biased. The estimated slope of both profiles depends on the estimation strategy, i.e. whether or not we account for potential endogeneity. When we do this both profiles are again upward sloping. However, using instrumental variables increases the imprecision of the parameter estimates and the two profiles are not significantly different from each other.

6.2 Productivity-Pay Gap

As indicated before, equality between the age effects on relative labor cost and relative productivity can be tested by comparing the estimated coefficients of ξ_j and γ_j . We can do this directly by taking first differences of Eqs. (3) and (7) and estimate:

$$\Delta \ln(q_{it}) - \Delta \ln(\bar{w}_{it}) = \text{Gap}_{it} = \alpha_g + \gamma_g \Delta \ln(L_{it}) + \sum_j \gamma_{g,j} \Delta l_{jit} + \delta_g \Delta F_{it} + \Delta \varepsilon_{g,it} \quad (10)$$

where $\alpha_g = \alpha - \alpha_w$, $\gamma_g = \gamma - \gamma_w$, $\gamma_{g,j} = \gamma_j - \xi_j$, $\delta_{g,f} = \delta_f - \delta_{w,f}$ and $\varepsilon_{g,it} = \varepsilon_{it} - \varepsilon_{w,it}$.

Table 3 shows the parameter estimates for Eq. (10). Employment has a negative effect on the productivity-pay gap, i.e. if employment increases the gap becomes smaller. If we ignore potential endogeneity of the age shares, a higher share of workers below age 25 increases the gap between productivity and wage costs. None of the other age shares has a significant effect. If we account for potential endogeneity none of the age shares has a significant effect either.

For the estimates of the pay-productivity gap we cannot reject the null-hypothesis of the age structure being exogenous. Apparently, an exogenous productivity shock may affect both productivity and the age structure causing younger workers to be laid-off and introducing a spurious negative relation between age and productivity. However, the negative productivity shock also introduces a spurious negative relation between age and wage costs as young workers usually have a lower wage. If an exogenous shock affects both the change in productivity and the change in wage costs, the difference between the two remains unaffected other than through changes in the age structure.

6.3 Sensitivity Analysis

To illustrate the sensitivity of our parameter estimates we did some sensitivity analysis for manufacturing firms.¹⁵ Our first sensitivity analysis address the issue of whether or not it is important to use balanced data. Our main analysis is based on unbalanced data on all firms in our sample. By way of sensitivity analysis we present the parameter estimates based on balanced panel data, consisting of firms surveyed in all years. This subset consists of 2,944 firms. In the balanced dataset on average the firms are larger and the value added, labor costs and depreciation are higher. The average number of employees per firm is doubled compared to the unbalanced panel data set. Parameter estimates based on the balanced panel show that the differences with the full dataset are minor.

As an alternative to distinguishing 8 age groups, we also tried an alternative estimation in which we use average age of the workforce as explanatory variable. In this case we respecify in Eq. (1) $\lambda_{it} = \exp(\lambda_a A_{it})$, with A is the average age of the workforce. In this case an additional year of age affects output per worker always in the same proportion. We estimate

$$\ln(q_{i,t}) = \alpha + \gamma \ln(L_{i,t}) + \gamma_a A_{i,t} + \delta F_{i,t} + \varepsilon_{i,t} \quad (11)$$

¹⁵ Van Ours and Stoeldraijer (2011) provide these sensitivity analyses in detail.

with $\gamma_a = \gamma\lambda_a$. The wage cost equation is specified accordingly. The signs of the estimated coefficients of the mean age are the same for the production function and the labor cost specification. The results are in line with the baseline estimates. For the pooled cross-section we find that age has a positive effect on productivity. Once we introduce firm fixed effects the age effect becomes negative. Again, in the GMM estimates there is a positive age effect for both productivity and wage costs, but again, given the imprecision of the estimates they are not significantly different from each other.

Furthermore, by way of sensitivity analysis we also included capital as an explanatory variable. Some studies include capital input as an explanatory variable in a production equation while other studies do not. Including capital input is driven by data availability. Our data set lacks a proper variable about capital, but the depreciation on fixed assets is used as a rough indicator for capital input. For reasons of symmetry we also introduce capital in the wage costs function. Although there are differences in the magnitude of some parameter estimates the general pattern of the age effects on productivity and wage costs is very much the same.

In the baseline model, the percentages of the age and gender groups were calculated by using the number of full-time equivalent workers in the firm, because that gives a better representation of the actual time a worker worked at the firm. Another possibility is using the number of workers the firm employed during the year. This measure is less precise but often used in other studies. The use of the number of workers does not change the estimated parameters for the middle groups, but for the youngest age group they are somewhat larger. This may be caused by the fact that younger stay at a firm for a shorter period than other workers (see, for instance, the proportions of the age groups in the descriptive statistics of Table 1). Nevertheless, using a different specification for employment does not influence the age patterns a lot.

In addition to an analysis of manufacturing firms, we applied our baseline approach to data from other industries. Our analysis for other industries provided some idea about potential effects of differences in educational attainment. There are substantial differences in educational attainment across industries. At the low end of the educational distribution, whereas 13.4% of the workers in manufacturing have basic education this is about 10% in construction and commercial services. At the high end of the educational distribution, whereas 18.2% of the workers in manufacturing have a higher vocational or academic degree, this is 4.5% in construction and 31.4% in commercial services. We performed additional analyses for construction, wholesale trade, retail and commercial services. There are clear differences in the age-productivity patterns and the age-wage costs pattern within industries for different estimation procedures and between industries for the same estimation procedure. An important difference compared to the manufacturing estimates is that for the other industries in many cases we cannot reject the exogeneity of the age structure. Therefore, for the other industries we prefer the first difference estimates, which show flat age profiles for productivity and wage costs. We can only speculate why exogeneity of the age structure cannot be rejected. Compared to manufacturing the other industries have more small firms, whereas retail trade and commercial services have a higher share of young workers and higher shares of female workers. Nevertheless, also for the

other industries the productivity and wage costs profiles follow each-other closely, confirming the general results for the manufacturing sector.

7 Conclusions

Previous empirical studies on the effect of age on productivity and wages find contradicting results. Some studies find that if workers grow older there is an increasing gap between productivity and wages. Wages increase with age while productivity does not or does not increase at the same pace. However, other studies find no evidence of such an age related pay-productivity gap. We perform an analysis of the relationship between age, wage and productivity using a matched worker-firm panel dataset from Dutch manufacturing covering the period 2000–2005.

Although previous studies have their own peculiarities in terms of data, specification of dependent and independent variables and estimation methods used, we draw some comparisons of our results with those obtained in other studies. Although some studies that used matched worker-firm data find evidence of a wage-productivity gap at higher ages there are also studies which do not find such a gap. In terms of main findings our study is most closely related to [Aubert and Crépon \(2007\)](#) and [Dostie \(2011\)](#). [Aubert and Crépon \(2007\)](#) observe an age-labor cost profile which is similar to the age-productivity profile and hence they do not find a wage productivity gap. [Dostie \(2011\)](#) finds that wage and productivity profiles are concave, at their highest for the 35–55 age group and diminishing afterwards. He cannot reject the hypothesis that across age categories wages are equal to productivity.

Many previous studies have their limitations because they are based on cross-sectional information only or because they do not account for potential endogeneity of the age composition of a firms' workplace. Our study has its limitations too. For example, we only have a imperfect measure for the capital stock. Furthermore, our data do not contain information on educational attainment of workers, which complicates the analysis because older cohorts have less formal education than younger cohorts. To the extent that birth cohorts affect productivity independent of age, our parameter estimates may be biased. Nevertheless, However, a range of sensitivity analyses indicate the our parameter estimates are quite robust to the inclusion of capital or the application of our estimation procedures to other industries with a different skill distribution.

Our method of analysis is focused on firm-level productivity, identifying relationships on the basis of changes in age composition and changes in firm-level productivity. In a steady state situation we could never establish the relationship between age and productivity because level differences are absorbed by the firm fixed effects. Establishing a relationship between a composite age variable and an firm aggregate such as average productivity is very indirect. Nevertheless, except for special situations—piece rate work, homogeneous tasks, easy to monitor activities—there is no clear alternative to the approach used in this paper. Measurement errors may dominate the relationship thus underestimating the true effect of age on productivity. However, using instrumental variables should account for at least part of the measurement errors. Indeed in some of our estimates using instrumental variables turns a flat age-productivity profile into

an upward sloping one. Still, this upward sloping profile is almost identical for wage costs, so that we cannot reject the hypothesis that the two profiles are insignificantly different from each-other. There is also an issue of selective attrition. It could be that the best 50+ workers remain while the least productive older workers leave the firms and perhaps even the labor force. Although this could be an explanation for the lack of negative productivity effects for older workers, this selectivity effect would also apply to the wage cost profile in a similar way. Finally, our data cover a short time period. Within these limitations our findings are clear.

Using a variety of estimation methods we find some evidence of age affecting productivity. In particular when we perform a pooled cross-section time series estimate we find that workers between age 30 and 45 have the highest productivity, while the productivity of younger and older workers is lower. The lowest productivity concerns workers below age 25. So, there is clear hump-shape relationship between age and productivity. However, there is also a hump-shape relationship between age and labor costs. Nevertheless for young workers productivity is higher than labor costs while for older workers it is the reverse: productivity is lower than labor costs. These clear-cut results disappear when we introduce firm fixed effects. By doing this we remove spurious correlation between firm characteristics, productivity and age composition. Then, we relate changes in the age composition to changes in productivity and labor costs. Now the age profile of productivity is rather flat until age 50, after which productivity declines. Wage costs are flat over the entire age range. Again, this suggests that older workers are relatively overpaid. Finally, we performed estimates in which we account for potential endogeneity of the change in the age composition. If for example a firm is hit by an exogenous productivity shock and due to this shock younger workers are fired, firm productivity goes down while at the same time the share of older workers goes up. Then, we might wrongly conclude that the aging workforce was responsible for the drop in productivity while in fact there was just correlation between the drop in productivity and the increase in the share of older workers. In these final estimates both productivity and wage costs increase with age, but we cannot reject that age has the same effect on productivity and wage costs. Clearly, productivity of workers will go down eventually at higher age. However, it is not clear where the turning point of the age-productivity profile is located.

Our main conclusion is that there is no evidence of an age-related productivity-pay gap. In addition to institutional reasons and related costs which explain why employers are reluctant to dismiss older workers, there is an additional rational. If the productivity-pay gap is absent or small, there is no urgent need to get rid of older workers. Once older workers lose their jobs, they are in a different position. It might well be that their productivity in a new job is below the pay they are supposed to earn which makes employers reluctant to hire older workers.

Appendix A: Details About Our Data

We combine four micro datasets from Statistics Netherlands using unique identifiers to match information of individual Workers to information about the firm who employs them: the General Firm Register (ABR), Production Statistics on Manufacturing

(PS—Ind), the Social Statistical Dataset on Jobs (SSB—Banen) and the Municipal Administration (GBA). All sources cover the period 2000–2005. The General Firm Register contains all firms in the Netherlands and serves as a sampling frame for the survey of firms. Information about the registration date of the firm is coupled to the information in the Production Statistics of Manufacturing. The Production Statistics on Manufacturing provide information to estimate the production and labor cost functions: value-added, employment and total labor costs. Manufacturing firms are defined as firms whereby mechanical, physical or chemical processing of materials, substances or components, new products are fabricated. The Production Statistics refer to a sample of firms. In total there are 21,685 different firms over the six years with an observation number of 61,562. Of these, 3,777 firms are observed for all years. The Social Statistical Dataset is an exhaustive database which registers all job-worker relations of a person in a continuous period. The dataset contains information on the number of working days and fiscal wage. The information is mainly coming from the Insured Workers Administration (VZA), but has also more detailed information from surveys. The surveys cover only 30% of all records in the data set and are therefore not used. Every record is identified by the persons identification number (based on the social security number), the firm's identification number and an identification code. The firm's identification code makes it possible to link this file to the information on the firms.

The identification code indicates if the person is present in the Municipal Administration. This data source provides more detailed information on the persons in the Social Statistical Dataset on Jobs: year of birth, gender, nationality, etc. However, information on schooling or occupation is lacking. With this information the characteristics of groups of workers can be easily computed from the aggregation of information on all workers with similar characteristics. For each firm, eight age categories are considered: less than 25, 25–29, 30–34, 35–39, 40–44, 45–49, 50–56 and 57 and older. The category 35–39 years old is taken as the reference group in our estimates. Labor input can be measured in two ways: first by the persons worked at the firm during the year and second by the number of full-time equivalents. It is possible that a person which is present in the Social Statistical Dataset on Jobs cannot be found in the Municipal Administration. One reason for this is that the person does not live in the Netherlands. The Municipal Administration gives information about the gender and year of birth and if this information is lacking, the person can not be arranged in the corresponding group which is used in the model. About 19% of the observations of the firms have at least one employee of which the personal information is lacking. These firms are not only larger, but also the value-added per employee is larger compared to firms which have all the information. To solve this problem, only firms with more than 95% of the employees personal information, are taken into account.

Table 4 presents summary statistics for plant-level data. There are 61,652 observations of value added and 1,366 of these observations are non positive. There are 2,133 firms with a negative labor cost. The average employment consists of 88.2 persons equal to 69.9 full-time equivalent workers. The lower part of Table 4 presents the mean and standard deviation of the shares of age groups on the basis of number of employees and number of full-time equivalents. The youngest category, less than 25 years old, has a larger share in the number of employees than in the number of

Table 4 Descriptive statistics of the complete dataset of manufacturing firms

	Obs.	Mean	SD	Min	Max
Value added	61562	5082	36455	-1926903	2693205
Log value added	60286	7.06	1.66	0	14.81
Total cost of labor	62556	3015	18738	-390	1649406
Log total cost of labor	59423	6.76	1.60	0	14.32
Depreciation	61554	729	6654	1	41648
Log depreciation	57607	4.75	1.81	0	13.14
Employment (NR)	55510	88.2	452.8	1	41648
Employment (FTE)	55510	69.9	373.5	0.0	34662.9
Establishment size (NR) %:					
1 employees	0.02	50–99 employees		0.19	
2–4 employees	0.06	100–199 employees		0.11	
5–9 employees	0.08	200–499 employees		0.06	
10–19 employees	0.13	500+ employees		0.02	
20–49 employees	0.34	Total		1.00	
	NR		FTE		
	Mean	SD	Mean	SD	
Age categories:					
–25	0.16	0.15	0.12	0.13	
25–29	0.11	0.10	0.11	0.11	
30–34	0.14	0.10	0.15	0.11	
35–39	0.15	0.10	0.15	0.11	
40–44	0.13	0.10	0.14	0.11	
45–49	0.11	0.09	0.11	0.10	
50–56	0.12	0.11	0.13	0.12	
57+	0.07	0.10	0.08	0.10	
No information	0.01	0.04	0.01	0.04	
Total	1.00	1.00	1.00	1.00	
Proportions female					
	0.26	0.20	0.24	0.20	

Value added, labor costs and depreciation in 1000 €; NR = number of workers; FTE = number of full-time equivalents

full-time equivalents, which indicates that younger persons work a shorter time at a firm compared to the other categories. The group of which there was no information is on average a little less than 1%.

In the cleaning process we lost observations for a variety of reasons. For some firms we have no information about value added, total labor costs, depreciation of fixed assets or the number of employees. There are also firms with negative values for value added, total labor costs or depreciation of fixed assets. And there are firms for which for more than 5% of the workers information about age or gender is missing.

Further cleaning of the dataset was done because of extreme changes in the number of employees in subsequent years. If the number of employees grew more than ten times as large or was less than 10% of the previous year, the firm was excluded from the sample. Also, firms which had a change in their main industry, based on the ISIC code, were excluded. Finally, we removed firms with less than 5 employees. The table below shows the transition from the gross sample to the net sample in more detail:

	Observations	Firms
Gross sample	61,562	21,684
After removing		
Value added (VA) negative	60,286	21,164
Labor costs (LC) negative	58,902	20,225
Depreciation (D) negative	56,019	19,368
Missing info on VA, LC, D	52,027	17,154
Missing info on more than 5% of workers	49,602	16,614
Big shifts in employment	48,974	16,435
Change of industry code	47,984	16,226
Firms with less than 5 workers	44,371	13,941

Note that when we remove firms for which there is no information over each of the years 2000–2005 a balanced sample of 2,944 firms with 17,664 observations remains

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

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., & Crépon, B. (2007). *Are older workers less productive? Firm-level evidence on age-productivity and age-wage profiles*. mimeo.
- Barth, M. C., McNaught, W., & Rizzi, P. (1993). Corporations and the aging workforce. In P. H. Mirvis (Ed.), *Building the competitive workforce: Investing in human capital for corporate success*. New York: Wiley.
- Baum, C. F., Schaffer, M. E., & Stillman, S. (2003). Instrumental variables and GMM: Estimation and testing. *Stata Journal*, 3, 1–31.
- Baum, C. F., Schaffer, M. E., & Stillman, S. (2007). Enhanced routines for instrumental variables/GMM estimation and testing. *Stata Journal*, 7, 465–506.
- Crépon, B., Deniau, N., & Pérez-Duarte, S. (2003). *Wages, productivity, and worker characteristics: a French perspective*. Working Paper CREST-INSEE, Paris.
- Deaton, A., & Paxson, C. (1994). Saving, growth, and aging in Taiwan. In D. Wise (Ed.), *Studies in the economics of aging* (pp. 331–357). Chicago: University of Chicago Press.
- De Ree, J., & Alessie, R. (2010). *Life satisfaction and age: Dealing with under-identification in age-period-cohort models*. Discussion paper NETSPAR,10/2010-056.
- Dostie, B. (2011) Wages, productivity and aging. *De Economist*, 159. doi:10.1007/s10645-011-9166-5.
- Garibaldi, P., Oliveira Martins, J., & van Ours, J.C. (2010). *Health, longevity and productivity: the economics of increased life expectancy*. Oxford: Oxford University Press.
- Göbel, C., & Zwick, T. (2009). *Age and productivity—Evidence from linked employer employee data*. Working Paper.
- Haegeland, T., & Klette, T. J. (1999). Do higher wages reflect higher productivity? Education, gender and experience premiums in a matched plant-worker data set. In J. Haltiwanger, J. Lane, J.

- R. Spletzer, J. Theeuwes, & K. Troske (Eds.), *The creation and analysis of employer-employee matched data*. Amsterdam: North-Holland.
- Hellerstein, J. K., David, N., & Troske, K. R. (1999). Wages, productivity and worker characteristics: Evidence from plant level production function and wage equations. *Journal of Labor Economics*, 17, 409–446.
- Hellerstein, J. K., & Neumark, D. (2004). *Production function and wage equation estimation with heterogeneous labor: Evidence from a new matched employer-employee data set*. NBER Working Paper Series, 13, 345–371.
- Ilmakunnas, P., & Maliranta, M. (2005). Technology, worker characteristics, and wage-productivity gaps. *Oxford Bulletin of Economics and Statistics*, 67, 623–645.
- Ilmakunnas, P., & Maliranta, M. (2007). *Aging, labor turnover and firm performance*. Discussion Paper 1092, The Research Institute of the Finnish Economy.
- Johnson, P. (1993). Aging and European economic demography. In P. Johnson & K. F. Zimmermann (Eds.), *Labor markets in an ageing Europe*. Cambridge: Cambridge University Press.
- Lallemand, T., & Rycx, F. (2009). Are older workers harmful for firm productivity? *De Economist*, 157, 273–292.
- Lazear, E. (1979). Why is there mandatory retirement. *Journal of Political Economy*, 87, 1261–1284.
- Malmberg, B., Lindh, T., & Halvarsson, M. (2008) Productivity consequences of workforce ageing—Stagnation or a Horndal effect?, *Population and Development Review*, Supplement to vol. 34, Population Aging, Human Capital, Accumulation and Productivity Growth, 238–256.
- Skirbekk, V. (2003). Age and individual productivity: A literature survey. In G. Feichtinger (Ed.), *Vienna yearbook of population research 2004* (pp. 133–153). Vienna: Austrian Academy of Sciences Press.
- Van Ours, J. C. (2009). Will you still need me—When I'm 64? *De Economist*, 157, 441–460.
- Van Ours, J. C., & Stoeldraijer, L. (2011). Age, wage and productivity in Dutch Manufacturing. *De Economist*, 159. doi:10.1007/s10645-011-9159-4.
- Van Vuuren, D., & de Hek, P. (2009). Firms, workers, and life-cycle wage profiles. In R. Euwals, de R. Mooij, & van D. Vuuren (Eds.), *Rethinking retirement*. The Hague: CPB Netherlands bureau for economic policy analysis.