

MAGNUS CARLSSON & STEFAN ERIKSSON
2017:4

The Effect of Age and Gender on Labor Demand

Evidence from a Field Experiment



Linnæus University

Discrimination and Integration Studies

Centre

The Effect of Age and Gender on Labor Demand Evidence from a Field Experiment*

Magnus Carlsson[†] and Stefan Eriksson[‡]

June 26, 2017

Abstract. In most countries, there are systematic age and gender differences in labor market outcomes. Older workers and women often have lower employment rates, and the duration of unemployment increases with age. These patterns may reflect age and gender differences in either labor demand (i.e. discrimination) or labor supply. In this study, we investigate the importance of demand effects by analyzing whether employers use information about a job applicant's age and gender in their hiring decisions. To do this, we conducted a field experiment, where over 6,000 fictitious resumes with randomly assigned information about age (in the interval 35-70) and gender were sent to employers with a vacancy and the employers' responses (callbacks) were recorded. We find that the callback rate starts to fall substantially early in the age interval we consider. This decline is steeper for women than for men. The negative age effect prevails in all seven occupations we include. These results indicate that age discrimination is a widespread phenomenon affecting workers already in their early 40s. Ageism and occupational skill loss due to aging are unlikely explanations of these effects. Instead, our employer survey suggests that employer stereotypes about other worker characteristics – ability to learn new tasks, flexibility/adaptability, and ambition – are important. We find no evidence of gender discrimination against women on average, but the gender effect is heterogeneous across occupations and firms. Women have a higher callback rate in female-dominated occupations and firms, and when the recruiter is a woman. These results suggest that an in-group bias affects hiring patterns, which may reinforce the existing gender segregation in the labor market.

JEL classification: J23, J71

Keywords: Age, Gender, Discrimination, Field experiment, Labor market

* We are grateful for helpful comments from Per Johansson and seminar participants at the ESPE 2017 Conference in Glasgow, the UCLS Members Meeting, Linnaeus University, and Members Meetings in the project “The Effects of an Ageing Population: A Life-time Perspective on Work, Retirement, Housing and Health”. We would like to thank Josefine Andersson for her work with the survey. Financial support from the Swedish Research Council for Health, Working Life and Welfare (FORTE) is gratefully acknowledged.

[†] Centre for Labor Market and Discrimination Studies, Linnaeus University, SE-391 82 Kalmar, Sweden, Magnus.Carlsson@lnu.se

[‡] Department of Economics, Uppsala University, PO Box 513, SE-751 20 Uppsala, Sweden, Stefan.Eriksson@nek.uu.se

1. Introduction

In most countries, there are systematic age and gender differences in key labor market outcomes. Older workers and women often have lower employment rates and total labor market earnings. Also, unemployment durations increase substantially with age. Understanding the reasons behind these observations is an important policy-relevant question. They could be the result of age and gender differences in either labor demand (i.e. discrimination) or labor supply. In this study, we investigate to what extent a worker's age and gender affect employers' hiring decisions, whether the effect of age differs between female and male workers, and what mechanisms may explain why employers discriminate. To do this, we conducted a field experiment on the demand side of the labor market, where over 6,000 fictitious resumes with randomly assigned information about age (in the interval 35-70) and gender were sent to employers with a vacancy. We recorded the responses, and use this information to estimate how labor demand depends on age and gender. To understand the mechanisms behind our experimental results, we also conducted an employer survey.

The existence of age and gender discrimination in the labor market could have serious consequences for both society as a whole and the individuals affected. Discrimination could make it difficult to address the demographic challenge that many Western countries face due to an aging population, which is expected to put severe pressure on public finances with increased expenditures. Avoiding large deficits will probably require an increase in labor force participation and employment among older workers, especially among women (cf. OECD 2006). However, if there is discrimination in the labor market, it will be difficult for policymakers to succeed in delaying people's retirement through policy reforms. Another consequence of discrimination could be lower worker mobility. This is likely to result in lower employer-worker match quality, and hence reduced labor market efficiency. Gender discrimination may also reinforce the existing gender segregation in the labor market. At the individual level too, discrimination is likely to have adverse effects. Many older workers have better health than previous generations (cf. OECD 2013), and hence could be both willing and able to stay longer in the labor market. However, this may require that they are given the opportunity to work under different conditions in the last years of their working life, e.g. change job or occupation. This may be difficult to achieve if there is age discrimination in hiring. Instead, they could be forced to either stay in their current job or exit the labor market. Gender discrimination could also have negative effects for individual workers, e.g. by creating barriers that prevent women and men to get a job in certain occupations.

For both age and gender, there are several theoretical reasons why employers may discriminate (cf. Becker 1957; Arrow 1973). For age, they may perceive older workers as less productive than younger workers. Employers may believe that older workers do not have the same ability to learn new tasks or the same occupational skills, or that they are less flexible and able to adapt to changes in the workplace. Employers could also perceive older workers as less ambitious and willing to work hard. In contrast, employers may prefer older workers because they are more experienced. Since many of these factors are difficult to observe before hiring, there could be statistical discrimination based on age, which is more easily observed. In addition, employers may discriminate due to ageism (i.e. a dislike of older workers).

For gender, there could be statistical discrimination for many reasons, and these considerations may vary by age. One reason is gender stereotypes among employers about which occupations are “male” and “female” jobs, and these could be stronger or weaker depending on the worker’s age. For younger workers, a possibility is that employers believe that women are more productive than men, especially in Western countries, where women nowadays have more education and better grades than men (cf. Blau et al. 2014). On the other hand, employers could perceive younger female workers as less productive than younger male workers, e.g. because younger women have more labor market disruptions due to family formation and child care. There could also be differences in how employers perceive older female and male workers. Women have a higher life expectancy and better health than men (cf. OECD 2013). Employers may, therefore, view older female workers as more productive than older male workers. Finally, employers could have preferences against hiring women and/or older women.

In 2015 and 2016, we conducted a large-scale field experiment to investigate to what extent age and gender discrimination contribute to the observed differences in labor market outcomes, and why employers may discriminate. More than 6,000 fictitious resumes for female and male applicants aged 35-70 were sent to employers with a vacancy. The measured outcome was responses from employers in the form of callbacks (e.g. invitations to job interviews). In this type of field experiment, we can identify a causal effect of age and gender since these characteristics are randomly assigned to the fictitious resumes. This is in sharp contrast to studies using survey or administrative data, which seldom have access to information about all productivity-relevant applicant characteristics that are observed by recruiting employers. If these characteristics are correlated with age or gender, a classical omitted variable bias arises. Field experiments of discrimination were developed in response to such concerns, and have frequently been used to study ethnic discrimination (e.g. Riach and

Rich 2002; Bertrand and Mullainathan 2004; Carlsson and Rooth 2007; Rich 2014; Neumark 2016)), but less to study age and gender discrimination (see below).

Our main result for age is striking. The callback rate starts to fall substantially early in the age interval we consider. One year of aging reduces the callback rate with around 0.5 percentage points. Hence, ten years of aging leads to an approximately five percentage points lower callback rate, which is a substantial effect. Our results suggest that age discrimination is a widespread phenomenon, affecting workers much younger than the age where employers consider them as old (which occurs at age 54 according to our employer survey).

Our main results for gender and the interaction of age and gender are also striking. At 35, women have an around five percentage points *higher* callback rate than men, but women's callback rate drops faster with age. Close to the retirement age women and men have a similar, and very low, callback rate. Otherwise, the callback-rate-age profiles have a similar pattern. Hence, there is no evidence of gender discrimination in hiring against women on average, but women clearly suffer more from age discrimination.

We explore the mechanisms behind the discrimination in several ways. First, we investigate the importance of perceived productivity differences. We do this by randomly varying signals in the resumes about employment status and being positive to participate in on-the-job-training. Here, we find no conclusive results. Second, we explore where in the age interval 35-70 the age effect is strongest, and find that this occurs rather early in the interval. Third, we consider heterogeneity across occupations and firms with different characteristics. We find that the age effect is very stable, while the gender effect is more heterogeneous. For gender, there is evidence of an in-group bias, where recruiters in female-dominated occupations and firms, and firms with a female recruiter prefer female applicants. This may contribute to reinforce the existing gender segregation in the labor market. Fourth, we conducted a survey on a representative sample of Swedish employers, who were asked questions about how they perceive that different skills vary by age and what their attitudes are towards older workers. The survey answers suggest that employer stereotypes about how the ability to learn new tasks, flexibility/adaptability, and ambition vary by age are important. Employers seem to worry that workers already in their 40s are starting to lose these abilities, which could result in statistical discrimination. In contrast, ageism and occupational skill loss due to aging seem to be unlikely explanations of the age discrimination we find.

The literature on age and gender discrimination using field experiments is rather limited, and we believe that we make several important contributions to it. Two of the most recent large-scale field experiments of age discrimination in the labor market are Lahey (2008) and

Neumark et al. (2015), both conducted in the US. Lahey (2008) finds that female job applicants aged 62 are less likely to be invited to a job interview than female applicants aged 35. Neumark et al. (2015) find that applicants around age 50 and 65 are less likely to be invited to a job interview than those around age 30, the difference is larger for those around age 65. Other studies are Bendick et al. (1996, 1999), Riach and Rich (2010), Albert et al. (2011), Ahmed et al. (2012), and Baert et al. (2016). Gender discrimination is studied using field experiments in e.g. Neumark et al. (1996), Riach and Rich (1987, 2006), Weichselbaumer (2004), Petit (2007), Carlsson (2011), and Baert et al. (2016). The results in this literature are mixed. There tend to be small gender differences on average, but sometimes larger differences for certain occupations or applicants with certain characteristics.

Our first contribution is that we include age as a continuous variable. This is in contrast to previous studies, which compare groups of applicants in specific age categories, usually two or three groups (e.g. Lahey, 2008; Neumark et al., 2015). The advantage of including age as a continuous variable is that it allows us to study the dynamic pattern of age discrimination. We can measure the effect of one year of aging in a wide age interval, and investigate at what age discrimination starts. Our second contribution is that we explicitly analyze the gender dimension of age discrimination. This is important since, on theoretical grounds, there are several reasons why age discrimination may vary by gender. The previous large-scale studies of age discrimination have either focused on only one gender (e.g. Lahey 2008), or have not applied with both female and male applicants in the same occupations (e.g. Neumark et al., 2015).¹ Our third contribution is that we study gender discrimination in a wide age interval, including older workers. Previous field experiments on gender discrimination often consider relatively young applicants, typically in the age interval 25-35 (e.g. Weichselbaumer 2004; Carlsson 2011). Our fourth contribution is that we conducted our own employer survey, with questions designed to shed light on why employers discriminate. This helps us interpret the experimental results.

In Section 2, we show that there are systematic age and gender differences in key labor market outcomes. In Section 3, we describe the field experiment. In Section 4, we present our results. In Section 5, we explore mechanisms that may explain the experimental results. Finally, Section 6 concludes.

¹ A few studies of age discrimination consider both female and male searchers in specific age categories (e.g. Baert et al. 2016). However, to fully analyze the age-gender interaction, age must be included as a continuous variable. Also, the small scale of most of these experiments makes it difficult to draw precise conclusions about the age-gender interaction.

2. The labor market for women and men at different ages

Using administrative data on real-world labor market outcomes, what reasons are there, if any, to suspect age and gender discrimination?

Figure 1 illustrates total labor market earnings in Sweden by age 35-65 for women and men in 2012.² For both groups, the age-earnings-profile has a positive slope until a peak at around age 45-50. After age 55, the profile starts to decline markedly. This hump-shaped relationship between total labor market earnings and age is well-documented in the literature. Earnings tend to increase in the early stages of a worker's career, flatten out when the worker is middle-aged, and then drop when the worker gets closer to the retirement age (cf. Willis 1985). Studies suggest that it is hours worked rather than hourly wages that fall with age (e.g. Johnson and Neumark 1996; Rupert and Zanella 2015), so the decline in the graphs should reflect reduced hours worked. This interpretation is supported by the graphs in Appendix Figure A2, which show a very similar pattern for the employment rate. The same pattern for the employment rate can be seen in CPS data for the US (Appendix Figure A5). It should also be noted that the shapes of the earnings and employment profiles are similar for women and men, but the levels are always higher for men (cf. Blau and Khan 2016).

Appendix Figure A3 show unemployment durations for male and female workers registered at the Swedish Public Employment Service in 2011-2012. Durations increase almost linearly with age, and the increase is somewhat steeper for women. That durations increase with age is observed in most other Western countries (cf. OECD 2016). Appendix Figure A6 uses CPS data to show that this is the case in the US, especially for men.

Worker mobility could also be affected by discrimination. If workers expect discrimination, they may not search for a new job or be unable to find a new job. Appendix Figure A4 shows the probability of changing job (between 2011 and 2012) by age for Swedish workers. This measure of worker mobility declines almost linearly with age. Mobility is very low among older workers and always lower for women than for men. That worker mobility falls with age and tenure is the case in most other Western countries, including the US (e.g. Farber 1999; Theodossiou and Zangelidis 2009).

Whether demand effects play an important role in explaining these patterns is what we investigate in the field experiment.

² For visual clarity in the graphs, we exclude individuals aged over 65, where most labor market outcomes change dramatically. The details of the analysis is described in the notes below the figures. In the experiment, we focus on low- and medium-skilled occupations. In Appendix Figure A7-10, we show that the patterns for low educated workers are very similar to all workers.

3. Experimental design

3.1 Age, gender, and other worker characteristics

In the experiment, age is included as a continuous variable in the age interval 35-70.³ This has several advantages. First, we can study the effect of one year of aging in a wide age interval rather than only comparing differences between a few age groups. Second, we can investigate at which age discrimination starts. Third, it allows us to construct a graph of the callback-rate-age profile, which can be compared to patterns in administrative data. Finally, we can investigate age discrimination close to, and even above, the current retirement age, which is important for pension reforms. Age is randomly assigned to the resumes and is drawn from a uniform distribution.

Gender is signaled through the name of the applicant. In Sweden, there is usually a clear distinction between female and male names. Names are randomly assigned to the resumes, and around 50 percent were assigned a typical female and male name, respectively. We chose three of the most common female and male names for people in the age interval 35-70 according to Statistics Sweden's name register.⁴

In the resumes, we also use two characteristics to signal productivity, employment status and being positive to participate in on-the-job-training. We include these signals mainly to interpret age discrimination, but potentially they could also help us interpret gender discrimination (e.g. if employers perceive women's productivity as more uncertain than men's).

We signal employment status by randomly assigning the job applicants a spell of unemployment in the interval 0-36 months.⁵ The idea is to investigate whether older workers are statistically discriminated because employers worry that they have worse unobserved skills, e.g. insufficient occupational skills. We do this by analyzing if being employed (or short term unemployed) rather than (long-term) unemployed raises the callback rate more for

³ There are several reasons why we choose 35 and 70 as the limits of the age interval. First, most workers at age 35 have both finished their education and worked for a number of years. This is important since all workers in the experiment are assigned at least ten years of relevant work experience. Second, although the current retirement age in Sweden is 61-67, there is a discussion that people should work a few additional years.

⁴ We used lists of the most common first names for workers born in the same years as our fictitious applicants. We also compiled a list of the most common surnames. Then, we randomly combined the first names and surnames. The male names are Anders Karlsson, Lars Johansson, and Peter Nilsson. The female names are Anna Eriksson, Eva Olsson, and Lena Persson.

⁵ In official unemployment statistics, a duration of unemployment of 12 months or more is classified as long-term unemployment. We chose 36 as the upper limit to include long-term unemployed workers with long durations, which is rather realistic among older unemployed.

older than younger job applicants.⁶ If this type of statistical discrimination is important, older *employed* workers should to a larger extent be perceived as having occupational skills that are comparable with younger workers than if we compare *unemployed* older and younger workers. Employment status is randomly assigned so that 1/3 of the applicants are on-the-job searchers and 2/3 of the applicants are uniformly distributed over the interval 1-36 months of unemployment.

We express that an applicant is positive to participate in on-the-job-training by stating a willingness to take part in such training in the resume. The idea is to investigate whether employers discriminate older workers because they are not perceived as adaptable to workplace changes. Survey evidence indicate that not being flexible and adaptable are among the main concerns employers have about older workers (cf. Section 5.4). We randomly assigned half of the resumes the sentence signaling this characteristic.

3.2 Employment histories for younger and older job applicants

The construction of employment histories in a field experiment on age discrimination must inevitably consider the fact that older workers have lived for more years than younger workers. Previous studies have used three alternative designs to handle this fact. We implement all of these alternatives.

The first design fills the resumes with a full history of relevant work experience. The argument for this design is that it is the most reasonable, since older workers tend to have more relevant work experience than younger workers (cf. Riach and Rich 2010). The argument against this design is that it does not capture the pure age effect, since older workers may be favored by the fact that they always have more relevant work experience than younger workers.

The second design assigns both younger and older workers a fixed number of years of relevant work experience, and then leaves out the earlier part of the employment history (e.g. Lahey 2008). This is the standard approach in field experiments on discrimination on other grounds, e.g. ethnicity, where the job applications are designed to be qualitatively identical. In the case of age discrimination, this design can make the age effect difficult to interpret (cf. Riach and Rich 2010). Employers may assume that older workers have more relevant work experience than the resume reveals, which potentially could generate statistical discrimination that favors older workers. In our experiment, this design means that we assign all job

⁶ Previous research has shown that employers prefer not to recruit long-term unemployed workers (e.g. Kroft et al. 2013; Eriksson and Rooth 2014).

applications ten years of relevant work experience, which is realistic for workers in the age interval we study.⁷

The third design attempts to hold perceived relevant work experience constant by explicitly signaling that older applicants do not have more such experience. This is typically done by including a complete employment history, in which it is stated that the older applicant has been engaged in a previous activity that is assumed not to affect the applicant's productivity in the current occupation. Examples of such activities are working in a very different (unqualified) occupation, which is the design we use, taking care of children at home, and being active in the military (e.g. Ahmed et al. 2012). A potential problem with this approach is that it creates a correlation between being older and having experience in a particular previous activity. Therefore, this strategy will only work if this activity is irrelevant for the recruiting employers.

The advantage of using all three designs in the experiment is that we can investigate whether the (theoretical) arguments that have been raised against each of the designs actually are important when conducting a field experiment.

Our approach is similar to Neumark et al. (2015), who use the first two designs by including both older and younger job applicants that have an equal amount of previous work experience as well as older job applicants that have a full CV history of previous work experience (i.e. more experience than the younger applicants). They find that the choice of design does not affect the age results, which is also what we find.⁸

3.3 Generating resumes

To create realistic resumes, we combined information from a large number of real job applications available in a database at the Swedish Public Employment Service. We also consulted colleagues and used our previous experience of conducting field experiments. Each generated resume consists of two parts – a cover letter and a CV (Appendix Figure A11 and A12 shows an example). The cover letter starts with a short presentation of the applicant that includes the applicant's name and age, a description of work experience, and some information about personal interests. The CV includes information about the applicant's full

⁷ Job applicants in Sweden are often given the advice to only include the most relevant work experience in their resumes. Lahey (2008) argues strongly that including ten years of work experience is what US firms prefer.

⁸ Baert et al. (2016) also analyze this issue and label it "the post educational years problem". They consider job applicants aged 38, 44, and 50. They assign all applicants the same amount of work experience in the relevant occupation immediately after leaving school as well as immediately before applying to the job in the experiment. However, for an intermediate period, they either leave the period empty, add a job that is assumed to be unrelated to the productivity in the job applied for, or add a job in the relevant occupation.

name, date of birth, contact details, work experience, education, computer skills, driving license, and some occupation-specific certificates.⁹

All resumes were generated before the experiment started. The first step in generating the resumes was to create templates, which specify the structure, layout, and typeface of the resume, and contain some general phrases and information. We use a specific template for each occupation.

In the second step, the resume templates were filled with content, which depended on the randomized variables and the occupation and city of the advertised vacancy. The first two randomized variables were the name and age of the applicant, which are stated both in the cover letter and the CV. In Sweden, it is very common that job applicants explicitly state their age in their resumes.¹⁰ The third random variable was employment status (duration of unemployment), which is signaled in the CV part of the resume. The CV lists previous jobs and an unemployed worker has a gap in the employment history. The fourth randomly assigned variable was the signal of being positive to participate in on-the-job-training. This is signaled by a sentence in the cover letter part of the resume that states “I enjoy taking courses and participating in on-the-job-training to improve my skills” (as opposed to no information about this). We also randomized the design of the employment history, as described above, in the CV part of the resumes. The education and work experience of the applicant was determined by the occupation. All applicants were given an occupation-relevant high school education, and at least ten years of occupation-relevant work experience (two jobs with random tenure summing up to ten years). Since we know from previous field experiments that employers very seldom respond by surface mail, we used a fictitious postal address in the same city as the job. The addresses were located in similar suburbs (in terms of the socio-economic characteristics of the residents) not too far from the city centers.

To enable employers to contact the job applicants, an email address and a telephone number (with an automatic answering service) were included in the applications. These were registered at a large Internet provider and phone company, respectively. Each of the six names we used in the experiment had a separate email address and telephone number.

There is an efficiency argument for sending several resumes to each employer, since a given number of observations can then be collected using less resources. However, the risk of making employers suspicious also increases when more resumes are sent to the same

⁹ All applicants had a driving license for a private car. In addition, truck drivers had a driving license for a truck. Truck drivers were also given a few other occupation-required certificates.

¹⁰ We have access to a large database of job applications, and in that database the majority of the job applicants mention their age in both the cover letter and the CV.

employer. We had to trade off these two issues and decided that it was reasonable to send three resumes to each employer. This design means that we had to construct three types of resume templates, which were different in terms of their structure, layout, typeface, and general phrases in order to not make employers suspicious. The resume type was randomly assigned, and hence cannot affect our measure of discrimination.

3.4 Sampling of vacancies

We planned for a sample of around 6,000 job applications, which would be sent in pairs of three to around 2,000 employers. The size of the experiment was determined by power calculations¹¹, the resources available for this project, and our previous experience of conducting field experiments. Vacancies were sampled in the three major cities in Sweden – Stockholm, Gothenburg and Malmö – which together cover a clear majority of all Swedish vacancies. In these cities, we randomly sampled firms in seven occupations that had an advertisement on the website of the Swedish Public Employment Service.¹²

When choosing the occupations, our aim was to include a sufficient number of common occupations to get a representative picture of the Swedish labor market. We started from a list of all occupations in Sweden and sorted it with respect to the number of workers employed in each occupation.¹³ Starting from the top of this list, we included each occupation that fulfilled a number of criteria.

The first criterion was that it must be possible to apply by email. Many employers in the private sector accept applications by email. In contrast, most employers in the public sector use web-based recruitment systems, where job applicants have to state their social security number. Therefore, we could not include occupations in health care and teaching, which in Sweden are dominated by the public sector. This criterion excluded ten occupations on the top 25 list.

The second criterion was that the occupation is not high-skilled, i.e. requiring a university education. For such jobs, employers are likely to use the Internet to screen job applicants, e.g.

¹¹ Before conducting the experiment, we performed simulation exercises to determine the appropriate sample size. Suppose we want to be able to detect an age effect of five percentage points for each ten years of aging with 90 percent certainty. We make 1,000 draws from a population characterized by such an age effect, an average callback rate of 10 percent (which is similar to the callback rate in Bertrand and Mullainathan, 2004), and an error term that also affects the probability of a callback. Then, we need a sample size of 400 to get p-values less than .05 for the age estimate in more than 90 percent of the cases. To be able to detect such an age effect in each of the seven occupations, and for female and male job applicants separately, a sample size of at least 5,600 (400*7*2) job applications is required.

¹² The Employment Service estimate that 30-40 percent of all vacancies are reported to them.

¹³ We define an occupation at the most detailed level in the Swedish Standard Classification of Occupations (SSYK 2012 which is similar to ISCO-08). The total number of categories is 429.

by using LinkedIn. However, the fictitious job applicants will not be found on the Internet, which could make employers suspicious. High-skilled occupations also often require more elaborate resumes, tailored to a specific advertisement, which are difficult to generate in an automated way. This criterion excluded two more occupations on the top 25 list.

The third criterion was that the number of advertisements in an occupation posted on the website we use is large enough to make each occupation separately analyzable in terms of statistical power. We checked this by counting the number of new advertisements available per day in each occupation. This criterion did not exclude any further occupations.

After these three criteria were applied, we had 13 possible occupations. However, we decided that it is reasonable to treat office assistants, accounting assistants, and secretaries as the same occupation, since the job tasks are similar. We label this occupation administrative assistants. We also treated food and non-food retail salespersons as one occupation, which we label retail sales persons and cashiers. This reduced the number of occupations to ten.

Finally, to avoid too much heterogeneity (and hence noisy estimates), we did not want to include extremely male- or female-dominated occupations. For male-dominated occupations, there were four occupations on our list where at least around 80 percent are men. We kept the largest of these occupations (truck drivers), and excluded the other three (warehouse workers, janitors, and construction workers). For female-dominated occupations, there is only one occupation with at least 80 percent women, which we kept (administrative assistants). We also checked for extreme deviations in the age distribution of the occupations, but did not find any reason to exclude any further occupations.

In the end, the following seven occupations are included: administrative assistants, chefs, cleaners, food serving and waitresses, retail sales persons and cashiers, sales representatives, and truck drivers. These are occupations that are common not only in Sweden, but also in most other Western countries, including the US. Appendix Table A1 presents descriptive statistics about the number of workers employed, the age and gender distributions, and the share of reported vacancies in these occupations. Each occupation employ between one and four percent of the total number of workers, which are relatively large numbers given the fact that there are 429 occupations at the four-digit level. The fractions of vacancies reported to the Employment Service in these occupations seem representative of the size of the occupations, i.e. the fractions are of similar magnitude as the share of workers employed in each occupation.

3.5 Conducting the experiment and recording responses

Between August 2015 and March 2016, we sent 6,066 resumes to 2,022 employers with an advertisement on the chosen website. For each advertisement, three resumes were randomly selected and sent in random order to the employer with a one day delay in-between. The employers mostly replied by email, or in some cases by leaving a voicemail message.¹⁴ After having recorded a reply, we promptly declined all invitations to job interviews.

From the recorded replies, we constructed a callback dummy variable which is coded as one in the case of a positive response (i.e. any response that can be interpreted as expressing an interest in the applicant) and zero in the case of a no response or a negative response. An alternative would be to use an indicator which takes the value one only for explicit invitations to a job interview. However, it is likely that many employers do not invite to interviews in an e-mail (or by leaving a message on an answering service), although such contacts may eventually lead to an interview. We consider the case with only explicit invitations to interviews as a robustness check, and get very similar results.

Table 1 shows the number of resumes sent in each of the seven occupations (column 1). For each occupation, we sent between around 500 and 1,100 resumes. Column 2 shows the corresponding shares of sent resumes, which are between 8 and 19 percent. The last column shows the callback rates in the seven occupations, which varies between around 4.0 and 15.8 percent with an average of around 8.7 percent. The variation in the number of resumes and callback rates is likely to reflect differences in labor demand between the occupations.

4. Results of the field experiment

4.1 The effect of age and gender on the callback rate

For age, our main result is striking. For both women and men, the callback-rate drops substantially early in the age interval we consider and there is a clear negative relationship between the callback rate and age, which Panel A and B in Figure 1 clearly show. Also, close to the retirement age the callback rate is almost zero, only around 2-3 percent. Table 2 (the first row in column 1) reports the magnitude of the age effect.¹⁵ The coefficient of age is around $-.0048$, and is statistically significant at the one-percent level. Hence, the callback rate falls by approximately five percentage points for each ten years of aging, which shows that

¹⁴ We were able to match almost all responses to a particular application (less than ten responses were impossible to match).

¹⁵ The coefficients are obtained from estimating a linear probability model with the callback indicator as the dependent variable and with age, gender, all other randomly assigned applicant characteristics, and fixed effects for the occupations and cities as explanatory variables. In the baseline regression, we include age linearly, but also consider other specifications (see Section 5.2).

the magnitude of the age effect is substantial. These results clearly establish that employers view age as a negative factor in the recruitment process, and suggest that age discrimination is important in the labor market.

For gender, there are a number of interesting results. First, female applicants have, on average, an approximately 1.4 percentage points *higher* callback rate than men, and this effect is statistically significant (at the ten-percent level; second row of column 1). In addition, the callback-rate-age profiles for women and men in Figure 1 are similar; both show the early fall in the callback rate and the clear negative relationship between the callback rate and age. However, there are also two highly visible differences. Early in the age interval, women have a higher callback rate than men, and the decline in the callback rate is steeper for women than for men. A 35-year-old woman has a callback rate that is several percentage points higher than a man of the same age, while close to the retirement age the callback rates are very similar, and close to zero. The second column in Table 2 shows the separate age coefficients for women and men. For applicants aged 35, women have a 4.8 percentage points higher (statistically significant) callback rate than men. The size of the age effect is around .0057 for women and .0038 for men, a difference which is statistically significant (p -value = .007). These results establish that, on average, employers do not view women negatively in the recruitment process, and suggest that gender discrimination against women in hiring is not important on average. However, the results show that the negative effect of age is stronger for women.

4.2 Robustness analysis

In the regressions in Table 2, we use an indicator of a positive response as the dependent variable. As mentioned above, an alternative is to instead use an indicator of an explicit invitation to a job interview as the dependent variable. Column 5 in Appendix Table A2 shows that our main results are very similar if we use this outcome variable. The coefficients are somewhat smaller in size, but this is as expected since the average callback rate is lower for job interviews.

Appendix Table A2 also reports the results of a number of other robustness checks of the main results for age and gender; excluding all other covariates than age and gender (column 2), including firm fixed effects (2,022 dummy indicators; column 3), and using the Probit model instead of the linear probability model (column 4). The results are very similar in all cases, which should be the case for experimentally generated data. The fact that we get very similar results with and without the other covariates also confirms that the random assignment

of the experimental variables has worked as intended, and hence that our estimates of age and gender represents causal effects. Random assignment also implies that no pair of random variables should show any sign of a substantial correlation. Appendix Table A3 shows that this holds for all our experimental variables.

Another important issue is whether the design of the applicants' employment histories affects the results. Table 3 presents the effects of the employment history variables and the interaction effects between these variables and age. The results show that the age coefficient is similar for all three CV gap categories; using an F-test we cannot reject the hypothesis of equal age slope coefficients (p -value = .351). Hence, these results suggest that the choice between the three employment history designs is not important for the results in a field experiment on age discrimination. As mentioned above, this is similar to the results in Neumark et al. (2015).

5. Analyzing the mechanisms

5.1 Do the signals of productivity matter?

As mentioned above, we included two signals in the resumes to test whether uncertainty about older workers' productivity is an important channel for age discrimination.¹⁶ As a first signal, we randomized the applicants' employment status. We construct the variable as a dummy indicator, which equals one if the applicant is long-term unemployed (i.e. unemployed twelve months or more) and zero if the applicant is an on-the-job searcher or short-term unemployed.¹⁷

Panel A in Table 4 reports results for this signal. The first column mirrors the main result in Table 2 for the age effect, but also shows that there is a strong negative effect of being long-term unemployed (the coefficient is -.018 and highly statistically significant). This result is in line with previous studies (Kroft et al. 2013; Eriksson and Rooth 2014), and shows that long-term unemployed workers face a challenge in finding employment. Importantly, this also confirms that the employers have noticed this signal in the resumes. In the second column, we investigate if the negative effect of long-term unemployment interacts with age. The results show that the age coefficients are rather similar in magnitude for those who are long-term unemployed and those who are not; an F-test cannot reject the hypothesis that the age coefficients are equal (p -value = .419). In column 3, we estimate separate age coefficients for

¹⁶ We have also interacted these signals with gender, but find no statistically significant effects.

¹⁷ We use this design because previous studies suggest that it is long-term unemployment, and not shorter spells of unemployment, that is perceived by employers as the strongest negative signal. We get a very similar result if we instead use a continuous variable measuring the duration of unemployment.

the four subgroups defined by employment status and gender. The results show that there is no evidence of a difference in the age coefficient between the employment status subgroups for either women (p-value = .174) or men (p-value = .680).

As a second signal of productivity, we use the sentence in the resume that is intended to signal being positive to participate in on-the-job-training. Panel B reports the results for this signal. Again, the first column mirrors the main result in Table 2 for the age effect, but also shows that there is no main effect of the signal. One interpretation of this is that employers do not value this characteristic, but this seems inconsistent with the literature which shows that flexibility and adaptability are considered as important by employers. However, being positive to participate in on-the-job-training is only one dimension of the more general concepts of flexibility and adaptability. These broader concepts include not only willingness to participate in training, but also other things such as being open to alternative solutions and accepting new tasks and roles. Hence, although we do not find a main effect of being positive to participate in on-the-job-training, flexibility and adaptability may still be important for employers. An alternative interpretation is that employers may not have noticed the signal in the resumes, although it is clearly expressed in a separate sentence in the cover letter. Regardless of why we find no main effect of this characteristic, this questions the value of analyzing this variable any further. Given this caveat, the second column shows separate age coefficients for those who signal being positive to invest in on-the-job-training and those who do not. The age coefficients are similar in magnitude; an F-test cannot reject the hypothesis that they are equal (p-value = .285). For women and men separately, there is no evidence either of a difference in the age coefficient between applications with and without this signal.

5.2 At what age is the age effect strongest?

Investigating in which age interval the negative age effect is strongest may help us interpret the main results. A natural way of exploring this is to include dummy variables for different age categories. We divide the age interval 35-70 into four categories of equal size, $36/4 = 9$ years.¹⁸ This gives us four age cutoffs: 44, 53, and 62. For each cutoff, we construct a dummy variable, which is equal to zero if the age of the applicant is strictly below the cutoff and otherwise equal to one. This allows us to estimate the change in the callback rate as result of a

¹⁸ We get similar results if we use other age categories of equal size (e.g., 36/3 or 36/6). An alternative way of investigating if the age effect varies across different ages is to include age squared in addition to the linear age term in the main regression. The result is an age coefficient of -.0075 (significant at the one percent level) and an age squared coefficient of .0001 (significant at the five percent level). This analysis confirms the picture above, i.e. that the age effect is strongest early in the age interval we consider and decreases (in absolute terms) in magnitude by age.

change in discrete age category, i.e. when changing from 35-43 to 44-52, from 44-52 to 53-61, and from 53-61 to 62-70.

The results of regressing the callback dummy on the age category dummies are shown in the first column in Table 5. The change in the callback rate for a change from age 35-43 to age 44-52 is sizable, a 7.6 percentage point drop in the callback rate, and strongly statistically significant. In comparison, the effect of changing from 44-52 to 53-61 and from 53-61 to 62-70 is smaller (but still statistically significant), less than half the effect of changing from 35-43 to 44-52. An F-test rejects the hypothesis of a constant age effect, i.e. equal age category coefficients in the model in column 1. Model 2 repeats the analysis, but estimates separate age category dummy variables for female and male applicants. All age category coefficients are larger for women, which is consistent with the larger age effect for female applicants in the main analysis. In an F-test of a constant age coefficient, the p-value is .014 for women and .105 for men.

5.3 Do the effects differ between occupations and firms?

An interesting issue is what characterizes occupations and firms where the age and gender effects are strong. We explore heterogeneity with respect to the occupations, the firms' gender composition, and the gender of the recruiters.¹⁹

Panel A in Table 6 shows results on heterogeneity across the occupations. We use a single regression, similar to the main specification, but where the occupations are interacted with age and gender. The regression does not include the covariates for age and gender. Therefore, the coefficients for the interactions should be interpreted as the age and gender effect for each occupation.

The first row shows the age effect across the occupations in the experiment, and there are at least two things worth noting. First, there is a statistically significant negative age effect in all occupations. This confirms that age discrimination is a widespread phenomenon, not limited to a few occupations. Second, the age effect differs across the occupations; an F-test strongly rejects that the age coefficients are equal (p-value = .000). The age effect is largest for truck drivers, chefs, and sales representatives.

¹⁹ In addition, we have tested for differences in the age and gender coefficients between firms located in Stockholm, which are around 70 percent of the sample, and outside Stockholm (i.e. Gothenburg or Malmö). We find no statistically significant differences for the age and gender coefficients. Finally, we have information on the firms' size from Statistics Sweden. We divided the sample of firms around the median number of employees (0-19 vs. >19). Again, we find no statistically significant differences for the age and gender coefficients.

The second row of Panel A shows the gender gap in the callback rate across the occupations. Although there is only a small gender difference on average, larger differences emerge at the occupational level. For administrative assistants and cleaners, women are preferred over men. These are the two most female-dominated occupations. Among truck drivers, chefs, and sales representatives, women have a slight disadvantage, but only for chefs is this result statistically significant. These are the most male-dominated occupations.

In Panel B, we estimate separate age coefficients for female and male applicants in each occupation. The results show that there is a statistically significant steeper age slope coefficient for women than for men only for administrative assistants and cleaners (see the F-tests in the last row). Hence, the reason why women, on average, are preferred in these occupations seems to be because younger female applicants have an advantage. The gender difference in the callback rate for applicants in these occupations at age 35 are 11.8 and 16.7 percentage points, respectively (see the first row of Panel B).

Panel A in Table 7 investigates heterogeneity with respect to the gender composition of the firms. The first column mirrors the main results, while the second column repeats the same specification for the subsample of 4,607 job applications where we have information about the gender composition of the firm.²⁰ The age and gender effects are rather similar for this subsample, which suggests that the subsample is representative. In the third column, we estimate separate age coefficients for firms that are male-dominated (0-50% females) and female-dominated (50-100% females). The age coefficient is very similar in the two groups. In the fourth column, we estimate separate gender coefficients for the two types of firms. The female coefficient differs substantially, and an F-test shows that the difference is statistically significant (p-value = .000). Women have an advantage in firms that are female-dominated, and a disadvantage in firms that are male-dominated. Since this analysis includes fixed effects for the seven narrowly defined occupations, the comparison is within occupations, and hence the results cannot be explained by occupational characteristics (e.g. stereotypes about which of these occupations are “female” and “male” jobs).

Panel B in Table 7 investigates heterogeneity with respect to the gender of the recruiter. We gathered information about the gender of the recruiter from the name of the contact person in the advertisement, and coded the gender based on the name. This contact person is likely to be the recruiter or someone who, at least, is involved in the recruitment process. A contact person was specified for a little more than half of the advertisements. Again, the first

²⁰ The information about the share of women at the firm is from Statistics Sweden and is missing for those firms not included in the subsample.

column mirrors the main specification, while the second column repeats the main specification for the subsample of 3,204 job applications where we have information about the gender of the recruiter. In this case too, the age and gender coefficients are rather similar in both samples. In the third column, we estimate separate age coefficients for firms that have a female and a male recruiter, respectively. The age coefficient is very similar in the two groups. Finally, in the fourth column, we estimate separate gender coefficients for the firms with female and male recruiters. The female coefficient differs substantially depending on the gender of the recruiter (p -value = .008). Women are preferred by female recruiters, while male recruiters appear indifferent between female and male job applicants. Again, the regressions include fixed effects for the occupations, so the results should be interpreted within occupations.

5.4 What do the employers in the survey say?

To learn more about why firms may discriminate, we conducted a survey, where employers were asked a number of questions about how they perceive workers at different ages.²¹ In the spring of 2016, a representative sample of Swedish employers was contacted by phone to find a person working with recruitment that could answer a web-based survey.²² For 3,937 employers such a person was found, and 1,344 (34 percent) of them responded.

One question in the survey asked about a number of worker characteristics at different ages. The question was: “Suppose that you are recruiting a new employee to a typical position in your workplace. To what extent, do you think that an average employee at age 20, 30, 40, 50, and 60 would have the following characteristics”. The characteristics were: i) self-sufficient, ii) able to learn new tasks, iii) flexible/adaptable, iv) ambitious, v) structured, vi) technical occupational skills, vii) communication skills, viii) reliable/loyal, ix) cooperation

²¹ Some of the questions in the survey were constructed specifically for this study, but the survey also contained other questions about aging and the labor market that are used in a larger project on aging. This means that all questions in the survey are about the effects of age.

²² A stratified sampling strategy was used, where the strata are based on the sector and the number of employees. The sectors are the government sector (10%), the municipal sector (30%), and the private sector (60%). In each of these three sectors, workplaces were included in a falling scale depending on the number of employees, i.e. all workplaces in the largest size categories and a random sample in the smallest size category. The company administrating the survey contacted 6,066 workplaces by phone to find a person working with recruitment that the survey could be sent to. For 3,937 workplaces such a person was found. A link to the web-based survey was sent by e-mail to these workplaces, including several reminders to non-responders.

skills, x) leadership skills.²³ The scale used goes from 0 (to a very small extent) to 10 (to a very large extent).

There are three characteristics which employers responded they believe deteriorate with age: able to learn new tasks, flexible/adaptable, and ambitious. The results for these questions are shown in Figure 2, which for workers aged 30, 40, 50, and 60 plots the average difference in the score compared to workers aged 20 (i.e. relative to the horizontal line). The first two characteristics peak at age 30 and then decline, while the last characteristic peaks at age 40 and then declines. The 95 percent confidence intervals show that all three declines are statistically significant. The fact that the perceived deterioration starts at age 30-40 suggests that, at least part of, the explanation why employers discriminate relatively young applicants is that they statistically discriminate based on age and their perceptions about these characteristics.

For the other seven characteristics, the answers to the survey suggest that employers do not expect that these factors deteriorate markedly with age. These characteristics either stay rather constant in the older age categories or decline only in the oldest age category (i.e. at age 60; see Appendix Figure A13). In particular, this is the case for occupational skills, which may explain why we did not find an interaction effect between age and employment status in the experiment. This signal is likely to mainly test employers' perceptions about occupational skills, and if the perception is that these skills do not decline much with age the interaction effect will not be important.

The results of other surveys on this topic are in line with our results. Most other such surveys ask employers about whether older workers have less of certain skills/abilities than younger workers (e.g. Taylor and Walker 1998; AARP 2000; Grey and McGregor 2003; Henkens 2005; Swedish Pensions Agency 2012). The results often show that employers worry that older workers are less able to learn new tasks, flexible/adaptable, and ambitious. Also, there is often no clear evidence that employers believe that older workers have less occupational skills or are less productive. Some surveys also find that employers worry about the physical strength and health of older workers. For our purposes, a potential caveat with these surveys is that they often define older workers quite loosely, e.g. as workers above age 50.

²³ Another question in the survey asked about how important, in general, these characteristics were for the employers. The scale is from 0 (unimportant) to 10 (crucial). The average score for the ten characteristics varied between 6.45 and 8.87.

Three other questions in our survey were intended to measure the employers' attitudes towards younger and older workers. The first of these questions asks about efficiency, and is intended to measure the importance of statistical discrimination. The answers show that most employers state that both younger and older workers contribute to an efficient production to a similar extent (see Appendix Figure A14). The second and third of these questions ask about how younger and older workers contribute to the working-environment, and attitudes to younger and older workers in general. These questions are intended to measure the importance of ageism. Again, most respondents state that both younger and older workers contribute to a good working-environment, and that none of the groups take up too much space in society (see Appendix Figure A15 and A16). A potential problem with these questions is social desirability bias. Another question asked about at what age employers start to consider a worker as old in the labor market. The answers suggest that this occurs around age 54.

5.5. How should the results be interpreted?

Explanations of age discrimination

One of our most striking results is that the negative age effect starts at such an early age. Our analysis shows that the marginal effect of age is strongest long before the age where workers are considered as old by employers. This suggests that ageism is unlikely to be the main explanation of age discrimination, since it is unreasonable to expect that ageism should affect workers in their early 40s. Our survey results point in the same direction, revealing no evidence of ageism.

The early onset of the age effect also makes it unlikely that stereotypes about occupational skills, physical strength, and health among older workers are important explanations of discrimination. In this case too, it is unreasonable to expect that workers already in their early 40s should lack important occupational skills, have low physical strength, or bad health. Our survey contains a question about occupational skills. The answers show that employers seem to believe that occupational skills actually improve by age, at least until age 50. This is also consistent with the absence of an interaction effect between age and being long-term unemployed, which should capture employer stereotypes about how occupational skills evolve as workers age.

The obvious question is what remaining mechanisms there are that potentially can explain the bulk of the substantial negative age effect. A candidate should both be relevant for workers in their early 40s, where the negative age effect starts, and be independent of

occupation, since there is a strong age effect in all occupations we consider. This is a difficult question to fully answer, but our analysis points out a few candidates as more plausible.

Our survey shows that there are three broad characteristics that employers report are important, and which they worry that workers over age 40 have started to lose. These are the ability to learn new tasks, being flexible/adaptable, and being ambitious. Hence, a potential explanation of the age discrimination is that employers statistically discriminate because they believe that workers lose these abilities rather early in the age interval we consider. These abilities are also general enough to be important in most occupations, albeit at different degrees.

Although our results suggest that some explanations of age discrimination are plausible, while others are less plausible, we cannot fully rule out any explanation. For example, our results show that the callback rate continues to fall for the oldest workers (say, among those aged 55-70), but at a declining rate. Employer perceptions about occupational skills, physical strength, and health may explain part of the decline in the callback rate among the oldest workers.

Alternative demand-side explanations of the age effect

An important issue is if there are other demand-side explanations than age discrimination that could explain the age effect.

A potential candidate is that it may be more costly to hire older workers than younger workers. This could be the case if wage-setting is based on seniority, so that employers have to pay older workers a higher wage irrespective of their productivity. Seniority-based wages are common in some European countries, but in Sweden age is not a relevant factor in the collective agreements that to a large extent determine wages for workers in low- and medium-skilled occupations. Moreover, most of the theoretical arguments for seniority-based wages are built on long-term implicit contracts for existing employees (cf. Lazear 1979). Therefore, it is not obvious that these considerations should be important for the entry wages of new hires. Another reason why it may be more costly to hire older workers is if there are other wage-related costs, e.g. pension contributions, which are higher for older workers. In Sweden, pension contributions for blue-collar workers are not age-dependent. Hence, for most of our occupations this is not a relevant issue. A third reason why it may be more costly to hire older workers is if employers pay substantial training costs for their newly hired workers. Then, it may be argued that employers have incentives to avoid older workers, since such workers are expected to remain employed for a shorter period of time. However, for the occupations we

study, we do not expect training costs to be important, especially since all job applicants have at least ten years of occupation-relevant work experience. Also, worker mobility is lower for older workers (cf. Appendix Figure A4), which means that it is not obvious that expected tenure is shorter for older workers.

Of course, we cannot rule out the possibility that one reason that some employers avoid older workers is because they *perceive* them as more costly to hire. However, if these perceptions are false, this is just an example of statistical discrimination. We conclude that it is unlikely that other demand-side considerations than discrimination are the main explanation of the age effect we find.

The heterogeneous gender effect

An interesting question is how we should interpret the heterogeneous gender effect. A quite coherent picture emerges on how the callback rate for female applicants varies with the occupation, the gender composition of the firm, and the gender of the recruiter. Female applicants are to a larger extent preferred in the most female-dominated occupations, in firms with a high share of female employees, and in firms where the recruiter is a woman.²⁴ Moreover, the latter two findings are within narrowly defined occupations (i.e. from a regression with occupational fixed effects), and hence are unlikely to be explained by stereotypes about which of these occupations are seen as “female” and “male” jobs, respectively. Instead, these results suggest an in-group bias, where women prefer female applicants.²⁵ This may contribute to reinforce the gender segregation in the labor market.

6. Conclusions

In most countries, there are systematic age and gender differences in key labor market outcomes. These differences may reflect differences in either labor demand (i.e. discrimination) or labor supply. In this study, we investigate the importance of demand effects by analyzing to what extent employers use information about a job applicant’s age and gender in their hiring decisions. We conducted a field experiment, where over 6,000 fictitious resumes with randomly assigned information about age (35-70) and gender were sent to

²⁴ Some other recent field experiments on gender discrimination also find that female applicants are preferred in female-dominated occupations, e.g. Carlsson (2011), Riach and Rich (2006), and Weichselbaumer (2004). Baert et al. (2016) do not find this pattern. Petit (2007) focuses on the financial sector and do not analyze this dimension.

²⁵ Hewstone et al. (2002) discuss theories that may explain in-group bias. Empirical evidence of in-group bias in hiring is found in e.g. Davison and Burke (2000). Reasons why in-group bias may be stronger among women is discussed in e.g. Rudman and Goodwin (2004).

employers with a vacancy and their responses (callbacks) were recorded.

We find that there is a strong negative age effect in all occupations that we investigate. The callback rate starts to fall substantially early in the age interval we consider, and close to the retirement age the callback rate is close to zero. The early fall in the callback rate suggests that the main story of age discrimination in the labor market is not about being old, say above age 55, but rather about not being young, say below age 40-45. Employer stereotypes about the ability to learn new tasks, being flexible/adaptable, and being ambitious appear to be plausible explanations of the age discrimination, while ageism and expectations about occupational skill loss due to aging are less plausible explanations.

For gender discrimination, we find that women aged 35 have an around five percentage points higher callback rate than men of the same age, but women's callback rate drops faster with age. The gender effect is heterogeneous across occupations and firms; women have a higher callback rate in female-dominated occupations, female-dominated firms, and firms where the recruiter is a woman. This heterogeneity suggests that in-group bias affects hiring patterns, which may reinforce the gender segregation in the labor market.

An important issue is the external validity of our results, i.e. to what extent they are valid beyond the specific search channel, occupations, and cities we consider. It is well-known that many vacancies in both Sweden and other countries are filled through informal search channels. However, there are no obvious reasons to expect that an applicant's age or gender should matter less in informal recruitments. Rather, it may be argued that discrimination could be more important when hiring is less transparent. Concerning the occupations, they are rather representative of the low- and medium-skilled parts of the labor market both in Sweden and in most other Western countries. If the results are valid for high-skilled occupations are less certain, but those occupations are difficult to study with this method. The fact that experience may be more important in high-skilled occupations suggests less age discrimination.²⁶ However, the similarity of the age profiles for the duration of unemployment in both the whole labor market and the labor market for low-skilled workers suggests that our results may be representative.

Another aspect of external validity is if our results are specific for Sweden or valid for other countries as well. In an international comparison, labor force participation among older people and women is high in Sweden (in 2015, the employment rate was 74.6 percent for workers aged 55-64, and 75.1 percent for women aged 16-64). This implies that our results

²⁶ Kuhn and Shen (2013) and Kuhn et al. (2016) report evidence that advertisements for jobs with higher skill requirements are less likely to request workers of a specific age or gender.

are more likely to be representative for other countries with high labor force participation in these groups. For countries with lower labor force participation among older people and women, the negative effects may be even larger, since employers in these countries probably have much less experience of hiring such workers.

An interesting question is if the strong negative demand effect of age that we find evidence of leaves a trail in real-world economic data. In Section 2, we used administrative data to show that there is an almost linear increase in unemployment durations from age 35. This pattern is consistent with our experimental results. The callback-rate-age-profiles in the experiment are almost the inverses of the profiles for unemployment durations. Of course, labor supply could also matter, especially for workers above age 60, where retirement is a realistic option. The increase in unemployment durations by age is somewhat steeper for women than for men, which also is consistent with the steeper decline in the callback-rate-age profile for women in the experiment. Job mobility also declines by age in a way that is in line with the callback-rate-profiles in the experiment. Of course, the main reason why job mobility is lower among older workers is most likely because many of them are well-matched, but it may also reflect that some workers expect discrimination, and therefore do not search for a new job or are unable to find a new job.

Our results suggest that policymakers working with pension reforms face a twofold challenge. The negative demand effect is strongest for the oldest workers and, in addition, there is likely to be a negative supply effect among older workers. Therefore, if policymakers are to succeed with increasing employment among older workers, they must design measures that both increase labor supply and combat age discrimination.

References


- AARP (2000), American business and older workers: A summary of findings, AARP, Washington.
- Ahmed, A. M., Andersson, L., & Hammarstedt, M. (2012), Does age matter for employability? A field experiment on ageism in the Swedish labour market, *Applied Economics Letters*, 19, 403-406.
- Albert, R., Escot, L., & Fernandez-Cornejo, J. A. (2011), A field experiment to study sex and age discrimination in the Madrid labour market, *International Journal of Human Resource Management*, 22, 351-375.
- Arrow, K. J. (1973), The theory of discrimination, in Ashenfelter, O., Rees, A. (eds.), *Discrimination in labor markets*, Princeton University Press.
- Baert, S., Norga, J., Thuy, Y. & Van Hecke, M. (2016), Getting Grey Hairs in the Labour Market. A Realistic Experiment on Age Discrimination, *Journal of Economic Psychology*, 57, 86-101.
- Baert, S., De Pauw, A.-S. & Deschacht, N. (2016), Do employer preferences contribute to sticky floors?, *Industrial & Labor Relations Review*, 69, 714-736.
- Becker, G. (1957), *The economics of discrimination*, University of Chicago Press.
- Bertrand, M. & Mullainathan, S. (2004), Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination, *American Economic Review*, 94, 991-1013.
- Bendick Jr, M., Jackson, C. W., & Romero, J. H. (1996), Employment discrimination against older workers: An experimental study of hiring practices, *Journal of Aging & Social Policy*, 8, 25-46.
- Bendick Jr, M., Brown, L. E., & Wall, K. (1999), No foot in the door: An experimental study of employment discrimination against older workers, *Journal of Aging & Social Policy*, 10, 5-23.
- Blau, F. D., Ferber, M. A. & Winkler, A. W. (2014), *The Economics of Women, Men, and Work*, 7th ed. Upper Saddle River, NJ, Prentice Hall/Pearson.
- Blau, F. D. & Kahn, L. M. (2016), The gender wage gap: Extent, trends, and explanations, *Journal of Economic Literature* (forthcoming).
- Carlsson, M., & Rooth, D. O. (2007), Evidence of ethnic discrimination in the Swedish labor market using experimental data, *Labour Economics*, 14, 716-729.

- Carlsson, M. (2011), Does hiring discrimination cause gender segregation in the Swedish labor market?, *Feminist Economics*, 7, 71-102.
- Davison, H. K. & Burke, M.J. (2000), Sex discrimination in simulated employment contexts: A meta-analytic investigation, *Journal of Vocational Behavior*, 56, 225-248.
- Delgado Helleseter, M. and Kuhn, P. & Shen, K. (2016), Age and gender profiling in the Chinese and Mexican labor markets: Evidence from four job boards, NBER Working paper no. w22187.
- Eriksson, S. & Rooth, D.O. (2014), Do employers use unemployment as a sorting criterion when hiring? Evidence from a field experiment, *American Economic Review*, 104, 1014-1039.
- Farber, H. S. (1999), Mobility and stability: The dynamics of job change in labor markets, *Handbook of labor economics*, 3, 2439-2483.
- Grey, L. & McGregor, J. (2003), Human resource development and older workers: Stereotypes in New Zealand, *Asia Pacific Journal of Human Resources*, 41, 338-353.
- Henkens, K. (2005), Stereotyping older workers and retirement: The manager's point of view, *Canadian Journal on Aging*, 24, 353-66.
- Hewstone, M, Rubin, M. & Willis, H. (2002), Intergroup bias, *Annual Review of Psychology*, 53, 575-604.
- Johnson, R. W. & Neumark, D. (1996), Wage declines among old men, *Review of Economics and Statistics*, 78, 740-748.
- Kroft, K., Lange, F., & Notowidigdo, M. J. (2013), Duration dependence and labor market conditions: Theory and evidence from a field experiment, *Quarterly Journal of Economics*, 128, 1123-1167.
- Kuhn, P. & Shen, K. (2013), Gender discrimination in job ads: Evidence from China, *Quarterly Journal of Economics* 128, 287-336.
- Lahey, J. N. (2008), Age, women, and hiring an experimental study, *Journal of Human Resources*, 43, 30-56.
- Lazear, E. P. (1979), Why is there mandatory retirement?, *Journal of Political Economy*, 87, 1261-1284.
- Neumark, D., Bank, R., & Van Nort, K. (1996), Sex discrimination in restaurant hiring: An audit study, *Quarterly Journal of Economics*, 111, 915-941.
- Neumark, D., Burn, I. & Button, P. (2015), Is it harder for older workers to find jobs? New and improved evidence from a field experiment, NBER Working Paper No. 21669.
- Neumark, D. (2016), Experimental research on labor market discrimination, NBER Working

- Paper No. 22022.
- OECD (2006), *Live Longer, Work Longer*, OECD Publishing.
- OECD (2013), *Health at a Glance 2013: OECD Indicators*, OECD Publishing.
- OECD (2016), *Employment Outlook*, OECD Publishing.
- Petit, P. (2007), The effects of age and family constraints on gender hiring discrimination: A field experiment in the French financial sector, *Labour Economics*, 14, 371-391.
- Riach, P., & Rich, J. (1987), Testing for sexual discrimination in the labour market, *Australian Economic Papers*, 26, 165-178.
- Riach, P. & Rich J. (2002), Field experiments of discrimination in the market place, *Economic Journal*, 112, F480-F518.
- Riach, P. & Rich, J. (2006), An experimental investigation of sexual discrimination in hiring in the English labor market, *B.E. Journal of Economic Analysis and Policy*, 6(2).
- Riach, P. A., & Rich, J. (2010), An experimental investigation of age discrimination in the English labor market, *Annals of Economics and Statistics/Annales d'Économie et de Statistique*, 99/100, 169-185.
- Rich, J. (2014), What do field experiments of discrimination in markets tell us? A meta analysis of studies conducted since 2000, IZA Discussion Papers No. 8584, IZA.
- Rudman, L. A., & Goodwin, S. A. (2004), Gender differences in automatic in-group bias: Why do women like women more than men like men?, *Journal of Personality and Social Psychology*, 87, 494-510.
- Rupert, P. & Zanella, G. (2015), Revisiting wage, earnings, and hours profiles, *Journal of Monetary Economics*, 72, 114-130.
- Swedish Pensions Agency (2012), *Arbetsgivares attityder till äldre arbetskraft, Pensionsmyndigheten analyserar 2012:2*, Pensionsmyndigheten.
- Taylor, P. & Walker, A. (1998), Employers and older workers: Attitudes and employment practices, *Ageing & Society*, 18, 641-658.
- Theodossiou, I., & Zangelidis, A. (2009), Should I stay or should I go? The effect of gender, education and unemployment on labour market transitions, *Labour Economics*, 16, 566-577.
- Weichselbaumer, D. (2004), Is it sex or personality? The impact of sex-stereotypes on discrimination in applicant selection, *Eastern Economic Journal*, 30, 159-186.
- Willis, R. J. (1985), Wage determinants: A survey and reinterpretation of human capital earnings functions, in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, North Holland, Amsterdam.

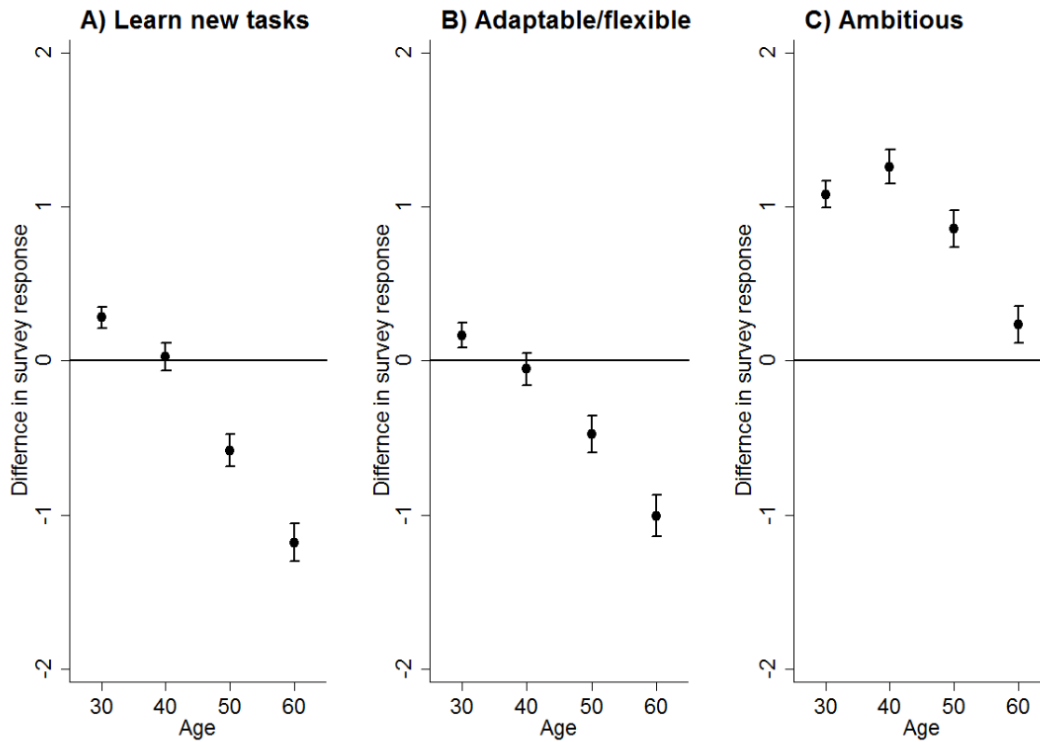
Figures and Tables

Figure 1. Callback rate by age



Note: The callback rate for each age is, on average, calculated based on $(6,066/2 \times 36) = 85$ job applications.

Figure 2. Difference in employers' perceptions about worker characteristics relative to a worker aged 20



Notes: The graphs are constructed from the survey answers to the following question: “Suppose that you are recruiting a new employee to a typical position in your workplace. To what extent, do you think that an average employee at age 20, 30, 40, 50, and 60 would have the following characteristics”. Employers’ perceptions are measured in scores on a scale from 0 (to a very small extent) to 10 (to a very high extent). The graphs show the average difference in employers’ perceptions compared to a worker aged 20 (the horizontal line at zero). The confidence intervals are at the 95 percent level.

Table 1. Descriptive statistics

Occupation	Number of resumes	Share of resumes	Callback rate
Administrative assistants	1,110	.18	.053
Chefs	1,059	.17	.158
Cleaners	789	.13	.091
Food serving and waitresses	1,137	.19	.048
Retail sales persons and cashiers	918	.15	.040
Sales representatives	558	.09	.091
Truck drivers	495	.08	.170
All	6,066	1	.087

Table 2. The probability of a callback

	(1)	(2)
Age	-.0048*** (.0004)	
Female	.0144* (.0076)	.0484*** (.0178)
Age x female		-.0057*** (.0005)
Age x male		-.0038*** (.0005)
Constant	.2651*** (.0283)	.2461*** (.0294)
p-value (test of equal age coeff.)		.0071

Notes: $N = 6,066$. All regressions include covariates for employment status, flexible, type of employment history, application type, application order, and fixed effects for occupations and cities. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.

Table 3. The probability of a callback and employment histories

No additional experience	.0149 (.0204)
Irrelevant experience	-.0202 (.0205)
Age x full relevant experience	-.0048*** (.0006)
Age x no additional experience	-.0055*** (.0006)
Age x irrelevant experience	-.0043*** (.0006)
p-value (test of equal age coeff.)	.351

Notes: $N = 6,066$. The regression include covariates for female, employment status, flexible, application type, application order, and fixed effects for occupations and cities. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.

Table 4. The probability of a callback, employment status and flexibility

	(1)	(2)	(3)
Panel A) Employment status			
Age	-.0048*** (.0004)		
Female	.0144* (.0076)	.0146* (.0076)	.0637*** (.0232)
Long-term unempl.	-.0182** (.0072)	-.0280* (.0169)	-.0091 (.0228)
Female x long-term unempl.			-.0360 (.0336)
Age x not long-term unempl.		-.0051*** (.0005)	
Age x long-term unempl.		-.0045*** (.0005)	
Age x not long-term unempl. x female			-.0063*** (.0007)
Age x long-term unempl. x female			-.0050*** (.0008)
Age x not long-term unempl. x male			-.0036*** (.0006)
Age x long-term unempl. x male			-.0040*** (.0007)
p-value (test of equal age coeff: all; female/male.)		.419	.174/.680
Panel B) Flexibility			
Age	-.0048*** (.0004)		
Female	.0144* (.0076)	.0145* (.0076)	.0482** (.0216)
Flexible	-.0014 (.0072)	.0119 (.0174)	.0128 (.0228)
Female x flexible			.0026 (.0351)
Age x not flexible		-.0046*** (.0004)	
Age x flexible		-.0053*** (.0006)	
Age x not flexible x female			-.0055*** (.0006)
Age x flexible x female			-.0061*** (.0009)
Age x not flexible x male			-.0034*** (.0006)
Age x flexible x male			-.0046*** (.0007)
p-value (test of equal age coeff: all; female/male)		.285	.588/.217

Notes: $N = 6,066$. All regressions include covariates for employment status, flexible, type of employment history, application type, application order, and fixed effects for occupations and cities. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.

Table 5. Change in callback rate resulting from a change in age category

	(1)	(2)
Age 44-52 relative age 35-43	-.0757*** (.0120)	
Age 53-61 relative age 44-52	-.0353*** (.0091)	
Age 62-70 relative age 53-61	-.0275*** (.0074)	
Female	.0132* (.0076)	.0449** (.0198)
Age 44-52 relative age 35-43 x female		-.0918*** (.0175)
Age 53-61 relative age 44-52 x female		-.0381*** (.0139)
Age 62-70 relative age 53-61 x female		-.0324*** (.0109)
Age 44-52 relative age 35-43 x male		-.0571*** (.0160)
Age 53-61 relative age 44-52 x male		-.0333*** (.0120)
Age 62-70 relative age 53-61 x male		-.0200** (.0097)
p-value (test of equal age coeff: all; female/male.)	.003	.0143/.1049

Notes: $N = 6,066$. The sample is divided into four age categories of a width of nine years. All regressions include covariates for employment status, flexible, type of employment history, application type, application order, and fixed effects for occupations and cities. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.

Table 6. The probability of a callback by occupation

	Adm. assistants	Chefs	Cleaners	Food serving	Retail sales	Sales represent- atives	Truck drivers	p-value (test of equal coeff.):
Regression A)								
Age x occupation	-.0029*** (.0006)	-.0078*** (.0011)	-.0046*** (.0009)	-.0037*** (.0007)	-.0022*** (.0007)	-.0059*** (.0012)	-.0083*** (.0016)	.000
Female x occupation	.0541*** (.0142)	-.0486** (.0246)	.0944*** (.0211)	.0190 (.0136)	.0113 (.0139)	-.0052 (.0226)	-.0443 (.0349)	.000
Regression B)								
Female x occupation	.1177*** (.0343)	-.0125 (.0500)	.1673*** (.0448)	.0283 (.0410)	.0445 (.0318)	-.0425 (.0609)	-.0483 (.0865)	
Age x female x occupation	-.0045*** (.0010)	-.0088*** (.0014)	-.0065*** (.0014)	-.0040*** (.0012)	-.0031*** (.0011)	-.0049*** (.0016)	-.0082*** (.0021)	
Age x male x occupation	-.0010 (.0008)	-.0066*** (.0017)	-.0022** (.0010)	-.0035*** (.0010)	-.0012 (.0009)	-.0070*** (.0018)	-.0084*** (.0026)	
p-value (equal age coeff.):	.009	.311	.014	.761	.142	.391	.948	

Notes: Both Panel A and B show results from a single regression with $N = 6,066$. The regressions do not include covariates for age and gender. Both regressions include covariates for employment status, flexible, type of employment history, application type, application order, and fixed effects for occupations and cities. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.


Table 7. The probability of a callback, share of females and gender of recruiter

	(1)	(2)	(3)	(4)
Panel A) Share of females				
Age	-0.0048*** (.0004)	-0.0050*** (.0004)		-0.0050*** (.0004)
Female	.0144* (.0076)	.0092 (.0088)	.0091 (.0088)	
Females 50-100%			.0125 (.0229)	-.0258** (.0131)
Age x females 50-100%			-.0051*** (.0006)	
Age x females 0-50%			-.0049*** (.0006)	
Female x females 50-100%				.0480*** (.0125)
Female x females 0-50%				-.0212* (.0122)
p-value (test of equal age coeff.):			.845	
p-value (test of equal female coeff.):				.000
Number of job applications	6,066	4,607	4,607	4,607
Panel B) Gender of recruiter				
Age	-0.0048*** (.0004)	-0.0049*** (.0005)		-0.0049*** (.0005)
Female	.0144* (.0076)	.0061 (.0112)	.0068 (.0113)	
Female recruiter			.0440 (.0301)	.0033 (.0168)
Age x female recruiter			-.0053*** (.0009)	
Age x male recruiter			-.0048*** (.0006)	
Female x female recruiter				.0474** (.0190)
Female x male recruiter				-.0155 (.0139)
p-value (test of equal age coeff.):			.659	
p-value (test of equal female coeff.):				.008
Number of job applications	6,066	3,204	3,204	3,204

Notes: Number of observations is smaller in model 2-4. All regressions include covariates for employment status, flexible, type of employment history, application type, application order, and fixed effects for occupations and cities. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.

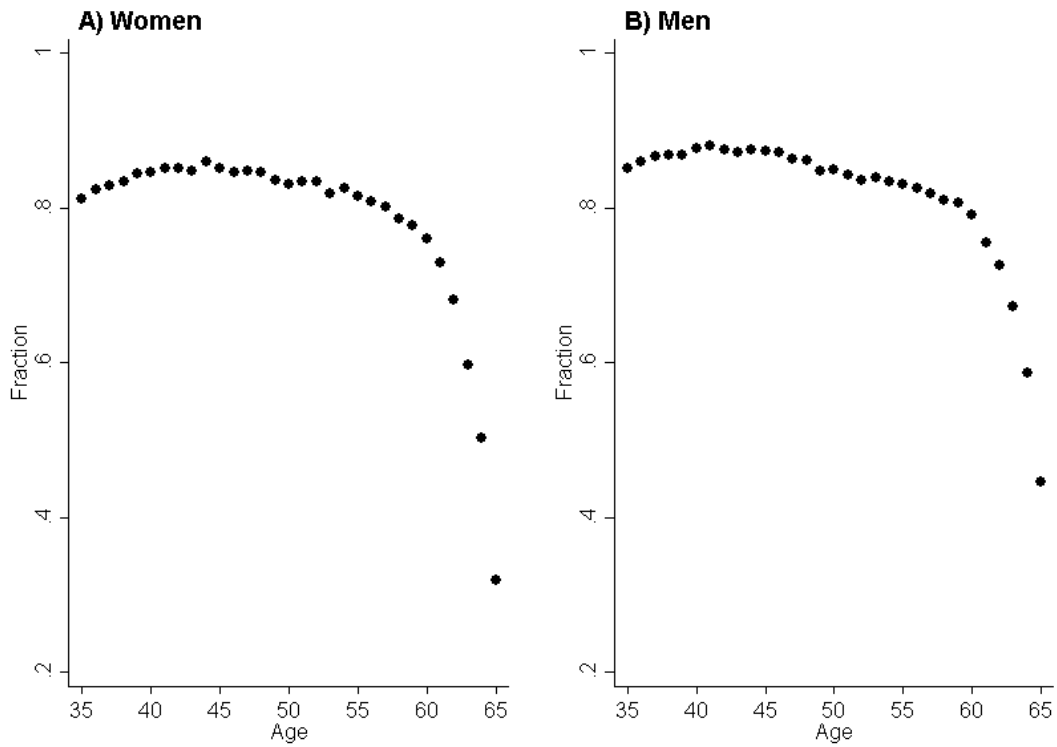
Appendix: Additional Figures and Tables

Appendix Figure A1. Yearly total labor market earnings by age




Notes: The graphs are constructed from the Longitudinal integration database for health insurance and labor market studies 2012 at Statistics Sweden which covers the whole Swedish population. We construct the graphs by first drawing a random sample of 20 percent of the individuals in the population and then excluding individuals outside the age interval 35-65 which leaves us with 755,938 individuals.

Appendix Figure A2. Employment rate by age



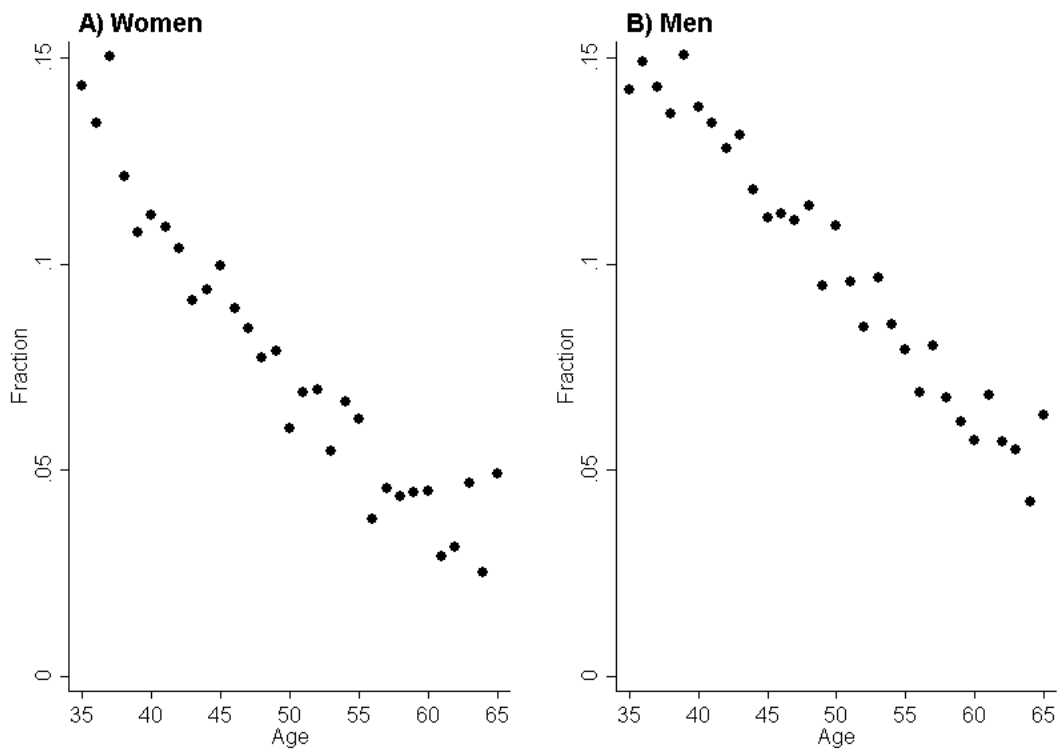
Notes: Details of the data used to construct these graphs are given in the notes below Figure A1. The employment rate is the fraction in the sample classified as employed according to Statistics Sweden's definition. Their method for measuring employment is similar to the method used by the International Labour Organization.

Appendix Figure A3. Unemployment duration by age




Notes: The graphs are constructed by calculating the average unemployment duration for each age. Data is obtained from the Swedish Public Employment Service, which contains information about all unemployed individuals registered at the Employment Service. The sample we use consists of individuals registered at the Employment Service from January 1, 2011, to February 18, 2012. Right censoring is not an important issue in these graphs, since all individuals are observed at least two years after the start of a period of unemployment because we observe individuals until February 18, 2014. The sample contains 46,201 unemployment periods and 44,258 individuals (a few persons have more than one period of unemployment).

Appendix Figure A4. Job mobility by age




Notes: The graphs are constructed from the Longitudinal integration database for health insurance and labor market studies and the register for Labour statistics based on administrative sources 2011-2012 at Statistics Sweden which covers the whole Swedish population. We construct the graphs by first drawing a random sample of 20 percent of the individuals in the population and then excluding individuals outside the age interval 35-65. We define a job change as a change in the workplace/firm where an individual is employed from 2011 to 2012. We are able to identify where an individual is employed, since we observe the workplaces/firms from which an individual has income. In the sample we use, there are 104,320 individuals. Then main reason why this sample is smaller than the samples used in Figure A1 and A2 is that here we conditional on being observed in both administrative registers and being employed both in 2011 and 2012.

Appendix Figure A5. Employment rate by age for United States



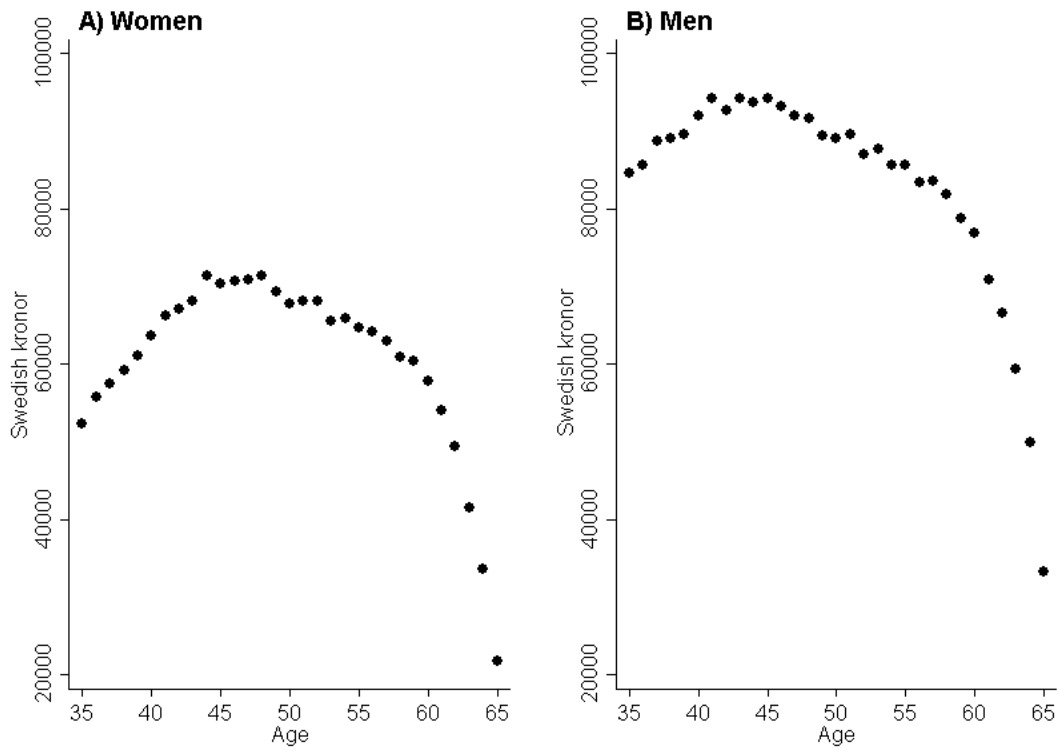
Notes: These graphs are constructed from aggregated labor force statistics from the Current Population Survey (2016, the average for quarter 1-4). Data can be downloaded here: <https://data.bls.gov/pdq/querytool.jsp?survey=ln>

Appendix Figure A6. Average unemployment duration by age for United States



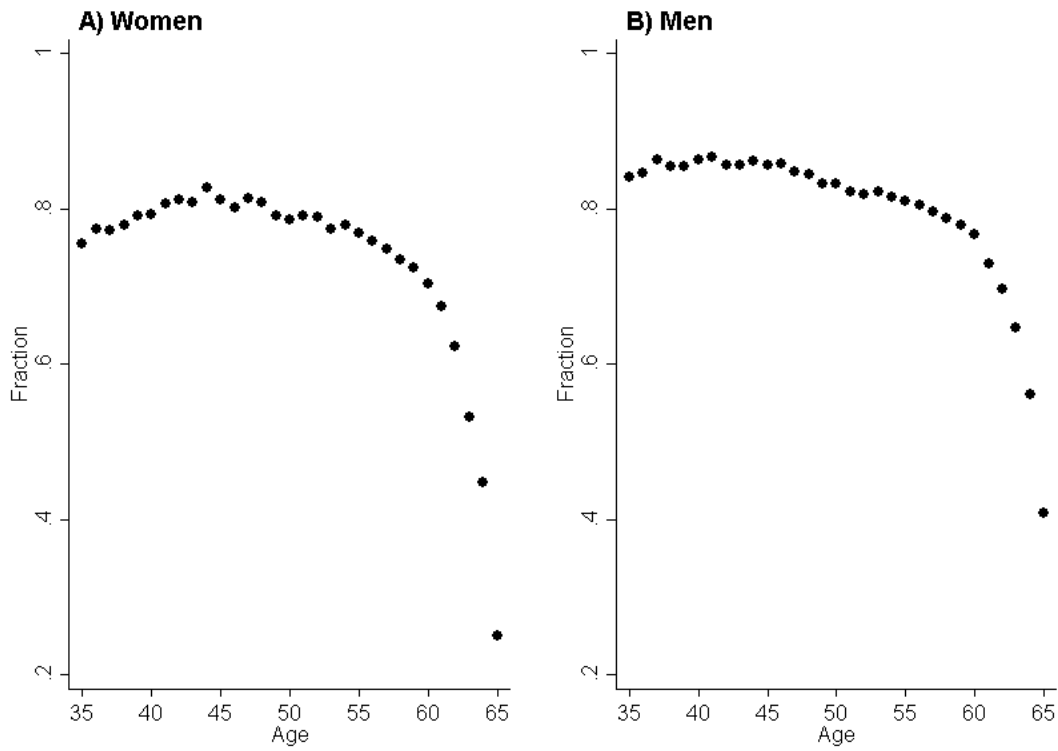
Notes: These graphs are constructed from aggregated labor force statistics from the Current Population Survey (February 2017). Data can be downloaded here: <https://www.bls.gov/web/empsit/cpseea36.htm>

Appendix Figure A7. Yearly total labor market earnings by age, low educated workers



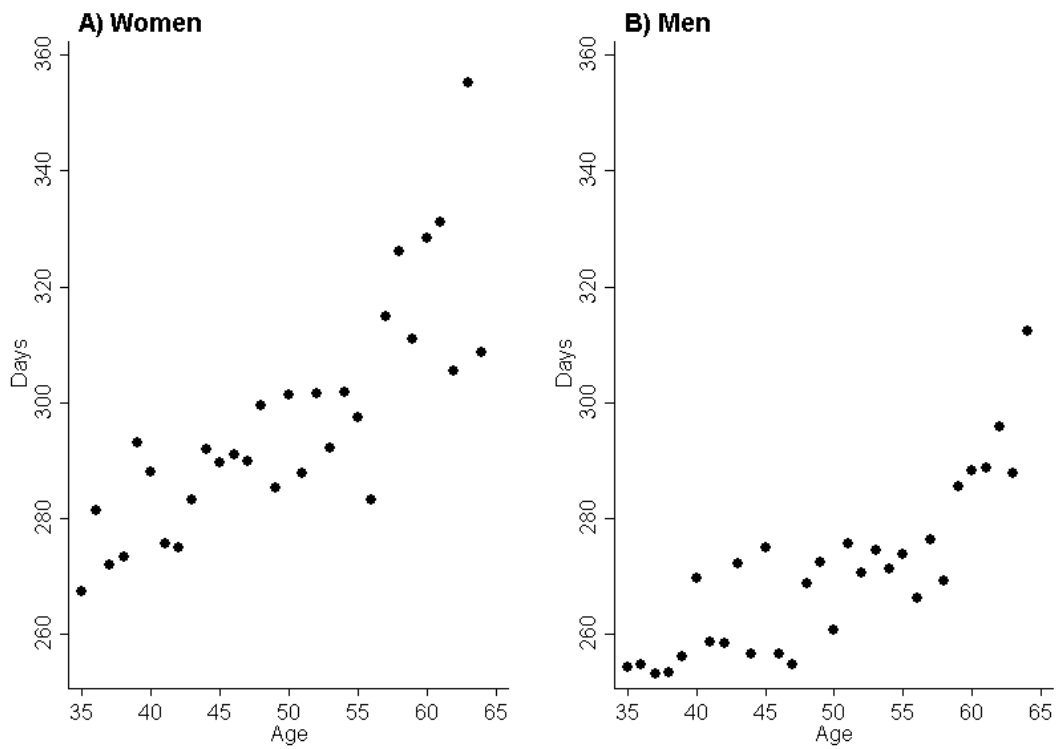
Notes: The sample used to construct these graphs contains 465,876 individuals. Low educated is defined as having a high school degree or less, i.e. no college education. Otherwise, see the notes below Figure A1.

Appendix Figure A8. Employment rate by age, low educated workers



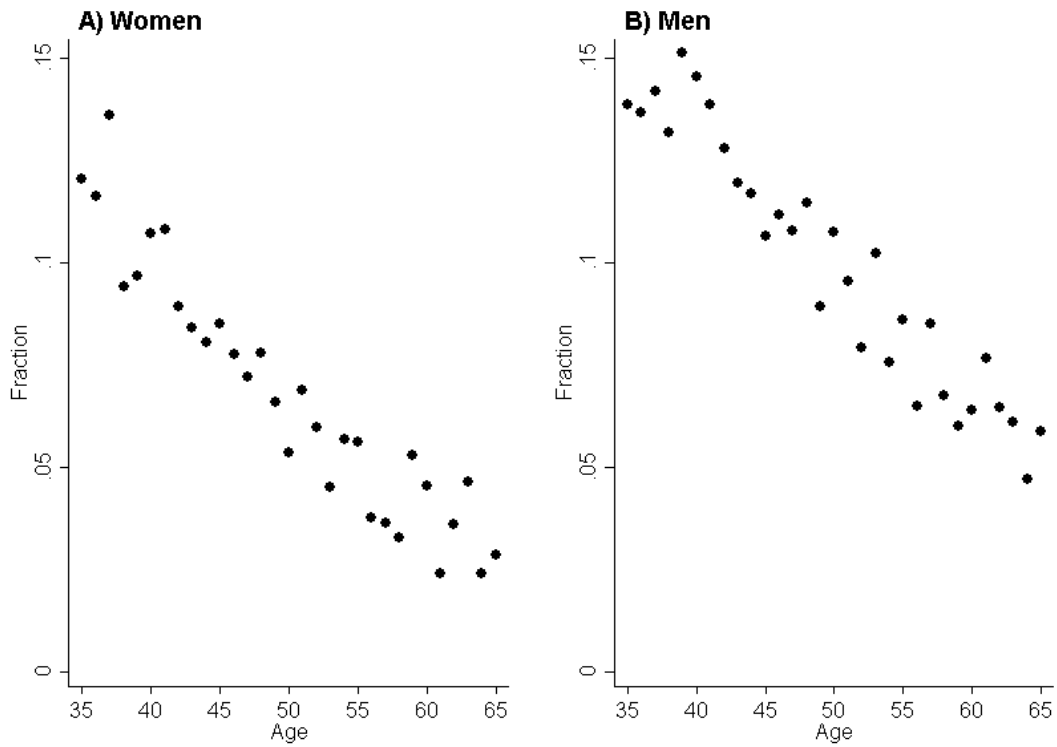
Notes: The sample used to construct these graphs contains 465,876 individuals. Low educated is defined as having a high school degree or less, i.e. no college education. Otherwise, see the notes below Figure A2.

Appendix Figure A9. Unemployment duration, low educated workers



Notes: The sample used to construct these graphs contains 30,677 individuals. Low educated is defined as having a high school degree or less, i.e. no college education. Otherwise, see the notes below Figure A3.

Appendix Figure A10. Mobility by age, low educated workers



Notes: The sample used to construct these graphs contains 61,026 individuals. Low educated is defined as having a high school degree or less, i.e. no college education. Otherwise, see the notes below Figure A4.

Appendix Figure A11. Example of cover letter

Hi,

The advertised job as an administrative assistant sounds very interesting and fits exactly what I would like to work with in the future.

To tell you something about myself, my name is Anna Eriksson and I am 38 years old. I have long experience of working as an administrative assistant. My last job was at Adecco, where my duties included accounting and financial statements, sales and purchase ledgers, and general office tasks at the various client companies.

As a person I am social, cooperative, careful, and used to a high work pace. Together with my partner, I spend my leisure time exercising and socializing with our many friends.

I would very much like to come and present myself more in an interview. Then I will bring my certificates and references from previous jobs.

Sincerely,
Anna Eriksson

Notes: In this example cover letter and CV (the CV is shown in Appendix Figure A6), the realized values of the random variables are as follows: (age = 38; female = yes; unemployment = 7 months; flexible and adaptable = no; employment history = 10 years of relevant work experience). 7 months of unemployment follows from the end date of the latest employment given by the CV, which is March 2015, and the fact that the job was applied for in October 2015 (not shown in the cover letter or CV). This resume and CV is for the administrative assistant occupation and was sent to a firm in Stockholm.

Appendix Figure A12. Example of CV

CV

Personal information:

Anna Eriksson
Kilsgatan 2, 123 44 Kista
0737-847914
anna.ia.eriksson@gmail.com
1976-12-03

Work experience:

March 2008 – March 2015	Adecco Sweden AB, accounting and financial statements, sales and purchase ledgers.
2005 – 2008	Telia Sonera AB, administrative duties
1994 – 2005	Several jobs at restaurants and cleaning companies

Education:

1992 – 1994	Distribution and office high school, Stockholm
1983 – 1992	Primary school, Stockholm

Language skills:

Swedish	Native speaker
English	Good skills

Computer skills:

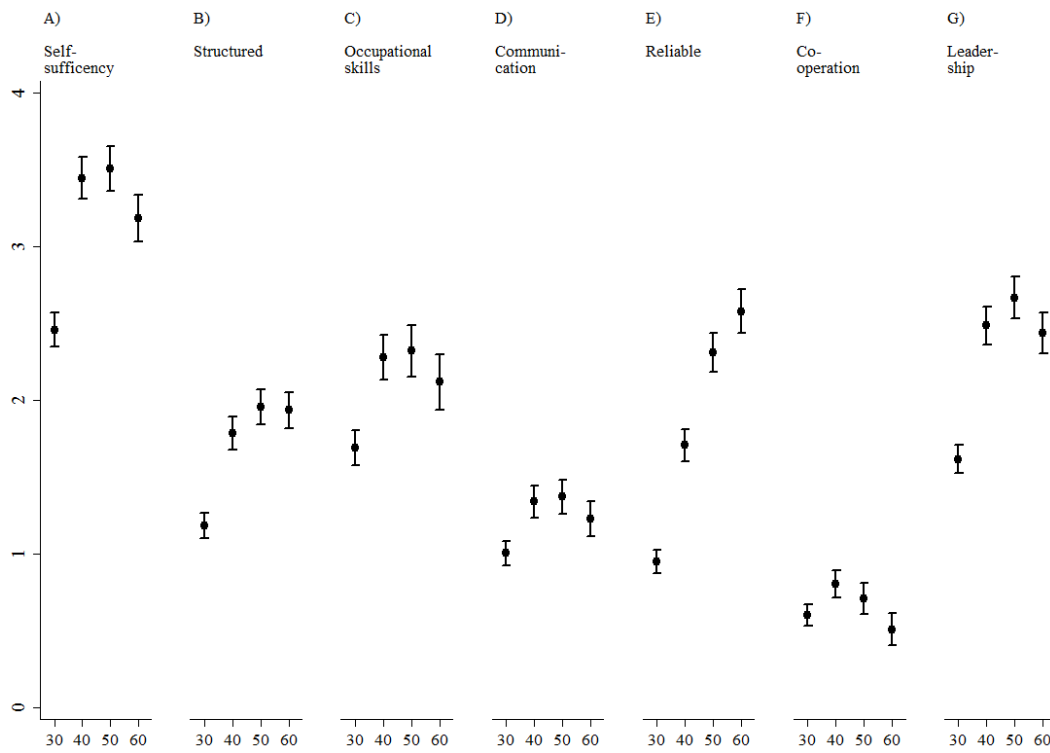
Word, excel, mail, Internet. Experience of wage and accounting programs.

Driver's license:

Private car

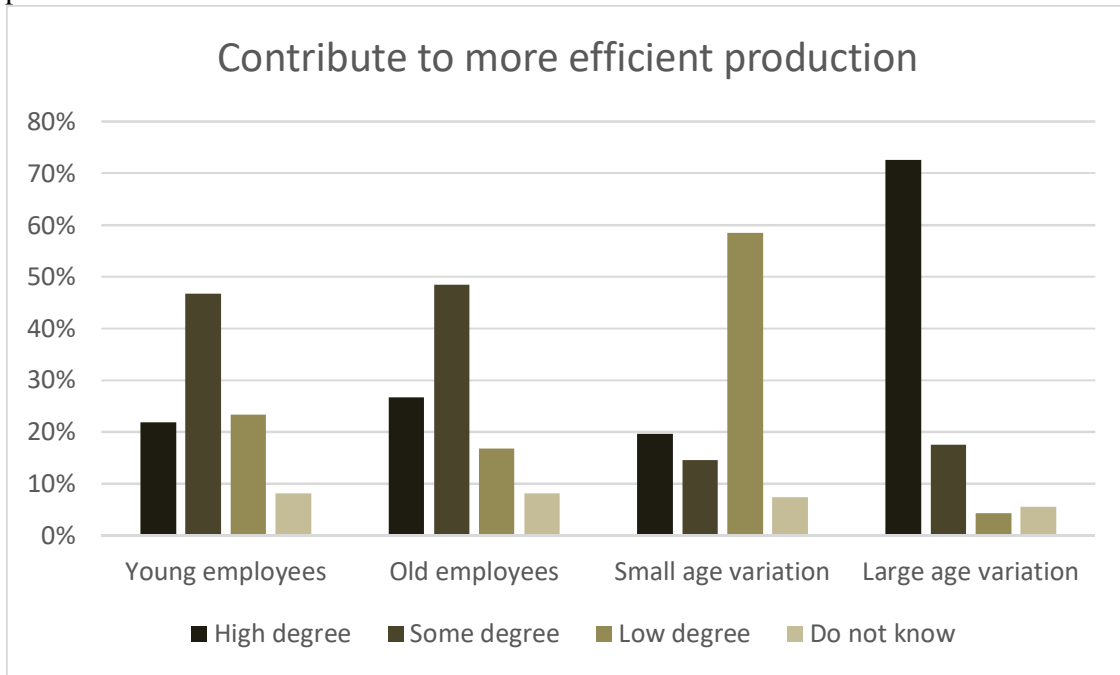
Notes: See the notes below Appendix Figure A11.

Appendix Figure A13. Employers' perceptions about worker characteristics at different ages



Notes: See the notes below Figure 2.

Appendix Figure A14. Employers' attitudes towards workers at different ages, efficient production



Notes: Based on the following question in the survey: "Efficiency is an important factor in most workplaces. To what extent do you agree with the following statements?"

(Scale: Fully/to a large extent/to some extent/to a low extent/not at all/do not know)

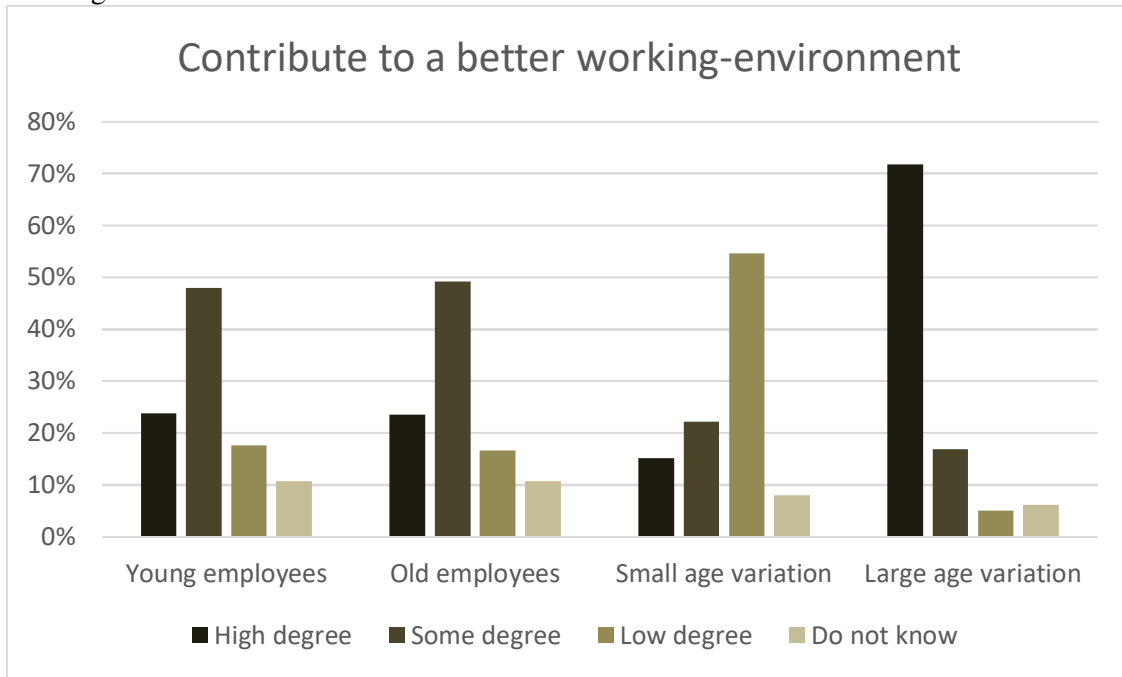
Younger employees contribute to more efficient production

Older employees contribute to more efficient production

Small variation in age among the employees contributes to more efficient production

Large variation in age among the employees contributes to more efficient production".

Appendix Figure A15. Employers' attitudes towards workers at different ages, better working environment



Notes: Based on the following question in the survey: "Another factor that is important in most workplaces is the working-environment (i.e. the atmosphere in the workplace).

To what extent do you agree with the following statements?

(Scale: Fully/to a large extent/to some extent/to a low extent/not at all/do not know)

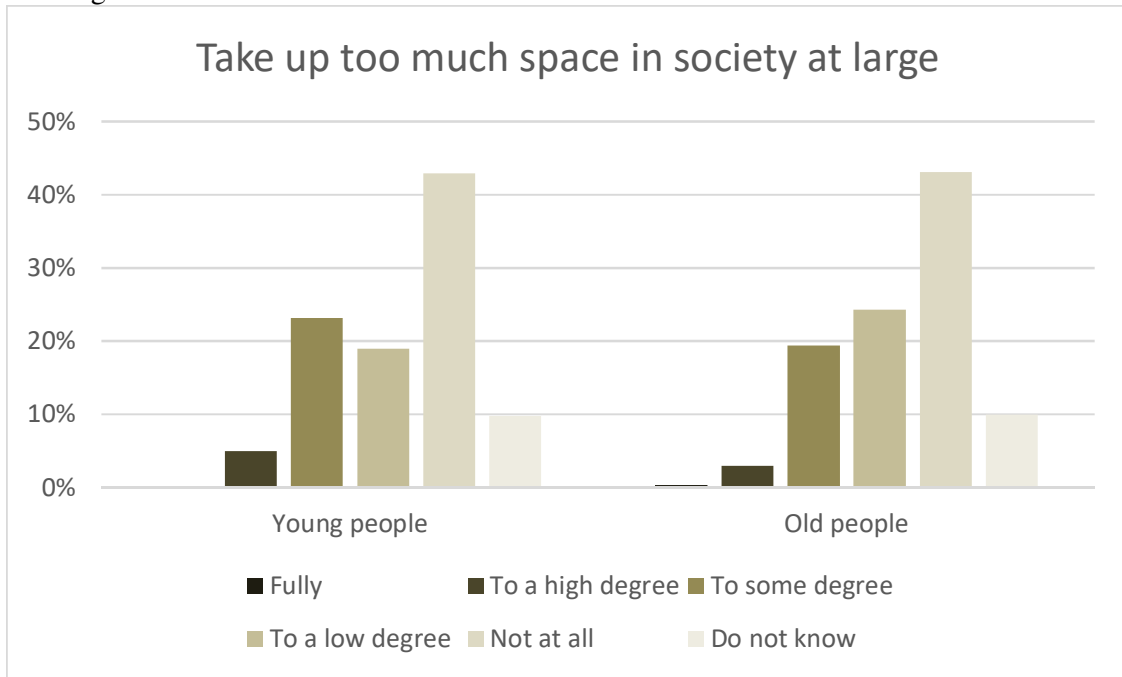
Younger employees contribute to a better working-environment

Older employees contribute to a better working-environment

Small variation in age among the employees contributes to a better working-environment

Large variation in age among the employees contributes a better working-environment".

Appendix Figure A16. Employers' attitudes towards workers at different ages, better working environment



Notes: Based on the following question in the survey: "To what extent do you agree with the following statements?"

(Scale: Fully/to a large extent/to some extent/to a low extent/not at all/do not know)

Younger people are taking up too much space in society at large

Older people are taking up too much space in society at large".

Appendix Table A1. Descriptive statistics of the occupations in the experiment.

Occupation	Number of workers	Share of all workers	Share men	Share women	Share age < 35	Share age 35-49	Share age 50-64	Share of all ES vac- ancies
Administrative assistants	141,672	.034	.17	.83	.26	.37	.37	.024
Chefs	37,220	.009	.46	.54	.37	.36	.27	.015
Cleaners	37,220	.016	.24	.76	.31	.37	.32	.015
Food serving and waitresses	84,291	.020	.32	.68	.60	.23	.17	.036
Retail sales persons, cashiers	173,663	.041	.37	.63	.59	.25	.16	.027
Sales representatives	91,176	.023	.71	.29	.26	.47	.28	.064
Truck drivers	55,216	.013	.94	.06	.35	.32	.33	.010

Notes: Data on employment in the occupations are from the Statistics Sweden's Occupational register, for year 2014, and workers aged 16-64. Administrative assistant are occupational codes 4111, 4112, 4119, and 4225. Chefs are occupational code 5120. Cleaners are occupational code 9111. Food serving and waitresses are occupational codes 5131, 9412, and 9413. Retail sales person and cashiers are occupational codes 5222, 5223, and 5230. Sales representatives are occupational code 3322. Truck drivers are occupational code 8332. Data on vacancies are from the Swedish Public Employment Service and measure the share of all reported vacancies for year 2015.

Appendix Table A2. Robustness.

	Baseline from Table 2 (1)	Excluding covariates (2)	Including firm fixed effects (3)	Probit (marginal effects) (4)	Interview as dependent variable (5)
Age	-.0048*** (.0004)	-.0049*** (.0004)	-.0047*** (.0005)	-.0044*** (.0003)	-.0025*** (.0003)
Female	.0144* (.0076)	.0170** (.0077)	.0212** (.0094)	.0135** (.0066)	.0033 (.0058)

Notes: $N = 6,066$. The regression in column 1 repeats the first column of Table 1. Column 2 exclude all other covariates than age and female. Column 3 include fixed effects for firm, i.e. 2,022 dummy indicators. Column 4 reports marginal effects from an estimated Probit regression. Column 5 uses an indicator of an explicit invitation to a job interview as the dependent variable instead of an indicator of a positive response. Standard errors are clustered by firm. *** significant at the 1% level, **significant at the 5% level; *significant at the 10% level.

Appendix Table A3. Correlation matrix.

	Age	Gender	Empl. status	Flexible	CV gap	Template	Order
Age	1						
Female	-.01	1					
Empl. status	-.02	.03	1				
Flexible	-.02	-.01	.01	1			
CV gap	.04	-.01	-.01	.06	1		
Template	.00	-.02	.01	-.01	.01	1	
Order	-.01	.00	-.02	-.03	-.00	.01	1